

Quantum®

# User's Guide

# StorNext 5



StorNext 5 User's Guide, 6-68042-01, May 2018, 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**

© 2018 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**

Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, FlexTier, GoVault, Lattus, NDX, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners.

Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners.

Quantum specifications are subject to change.



# Contents

---

Preface .....	xxi
<b>Chapter 1: Introduction .....</b>	<b>1</b>
StorNext Gateway Terminology .....	1
About the StorNext File System .....	2
About StorNext Storage Manager .....	3
About StorNext LAN Clients .....	3
StorNext Features .....	3
StorNext Limits .....	4
StorNext Use Cases and File System Restart Requirements .....	4
<b>Chapter 2: StorNext Graphical User Interface (GUI) Overview .....</b>	<b>5</b>
Accessing the StorNext GUI .....	5
StorNext Browser Support .....	5
Internet Explorer 9 Security Settings .....	6
The StorNext Home Page .....	8
StorNext Monitors .....	9
StorNext Home Page Dropdown Menus .....	11

<b>Chapter 3: The Configuration Wizard</b> .....	<b>20</b>
Configuration Wizard .....	20
High Availability Systems .....	21
Welcome .....	22
Licenses .....	22
System .....	24
Name Servers .....	30
Multiple fsnameservers Hosts and Redundant Metadata Networks .....	33
File Systems .....	34
Configuring Spotlight Proxy .....	35
Allocation Session Reservation .....	36
View a File System .....	37
Add a File System .....	38
Edit a File System .....	45
Delete a File System .....	51
Perform File System Actions .....	51
Storage Destinations Overview .....	53
Setting Up Libraries .....	54
Setting Up Storage Disks .....	57
Setting Up Object Storage Destinations .....	60
Change the Current State of Object Storage Destinations, Controllers, and I/O Paths .....	89
Setting Up Q-Cloud .....	90
Configure Data Replication .....	95
Configure Data Deduplication .....	97
Storage Policies .....	98
Adding a Storage Manager Policy .....	99
Viewing, Running, Editing, Deleting or Testing a Storage Policy .....	104
Adding a Replication or Deduplication Policy .....	105
Client-side Encryption .....	112

Linear Tape File System (LTFS) Media Format .....	118
Enhanced Control of Tape Drive Allocation .....	121
Email Server .....	124
Email Notifications .....	125
Done .....	127

## Chapter 4: File System Tasks ..... 129

Label Disks .....	130
Label a Device .....	130
Unlabeling a Device .....	132
Understanding Resource Allocation .....	132
About File System Expansion .....	132
About Stripe Group Movement .....	133
Expansion and Movement Steps .....	133
Using Resource Allocation From the Command Line .....	133
Checking the File System .....	134
Adding a Stripe Group Without Moving .....	134
Adding and Moving a Data Stripe Group .....	134
Moving a Metadata/Journal Stripe Group .....	135
Check File System .....	137
Viewing and Deleting a Check Report .....	138
File System Check Output Files .....	139
Affinities .....	139
Auto Affinities .....	140
Affinity Preference .....	140
Allocation Strategy .....	141
Example Use Cases .....	141
Adding a New Affinity .....	142
Deleting an Affinity .....	143

Migrate Data .....	144
Migrate Metadata and Journal Data .....	145
Migrate User Data .....	146
Truncation Parameters .....	147
Manage Quotas .....	148
Quota Limits .....	152
Quota Types .....	152
Rename a Standalone (unmanaged) StorNext File System .....	153
<b>Chapter 5: Storage Manager Tasks .....</b>	<b>155</b>
Storage Components .....	157
Drive Pools .....	159
Media Actions .....	161
Storage Exclusions .....	170
Truncation Exclusions .....	173
Tape Consolidation .....	175
Library Operator Interface .....	177
Software Requests .....	179
Scheduler .....	180
Alternate Store and Retrieval Location .....	183
Architecture .....	185
Components of the Alternate Store Location Feature .....	186
User Customizable Transfer Script .....	188
Commands .....	189
Log Files .....	190
Installation, Configuration and Upgrade .....	190
Operational Modes .....	191
Monitoring .....	191
Copying Pre-Existing Files .....	191

Disaster Recovery .....	192
Troubleshooting .....	192
Examples .....	193
Distributed Data Mover .....	195
Drive Replacement .....	203
Active Vault Policy .....	204
System Parameters .....	211
<b>Chapter 6: Replication and Deduplication .....</b>	<b>213</b>
Replication Overview .....	214
Replication Configuration Overview .....	214
Replication Process Overview .....	215
Files Excluded From Replication .....	216
Replication Terms and Concepts .....	217
Namespace Realization .....	217
Blockpool .....	217
Blackout Period .....	217
Replication Source Policy and Replication Source Directory .....	218
Replication Target Directory .....	218
Replication Schedule .....	218
Replication Copies .....	218
Bandwidth Throttling .....	218
Multilink .....	219
Virtual IP (VIP) .....	219
Some Replication Scenarios .....	219
Scenario 1: Simplest Replication .....	220
Scenario 2: Replicating Multiple Copies in the Same Target File System .....	220
Scenario 3: Replicating to Multiple Target Hosts / File Systems .....	222
Additional Replication Possibilities .....	222

Non-Supported Replication Between Source and Target .....	225
“Chained” Replication .....	225
Configuring Replication .....	226
Step 1: Create Source and Target File Systems .....	226
Step 2: Setting up the Blockpool .....	229
Step 3: Creating Replication Targets .....	230
Step 4: Create a Replication Storage Policy .....	232
Configuration Steps Summary .....	232
Scheduling Replication Blackouts (Optional) .....	233
Optional HA and Multilink Configuration .....	233
Running Replication Manually (Optional) .....	236
Replication Statuses and Reporting .....	237
Replication Reports .....	237
Replication Administration .....	237
StorNext Jobs .....	238
Replication Target Relocating Procedures .....	239
Replication Target Relocating Basics .....	239
Target Relocating Procedures .....	241
Perform Cross-Mount Initial Replication .....	243
Relocating the Target File System to the Target Site .....	244
Troubleshooting Replication .....	246
Data Deduplication Overview .....	247
How Deduplication Works .....	247
Deduplication and Replication .....	248
Setting Up Deduplication .....	248
Data Deduplication Functions .....	250
Replication / Deduplication Removal Procedures .....	251
Assumptions .....	251
Removal Procedures .....	252



Collect and Understand Replication/Deduplication Configurations ..... 252

Obtain Information from StorNext GUI ..... 252

Obtain Information from the Command Line ..... 253

Replication Removal on a Target Host ..... 255

Replication Removal on a Source Host ..... 261

**Chapter 7: Tools Menu Functions ..... 265**

Tools Menu Overview ..... 266

User Accounts ..... 268

Client Download ..... 272

    Downloading Client Software ..... 272

Firmware Upgrade ..... 285

Appliance Release Notes ..... 285

Dynamic Application Environment ..... 287

System Control ..... 290

Lattus Certificates ..... 291

File and Directory Actions ..... 303

S3 Buckets ..... 312

StorNext Connect Ports ..... 314

File Systems Overview ..... 315

    Label Disks ..... 315

    Check File System ..... 317

    Affinities ..... 318

    Migrate Data ..... 319

    Truncation Parameters ..... 321

    Manage Quotas ..... 322

Storage Manager ..... 324

Replication and Deduplication ..... 324

High Availability (HA) ..... 324

Upgrade Firmware .....	324
Appliance Release Notes .....	325
<b>Chapter 8: Service Menu Functions .....</b>	<b>326</b>
Service Menu Overview .....	326
Health Check .....	327
Capture State and Capture DSET .....	329
Capture State .....	332
System Backup .....	337
Admin Alerts .....	337
Tickets .....	339
Logging .....	342
Web Services (V2) .....	342
<b>Chapter 9: Converting to HA .....</b>	<b>344</b>
HA Overview .....	344
HA Terms and Concepts .....	345
Failover .....	345
Primary Node .....	345
Secondary Node .....	345
Virtual IP (vIP) .....	346
Virtual Netmask .....	346
HA Reset .....	346
Preparing for HA Conversion .....	346
Pre-Conversion Steps .....	347
HA and LAN Clients .....	347
StorNext LAN Clients in HA Environments .....	348
Converting to HA .....	348
HA Conversion Procedure .....	349

Manage .....	353
HA Statuses and Reporting .....	354
Troubleshooting HA .....	355
<b>Chapter 10: StorNext Reports .....</b>	<b>361</b>
Reports Menu Overview .....	362
Logs .....	364
Jobs .....	365
Files .....	369
Drives .....	369
Media .....	370
Q-Cloud Object Storage Media Usage .....	372
Relations .....	373
File Systems .....	374
SAN Devices .....	374
Tape Consolidation .....	375
SAN and LAN Clients .....	376
LAN Client Performance .....	377
Replication / Deduplication Reports .....	378
Policy Activity .....	378
Policy Summary .....	379
Progress Report .....	381
Distributed Data Mover .....	382
Hardware Status .....	383
Gateway Metrics .....	385
StorNext Metrics .....	389
<b>Chapter 11: Lattus Object Storage .....</b>	<b>401</b>
Audience .....	402

Overview .....	402
Object Storage Media versus Storage Manager Media .....	402
Converting an AXR Namespace to an S3 Bucket .....	403
Object Storage Features in the StorNext GUI .....	404
Configuring Object Storage .....	410
The fsobjcfg Command .....	411
How to Route Backup Files to Lattus for Easier Recovery .....	413
The MAX_STORE_SIZE System Parameter .....	413
Setting Up Object Storage Destinations .....	414
HTTPS Support for Object Storage .....	421
The FS_OBJSTORAGE_CACERT System Parameter .....	422
The FS_OBJSTORAGE_CAPATH System Parameter .....	422
The FS_OBJSTORAGE_SSL_VERIFY_PEERHOST System Parameter .....	422
Configure HTTPS on DDM Hosts .....	422
HTTPS Support for Q-Cloud .....	424
Changes to Existing CLI Commands .....	424
The fsfileinfo Command .....	424
The fsmedinfo Command .....	424
The fsmedread Command .....	425
The fsmedwrite Command .....	425
The fsmedscan Command .....	425
The fsobjcfg Command .....	425
The fsaddclass Command .....	425
The dm_info Command .....	425
The dm_util Command .....	426
Other Changes and Considerations .....	426
Object Storage Segment Size .....	426
snpolicyd .....	427
Storage Manager .....	427

<b>Chapter 12: Object Storage</b> .....	<b>429</b>
Audience .....	430
Overview .....	430
Object Storage Media versus Other Storage Manager Media .....	431
Object Storage Features in the StorNext GUI .....	431
System Parameters .....	437
Configuring Object Storage .....	438
The fsobjcfg Command .....	439
How to Route Backup Files to Object Storage for Easier Recovery .....	441
How to Route File Copies to a Specific Object Storage Namespace .....	441
The MAX_STORE_SIZE System Parameter .....	442
Setting Up Object Storage Destinations .....	443
Setting up Lattus Object Storage Destinations on a StorNext Configuration .....	443
Setting up S3COMPAT Object Storage Destinations on a StorNext Configuration .....	450
Setting up AWS Object Storage Destinations on a StorNext configuration .....	455
Setting up Azure Object Storage Destinations on a StorNext Configuration .....	468
Change the Current State of Object Storage Destinations, Controllers, and I/O Paths .....	472
HTTPS Support for Object Storage .....	473
The FS_OBJSTORAGE_CACERT System Parameter .....	473
The FS_OBJSTORAGE_CAPATH System Parameter .....	473
The FS_OBJSTORAGE_SSL_VERIFY_PEERHOST System Parameter .....	474
Configure HTTPS on DDM Hosts .....	474
HTTPS Support for Q-Cloud .....	475
Changes to Existing CLI Commands .....	475
The fsfileinfo Command .....	476
The fsmedinfo Command .....	476
The fsmedread Command .....	476
The fsmedwrite Command .....	476
The fsmedscan Command .....	476

The fsobjcfg Command .....	477
The fsaddclass and fsmodclass Commands .....	477
The dm_info Command .....	477
The dm_util Command .....	477
Other Changes and Considerations .....	477
Object Storage Segment Size .....	478
<b>Chapter 13: Q Cloud .....</b>	<b>479</b>
Overview of Q-Cloud .....	479
Ease of Use .....	479
Flexibility .....	480
Reliability and Security .....	480
Cost Effectiveness .....	480
Setting Up Q-Cloud .....	481
Q-Cloud Tuning Considerations .....	486
General Considerations .....	487
TCP/IP Tuning .....	487
Network Interface (NIC) Bonding .....	487
Number of Concurrent IO Streams .....	487
Quality of Service Configuration .....	488
Q-Cloud Time Stamp Requirement .....	488
Configure Storage Manager to Perform Continuous Retries .....	488
Compression or Encryption Considerations .....	488
<b>Appendix A: Operating Guidelines .....</b>	<b>490</b>
Gateway Server/Client Network and Memory Tuning .....	491
Gateway Server and Client Network Tuning .....	491
Gateway Server Memory Tuning .....	492
Configuring LDAP .....	492

Using LDAP .....	492
UNIX File and Directory Modes .....	493
LDAP Refresh Timeout .....	493
Setting Up Restrictive ACLs .....	493
Default Single-Path I/O Retry Behavior .....	493
Event Handles for fsm.exe on a Windows Metadata Server .....	494
FSBlockSize, Metadata Disk Size, and JournalSize Settings .....	494
Metadata Disk Size Setting .....	494
JournalSize Setting .....	495
Disk Naming Requirements .....	495
Changing StorNext's Default Session Timeout Interval .....	496
Configuring a Data Partition for Use with Spectra Logic T-series Tape Storage Libraries .....	497
Basic Secure Sockets Layer (SSL) Guidelines .....	497
Example of a server.pem File .....	499
Example of a public.pem File .....	500
Name Limitations .....	500
Ports Used By StorNext .....	501
Log Rolling and Disk Space Health Check .....	503
Log Rolling .....	503
Disk Space Health Check .....	505
General Operating Guidelines and Limitations .....	505

## Appendix B: Additional Replication and Deduplication Information ...521

Replication Configuration File .....	522
Replication Terminology and Conventions .....	522
Copies in Replication Versus Copies and Versions in Storage Manager .....	522
Context 1 .....	522
Context 2 .....	523
Context 3 .....	523

Context 4 .....	523
Context 5 .....	523
Context 6 .....	524
Context 7 .....	524
Replication Target Directories .....	524
Number of Replication Copies .....	524
Isolating a Replication Target Directory .....	526
Final Recommendation For Target Directories .....	527
StorNext snpolicyd Policies .....	527
Example .....	528
Replication Copies = 2 (Detail) .....	530
Digression .....	530
Second Replication .....	531
More About Replication Target Directories .....	532
Examples .....	533
Deduplication Overview .....	534
Enabling Deduplication .....	535
Deduplication Modification Time .....	536
Deduplication and Blockpools .....	536
Deduplication and Truncation .....	536
Enabling Deduplication and Truncation .....	537
Storage Manager Truncation .....	537
Replication, Deduplication and Truncation .....	537
Replication, Deduplication and Storage Manager .....	538
Replicating into a Storage Manager Relation Point .....	539
Truncation and Deduplication / Replication (with and without SM) .....	540
The snpolicyd Debug Log .....	546



<b>Appendix C: High Availability Systems .....</b>	<b>548</b>
High Availability Overview .....	549
HA Internals: HAmom Timers and the ARB Protocol .....	550
FSM Election, Usurpation and Activation .....	550
LAN Connectivity Interruptions .....	551
Autonomous Monitoring and HA Resets .....	551
Setting the Timer Value .....	552
Negotiated Timer Resets .....	552
Primary and Secondary Server Status .....	553
File System Types .....	553
The ha_peer and fsnameservers File .....	554
HA Manager .....	555
Configuration and Conversion to HA .....	559
Conversion to HA .....	560
SyncHA process .....	561
Managing High Availability in the StorNext GUI .....	562
Configuring Multiple NICs .....	563
LAN Configuration .....	563
High Availability Operation .....	564
Windows and Linux SNFS Installations Without the HaShared File System .....	565
Linux SNMS and SNFS Installations with the HaShared File System .....	565
HA Resets .....	570
HA Resets of the First Kind .....	570
HA Resets of the Second Kind .....	571
HA Resets of the Third Kind .....	571
Using HA Manager Modes .....	571
HA Tracing and Log Files .....	572
Single (Singleton) Mode .....	572
Replacing a HA System .....	573

FSM Failover In HA Environments .....	573
Failover Timing .....	573
Moving a HA Shared File System to a New Raid .....	576
Changing the IP Address of the MDC in an HA Pair .....	577
Installing StorNext Licenses For HA Configurations From the CLI .....	577
Installing StorNext Licenses From the CLI for a Running HA Configuration .....	577
Installing StorNext Licenses From the CLI for a Stopped HA Configuration .....	578
<b>Appendix D: Web Services API .....</b>	<b>579</b>
<b>Appendix E: Storage Manager Truncation .....</b>	<b>580</b>
Truncation Overview .....	580
Normal Truncation .....	581
Immediate Truncation .....	581
Daily Truncation .....	582
Space Management .....	582
LOSPACE Cycle .....	582
Emergency Cycle .....	583
Disabling Truncation .....	584
Truncation Feature Locking .....	584
Disable Truncation Commands .....	585
Common Problems .....	585
Files Are Not Truncated as Expected .....	585
"Old" Files Not Truncating According to Policy Class .....	585
Small Files Not Truncating .....	586
Miscellaneous Usage Notes .....	586
Scheduling Truncation Manually .....	586
Setting Truncation .....	587
Setting up the CRON Job .....	587

Considerations When Scheduling Truncation and Relocation Policies .....589

**Appendix F: Security .....591**

StorNext Security (for pre-StorNext 5 release 5.4.0 Systems) ..... 591

    ACLs on Windows (for pre-StorNext 5 release 5.4.0 Systems) .....592

    Enable ACLs on a StorNext MDC (for pre-StorNext 5 release 5.4.0 Systems) .....595

StorNext Security ..... 597

    Permission Enforcement Details ..... 599

    Display and Modification of File Permissions ..... 603

    ACL Inheritance ..... 606

    Identity Mapping ..... 607

    Configuration Examples ..... 609

Central Control ..... 612

    Controls ..... 612

Limitations ..... 613

Example ..... 614

Cross-Platform Permissions ..... 615

Config (.cfg) File Options ..... 616

**Appendix G: Troubleshooting .....619**

Troubleshooting StorNext File System ..... 619

Troubleshooting OS Issues ..... 626

Troubleshooting Replication ..... 628

Troubleshooting HA ..... 630

Troubleshooting StorNext Installation and Upgrade Issues ..... 635

Troubleshooting Other Issues ..... 636

Debugging StorNext for Object Storage Systems and Cloud Providers ..... 638

    Debugging Steps to Determine the Source of the Error ..... 638

<b>Appendix H: StorNext Offline Notification</b> .....	<b>641</b>
StorNext Offline Notification Overview .....	641
What is a StorNext Offline File? .....	642
Why Block Access to Offline Files? .....	642
Offline Notification Configuration Options .....	642
How the Notification Feature Operates .....	642
Installing the Notification Application .....	643
Installing onto a Standalone Machine .....	644
Starting the Notification Application .....	650
Configuring the Notification Application .....	650
Setting Application Options .....	651
Viewing Access Denied Files .....	652
Viewing the Application Log .....	653
Exiting Application Preferences .....	654
Configure StorNext Offline Notification with Anti-virus or Anti-malware Software .....	654
Uninstalling the Notification Application .....	654
 <b>Appendix I: RAS Messages</b> .....	 <b>656</b>
 <b>Appendix J: Repairing and Replacing StorNext Metadata Servers</b> ...	 <b>657</b>
Replacing an MDC in a non-HA environment(non-backup/restore method) .....	658
Replacing an MDC in non-HA environment(backup/restore method) .....	660
Replacing an MDC in an HA environment .....	662
Pre-Conversion Steps .....	662
Conversion Steps .....	663
Post-Conversion Steps .....	664



# Preface

This manual contains the following chapters:

- [Chapter 1: Introduction](#)
- [Chapter 2: StorNext Graphical User Interface \(GUI\) 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: Lattus Object Storage](#)
- [Chapter 12: Object Storage](#)
- [Chapter 13: Q Cloud](#)
- [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](#)
- [Appendix 1: MySQL File-per-table Conversion](#)

This manual is written for StorNext 5 (**FX** stands for **F**ile System for **X**san) operators, system administrators, and field service engineers.

## Notational Conventions


This manual uses the following conventions:

Convention	Example
User input is shown in bold monospace font.	<b>./DARTinstall</b>
Computer output and command line examples are shown in monospace font.	./DARTinstall
User input variables are enclosed in angle brackets.	<b>http://&lt;ip_address&gt;/cgi-bin/stats</b>
For UNIX and Linux commands, the command prompt is implied.	./DARTinstall is the same as # ./DARTinstall
File and directory names, menu commands, button names, and window names are shown in bold font.	<b>/data/upload</b>
Menu names separated by arrows indicate a sequence of menus to be navigated.	<b>Utilities &gt; Firmware</b>

The following formats indicate important information:

 **Note:** Note emphasizes important information related to the main topic.

 **Caution:** Caution indicates potential hazards to equipment or data.

 **WARNING:** Warning indicates potential hazards to personal safety.

- Right side of the system - Refers to the right side as you face the component being described.
- Left side of the system - Refers to the left side as you face the component being described.
- Data sizes are reported in base 10 (decimal) rather than base 2<sup>10</sup> (binary). For example:

10,995, 116,277,769 Bytes are reported as 11.0 TB (decimal/1000). In binary, this value is 10 TiB (binary/1024).

## Product Safety Statements

Quantum will not be held liable for damage arising from unauthorized use of the product. The user assumes all risk in this aspect.

This unit is engineered and manufactured to meet all safety and regulatory requirements. Be aware that improper use may result in bodily injury, damage to the equipment, or interference with other equipment.

- 
- ⚡ WARNING:** Before operating this product, read all instructions and warnings in this document and in the *Quantum Products System, Safety, and Regulatory Information Guide*.
- 
- ⚡ ADVARSEL:** Læs alle instruktioner og advarsler i dette dokument og i *Informationsvejledning vedrørende system-, sikkerheds- og lovbestemmelser for Quantum produkter, før produktet betjenes*.
- 
- ⚡ AVERTISSEMENT :** Avant d'utiliser ce produit, lisez toutes les instructions et les avertissements de ce document et du *Guide d'informations sur le système, la sécurité et la réglementation de Quantum*.
- 
- ⚡ WARNUNG:** Lesen Sie vor der Inbetriebnahme dieses Produkts alle Anleitungen und Warnungen in diesem Dokument und im *System-, Sicherheits- und Betriebsbestimmungen-Handbuch für Quantum-Produkte*.
- 
- ⚡ ADVERTENCIA:** Antes de hacer funcionar este producto, lea todas las instrucciones y advertencias de este documento y de la *Guía de información normativa, del sistema y de seguridad de los productos de Quantum*.
- 
- ⚡ VARNING:** Läs igenom alla instruktioner och varningar i detta dokument och i *Quantums produktsystem, säkerhet och reglerande informationsguide* innan denna produkt används.
- 
- ⚡ ВНИМАНИЕ!** Перед началом эксплуатации данного изделия прочтите все инструкции и предупреждения, приведенные в настоящем документе и в *Руководстве по системе, технике безопасности и действующим нормативам компании Quantum*.
- 
- ⚡ 警告：**本製品を使用される前に、本書と『*Quantum製品システム、安全、規制情報ガイド*』に記載されているすべての説明と警告をお読みください。
- 
- ⚡ 경고:** 본 제품을 작동하기 전에 본 문서와 *Quantum 제품 시스템, 안전 및 규제 정보 설명서*에 있는 모든 지침과 경고를 참조합니다.
- 
- ⚡ 警告：**在操作本产品之前，请阅读本文档和 *Quantum 产品系统、安全和法规信息指南*中的所有说明和警告。
- 
- ⚡ 警告：**操作此產品前，請閱讀本檔案及 *Quantum 產品系統、安全與法規資訊指南*中的指示與和警告說明。

---

**אזהרה:** לפני ההפעלה של מוצר זה, קרא את כל ההוראות והאזהרות הכלולות במסמך זה וכן במדריך המידע בנושא מערכת, בטיחות ותקינה עבור מוצרי Quantum.



For the most up to date information on StorNext 5, see:

<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

## Contacts

For information about contacting Quantum, including Quantum office locations, go to:

<http://www.quantum.com/aboutus/contactus/index.aspx>

## Comments

To provide comments or feedback about this document, or about other Quantum technical publications, send e-mail to:

[doc-comments@quantum.com](mailto: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:

- **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. Get started at:  
<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>
- **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. Get started at:

<http://www.quantum.com/customercenter/>

For further assistance, or for training opportunities, contact the Quantum Customer Support Center:



Region	Support Contact
North America	1-800-284-5101 (toll free) +1-720-249-5700
EMEA	+800-7826-8888 (toll free) +49 6131 324 185
Asia Pacific	+800-7826-8887 (toll free) +603-7953-3010

For worldwide support:  
<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

## 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>



# Chapter 1: Introduction

---

This guide is intended to assist StorNext users to 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 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

This chapter introduces StorNext and contains the following topics:

StorNext Gateway Terminology .....	1
About the StorNext File System .....	2
About StorNext Storage Manager .....	3
About StorNext LAN Clients .....	3
StorNext Features .....	3
StorNext Limits .....	4
StorNext Use Cases and File System Restart Requirements .....	4

---

## 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
<b>StorNext Gateway</b>	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
<b>StorNext LAN Client</b>	A LAN-connected computer attached to a StorNext Gateway that has shared access to a StorNext SAN.	StorNext DLC
<b>StorNext Gateway Metrics</b>	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.

**i 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 the StorNext File System

The 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.

---

**i Note:** The maximum supported file name 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 [Licenses on page 22](#), and the *StorNext 5 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 [StorNext Reports on page 361](#).

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 [Licenses on page 22](#).

# StorNext Limits

Below is a list of StorNext limitations:

- The maximum number of LUNs per file system is 512.
- The maximum number of LUNs per data stripe group is 127.
- The maximum number of stripe groups per file system is 256.
- The maximum number of tape drives is 256.
- The maximum number of SAN and LAN clients is 1500.
- The maximum number of LAN Gateways is 64.
- The maximum file name length is 255.
- The maximum file path length is 1023.
- The maximum number of files in a non-managed file system is 5 Billion.
- The maximum number of files in a managed file system is 1 Billion.
- The maximum number of managed files on a single MDC is 1 Billion.
- The maximum number of file system managers (FSMs) per MDC is 8.
- The maximum number of StorNext mount points on a host is 16.
- The maximum size of a single file is 2 Petabytes.

---

**i 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.

---

# StorNext Use Cases and File System Restart Requirements

Use cases such as the following require a restart of StorNext services:

- Expanding the file system
- Modifying file system settings
- Adding or removing a High Availability (HA) license



## Chapter 2: StorNext Graphical User Interface (GUI) Overview

---

This section describes how to access and navigate through the StorNext GUI.

**i 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.

This chapter contains the following topics:

Accessing the StorNext GUI .....	5
The StorNext Home Page .....	8

---

## Accessing the StorNext GUI

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

**i Note:** The MDC host-name must be resolvable on DNS or have a local entry in the local MDC `/etc/hosts` file. If the MDC host-name can not be resolved, you will not be able to log in to the StorNext GUI from the web browser client.

## StorNext Browser Support

StorNext browser requirements are listed in the *StorNext 5 Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

---

**i 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 **Security** tab and then click **Custom level...**
4. Enable the following parameters in the **ActiveX controls and plug-ins** section:
  - **Allow Scriptlets.**
  - **Only allow approved domains to use ActiveX without prompt.**
  - **Run ActiveX controls and plug-ins.**
  - **Script ActiveX controls marked save for scripting\*.**
  - Set **Download signed ActiveX controls** to **Prompt**.
5. Enable **Display mixed content** in the **Miscellaneous** section.
6. Click **OK**.
7. Click the **Advanced** tab and un-check **Do not save encrypted pages to disk** in the **Security** section.
8. Click **OK**.
9. Close Internet Explorer 9 and then restart the application.

Use the following 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.

---

**i Note:** Typically, the port number is 81. If port 81 is in use, use the next unused port number. (for example, 82, 83, etc.)

---

**i 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.

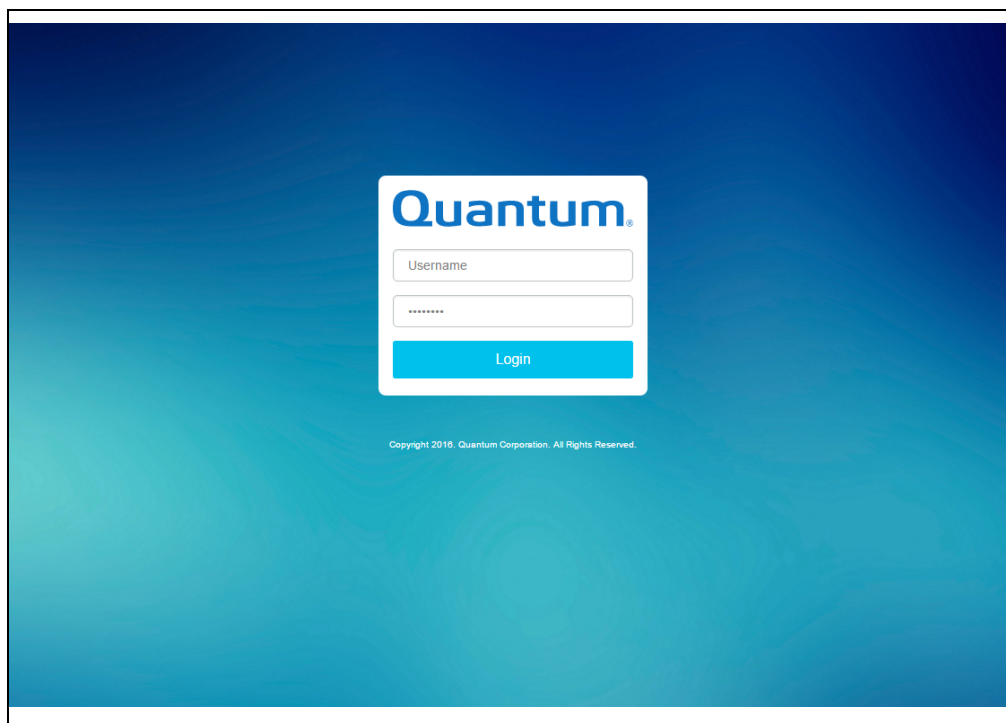
**For Internet Explorer:** Internet Explorer cannot display the web page.

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 Page



3. In the **User ID** field, type **admin**.
4. In the **Password** field, type **password**.
5. Click **Login**.

---

**i Note:** Depending on your browser configuration, you may receive a warning message informing you the web site's security certificate is not trusted. Refer to your browser documentation for procedures on how to accept and trust a self-signed SSL security certificate, or contact your local System Administrator for assistance.

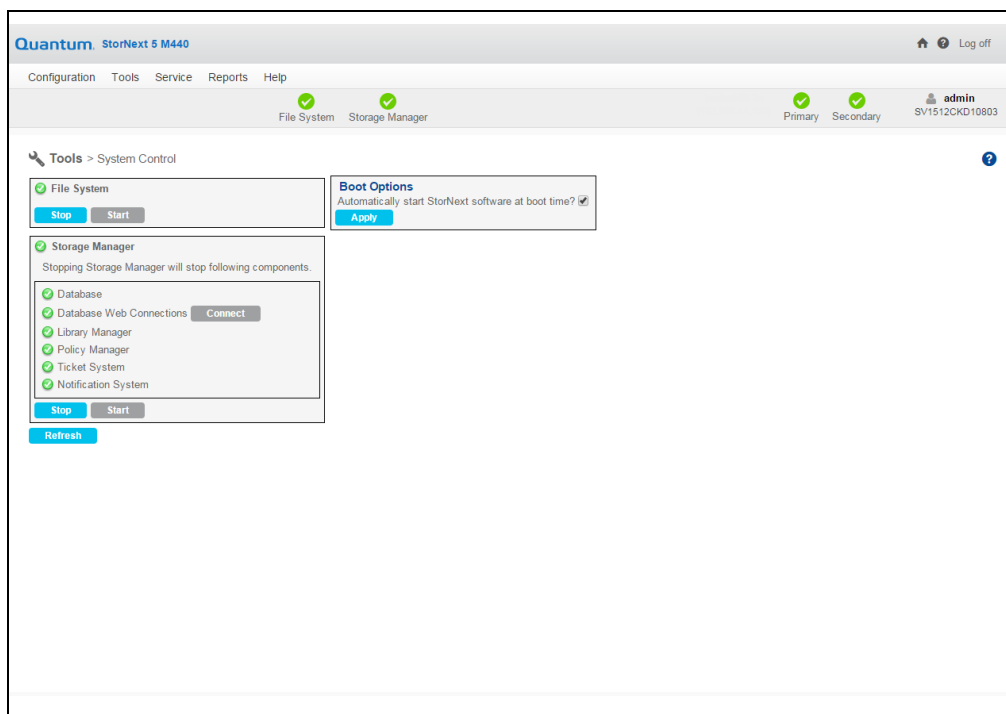


**Note:** When you start the StorNext GUI for the first time after installing StorNext, the GUI End User License Agreement is presented. Read the agreement and click **Accept** to start the GUI. After accepting the End User License Agreement, the **Configuration Wizard Welcome** page displays. The **Configuration Wizard** guides you step-by-step through the process of configuring StorNext. Proceed to [Configuration Wizard on page 20](#).

6. The StorNext **Home** page appears (refer to [The StorNext Home Page below](#)).

**Note:** If the StorNext File System and Storage Manager components are not currently started, the StorNext **Tools > System Control** page 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



## The StorNext Home Page

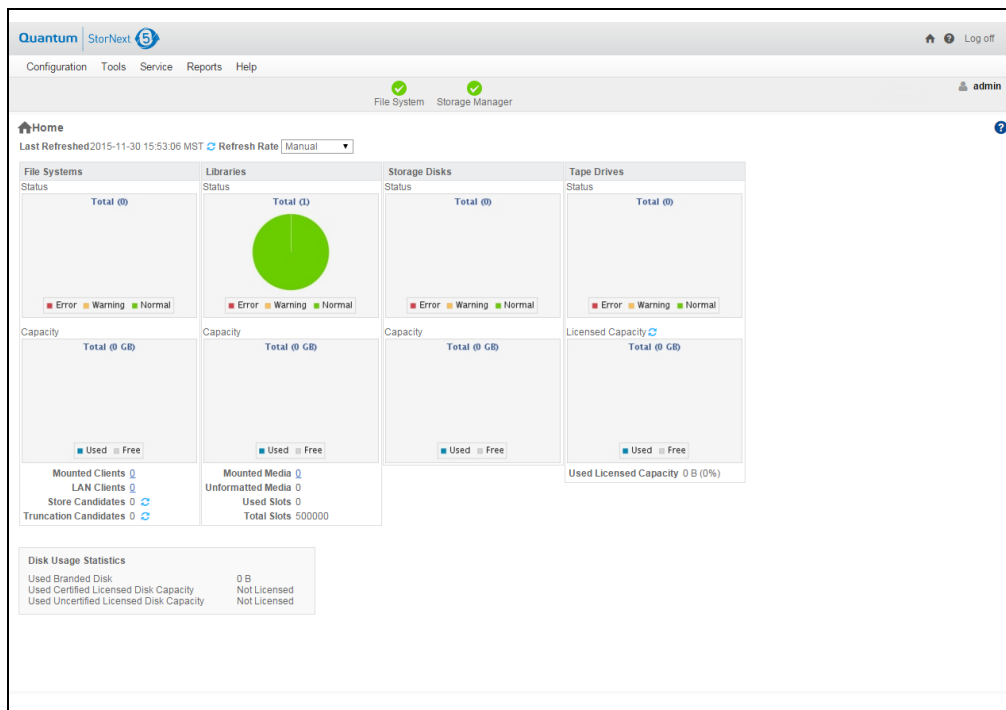
On the home page you will find the following:

- 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 on the next page](#)
- [Libraries Capacity Monitor on the next page](#)
- [Storage Disks Capacity Monitor on page 11](#)
- [Tape Drive Status on page 11](#)
- [Policy Capacity Monitor on page 11](#)

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.

---

**i 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

---

**i 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


The **Configuration** menu contains the following options that allow you to perform both initial and ongoing setup and configuration tasks for your StorNext system.

Option	Description
<b>Configuration Wizard</b>	Launch the StorNext Configuration Wizard.
<b>Licenses</b>	Enter or view license information for StorNext features.
<b>System</b>	Manage the cluster network interfaces between HA nodes, NTP, DNS, hostnames, gateway, domain and set bonding options. <b>i Note:</b> This option is only visible on StorNext M660, M440, M330, Pro Foundation, Artico Metadata Appliances, and Xcellis Workflow Director systems.
<b>Name Servers</b>	Enter and set order for servers used for StorNext file systems.
<b>File Systems</b>	Add a file system to your environment.
<b>Storage Destinations</b>	Add a library or storage disk, or set up data replication and deduplication. Additionally, you can add Object Storage Destinations, or Q-Cloud Storage Destinations.
<b>Storage Policies</b>	Add a storage policy to a file system.
<b>Email Server</b>	Configure the email server to use for notifications.
<b>Email Notifications</b>	Configure email notifications for Service Tickets, Admin Alerts, StorNext Backups, and Policy Class Alerts.

## The Tools Menu

The **Tools** menu contains options to control day-to-day operations of StorNext.

Menu Item	Description
<b>User Accounts</b>	Control user access to StorNext tasks.
<b>Client Download</b>	Download SNFS client software.
<b>Firmware Upgrade</b>	Perform a firmware upgrade for your HA system. <b>i Note:</b> This option is only visible on StorNext M660, M440, M330, Pro Foundation, Artico Metadata Appliances, and Xcellis Workflow Director systems.
<b>Dynamic Application Environment</b>	Displays the current status of your virtual machine environments. <b>i Note:</b> This option is only visible on StorNext M660, M440, Pro Foundation, Artico Metadata Appliances, and Xcellis Workflow Director systems.

Menu Item	Description
<b>System Control</b>	Stop or start the file system or StorNext Storage Monitor, and specify whether to automatically start StorNext at system startup.
<b>Object Storage Certificates</b>	View, create, import, convert, download, and delete Object Storage certificates.
<b>File and Directory Actions</b>	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.
<b>S3 Buckets</b>	Enables you to scan for, add new, and delete existing S3 buckets.
<b>StorNext Connect Ports</b>	Enables you to scan for, add new, and delete existing S3 buckets.  <b>Note:</b> This option is only visible on StorNext M660, M440, Pro Foundation, Artico Metadata Appliances, and Xcellis Workflow Director systems.
<b>File Systems: Label Disks</b>	Label disk drives.
<b>File Systems: Check File System</b>	Run a check on your file system before expanding the file system or migrating a stripe group.
<b>File Systems: Affinities</b>	Configure affinities for your file system.
<b>File Systems: Migrate Data</b>	Migrate the file system's stripe group(s).
<b>File Systems: Truncation Parameters</b>	Manage the file system's truncation parameters.
<b>File Systems: Manage Quotas</b>	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.
<b>Storage Manager: Storage Components</b>	View current status for libraries, storage disks, and tape drives; place one or more of these components online or offline.
<b>Storage Manager: Drive Pools</b>	Add, modify, or delete drive pools.
<b>Storage Manager: Media Actions</b>	Remove media from a library or move media from one library to another.
<b>Storage Manager: Storage Exclusions</b>	Specify types of file names to exclude from StorNext Storage Manager.
<b>Storage Manager: Truncation Exclusions</b>	Specify files or directories to exclude from the truncation process.
<b>Storage Manager: Tape Consolidation</b>	Enter parameters for automatically consolidating space on tape media.

Menu Item	Description
<b>Storage Manager: Library Operator Interface</b>	Enter or eject media from the Library Operator Interface.
<b>Storage Manager: Software Requests</b>	View or cancel pending software requests.
<b>Storage Manager: Scheduler</b>	Schedule file system events including Clean Info, Clean Versions, Full Backup, Partial Backup, and Rebuild Policy.
<b>Storage Manager: Alternate Store and Retrieval Location</b>	<b>Alternate Retrieval Location</b> allows you to specify a remote retrieval location to use in situations where files stored on tape or a storage disk cannot be accessed. <b>Alternate Store Location</b> provides an automatic system for copying files from a main instance of StorNext to a remote instance of StorNext at the same time as copies are made to tertiary storage at the main site.
<b>Storage Manager: Distributed Data Mover</b>	Spread the distribution of data across several machines rather than the primary server.
<b>Storage Manager: Drive Replacement</b>	Allows you to update the drive serial number mappings.
<b>Storage Manager: Client-side Encryption</b>	Lists the master keys that can be used for client side encryption.
<b>Storage Manager: System Parameters</b>	Allows you to set and modify StorNext system parameters.
<b>Replication/Deduplication: Administration</b>	View current replication process, or pause, resume, or stop replication.
<b>Replication/Deduplication: Replication Targets</b>	Add a host or directory for data replication, or edit existing replication targets.
<b>Replication/Deduplication: Replication Bandwidth</b>	Configure replication bandwidth limits and multi-link.
<b>High Availability: Convert</b>	Convert to a high availability configuration.
<b>High Availability: Manage</b>	Manage High Availability system parameters.

## The Service Menu

The **Service** menu contains options to monitor and capture system status information.


Menu Option	Description
<b>Health Check</b>	Perform one or more health checks on StorNext and view recent health check results.
<b>Capture State</b>	Obtain and preserve detailed information about the current StorNext system state.
<b>Capture DSET</b>	<p>Obtain and preserve detailed information about the current StorNext system state.</p> <p><b>i Note:</b> The information in this section is applicable for releases up to StorNext 5 release 5.3.x (in other words, prior to StorNext 5 release 5.4.0). Refer to <a href="#">Capture State on page 332</a> for the new information, beginning with StorNext 5 release 5.4.0.</p> <p>Obtain and preserve detailed information about the current state of your StorNext hardware.</p> <p><b>i Note:</b> This option is only visible on StorNext M660, M440, M330, Pro Foundation, Artico Metadata Appliances, and Xcellis Workflow Director systems.</p>
<b>System Backup</b>	Run a backup of StorNext software.
<b>Admin Alerts</b>	View informational messages about system activities.
<b>Tickets</b>	View, edit, or close service tickets generated for the system.
<b>Logging</b>	Enables robust debugging mode for advanced tracing.
<b>Web Services (V2)</b>	Specify your encryption and authentication options.

## The Reports Menu

The **Reports** menu contains options to view StorNext reports.

Menu Option	Description
<b>Logs</b>	Access logs of StorNext operations.
<b>Jobs</b>	View a list of pending and completed jobs on the system.
<b>Files</b>	View information about specific files, such as the owner, group, policy class, permissions, and copy information.
<b>Drives</b>	View information about the drives in your libraries, including the serial number and current state and status.



Menu Option	Description
<b>Media</b>	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.
<b>Q-Cloud Object Storage Media Usage</b>	View the usage report for Q-Cloud object store media.
<b>Relations</b>	View the names of the policy classes which correspond to the managed directories in your system.
<b>File Systems</b>	View file system statistics including active clients, space, size, disks, and stripe groups.
<b>SAN Devices</b>	View information about devices currently attached to your SAN, including disks/LUNs, libraries, and tape drives.
<b>Tape Consolidation</b>	View statistics on the tape consolidation (defragmenting) process.
<b>SAN and LAN Clients</b>	View statistics for StorNext clients, including the number of connected clients and LAN Clients, and client performance.
<b>LAN Client Performance</b>	View information about LAN Clients and servers, including read and write speed.
<b>Replication/Deduplication: Policy Activity</b>	View replication and deduplication performance statistics.
<b>Replication/Deduplication: Policy Summary</b>	View replication and deduplication information for each policy.
<b>Hardware Status</b>	<p>View up-to-date information about the system board and network ports for both nodes in your HA system, plus storage arrays.</p> <p> <b>Note:</b> This option is only visible on StorNext M660, M440, M330, Pro Foundation, Artico Metadata Appliances, and Xcellis Workflow Director systems.</p>
<b>Distributed Data Mover</b>	View activity related to the Distributed Data Mover feature.
<b>Gateway Metrics</b>	View information and activity related to your gateways, clients, and file systems.
<b>StorNext Metrics</b>	<p>View information and activity related to StorNext Metrics.</p> <p>The StorNext Metrics reports provide performance data logging and visual reporting and graphing features for StorNext systems. The StorNext Metrics reports are a visual reporting tool that combines comprehensive performance data logging with powerful visual reporting and analysis tools to help you identify potential problems and optimize system operations.</p>

## The Help Menu

The **Help** menu contains the following options to access StorNext documentation, find Quantum contact information, or detailed information about this version of StorNext.

## Documentation

On the **Help** menu, click **Documentation**.

## Support

On the **Help** menu, click **Support**. The **StorNext Support** page appears. The StorNext Support page includes the following information:

Label	Description
<b>Company</b>	Displays a hyperlink to the Quantum Corporation company web site. Click <b>Quantum Corporation</b> to open the company web site.
<b>Technical Assistance</b>	Displays telephone numbers for technical support around the world. Click <b>Quantum Technical Assistance Center</b> to open the Quantum Technical Support web site.

## About

The **Help > About** page displays various tabs, which display information about StorNext components. Depending on your configuration, the **Help > About** page displays the following tabs:

- [Software below](#)
- [Gateways on the next page](#)
- [Firmware on page 19](#)

## Access the Help > About Page

1. On the **Help** menu, click **About**. The StorNext **Help > About** page appears.
2. **(Optional)** Click **Refresh** to update the **Help > About** page.

## Software

The **Software** tab displays the following information about StorNext:

Component	Description
<b>StorNext</b>	The StorNext version and build number.

Component	Description
<b>User Interface</b>	The graphical user interface version and build numbers.
<b>Database</b>	The database version number.
<b>Perl</b>	The Perl programming language version and build number.
<b>Tomcat</b>	The Apache Tomcat version and build number.
<b>File System Server</b>	The version and build number of the file system server you are running.
<b>File System Client</b>	The version and build number of the file system client you are running.
<b>Library Manager</b>	The Library Manager version and build numbers.
<b>Policy Manager</b>	The Storage Policy Manager version and build numbers.
<b>Serial Number</b>	The StorNext serial number.
<b>Operating System</b>	The operating system.
<b>Patents</b>	Display Quantum Corporation patents.

## Gateways

The **Gateways** tab displays information of each gateway connected to the metadata controller (MDC).

Component	Description
<b>Server</b>	The Gateway IP address.
<b>System Serial Number</b>	If known, the gateway system serial number.
<b>StorNext Serial Number</b>	The StorNext software serial number running on the gateway, if specified in the license strings (optional).
<b>File System Server Version</b>	The StorNext software file system version running on the gateway.
<b>Gateway License</b>	Displays whether the gateway license is installed on gateway.  Yes: Signifies the gateway is running in unlimited LAN client connections mode available on certain gateways. No: Signifies the gateway client connections are limited to the number of licensed LAN client connections on the Gateway.

## Firmware

The **Firmware** tab displays the firmware levels for various hardware components in the system. This information is mostly for support purposes, if debugging a problem on a system.



# Chapter 3: The Configuration Wizard

---

This chapter contains the following topics:

Configuration Wizard .....	20
Welcome .....	22
Licenses .....	22
System .....	24
Name Servers .....	30
File Systems .....	34
Storage Destinations Overview .....	53
Storage Policies .....	98
Email Server .....	124
Email Notifications .....	125
Done .....	127

---

## 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. If you do not finish performing all the tasks, the wizard reappears whenever you return to the StorNext home page so you can resume completing tasks where you left off. For example, if you complete tasks 1 through 3, the next time the StorNext wizard appears you will be ready to complete task 4.

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

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

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

1. [Welcome on the next page](#): View disks and libraries currently available for StorNext usage.
2. [Licenses on the next page](#): Enter license information for StorNext features and components.
3. [System on page 24](#): Configure network settings. This step only applies to Xcellis, Artico, Pro Foundation, and StorNext Metadata Appliances.
4. [Name Servers on page 30](#): Specify and order the machines acting as StorNext name servers.
5. [File Systems on page 34](#): Add a StorNext file system.
6. [Storage Destinations Overview on page 53](#): Add a library, storage disks, and other storage destinations, such as **Object Storage** (see [Setting Up Object Storage Destinations on page 443](#)) and **Q-Cloud** (see [Setting Up Q-Cloud on page 481](#)).
7. [Storage Policies on page 98](#): Add a Storage Manager or replication storage policy.
8. [Email Server on page 124](#): Specify an email server to handle StorNext notifications.
9. [Email Notifications on page 125](#): Add email notifications recipients.
10. [Done on page 127](#): Signify that you are finished using the Configuration Wizard. You can also convert to a high availability (HA) system.

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

## High Availability Systems

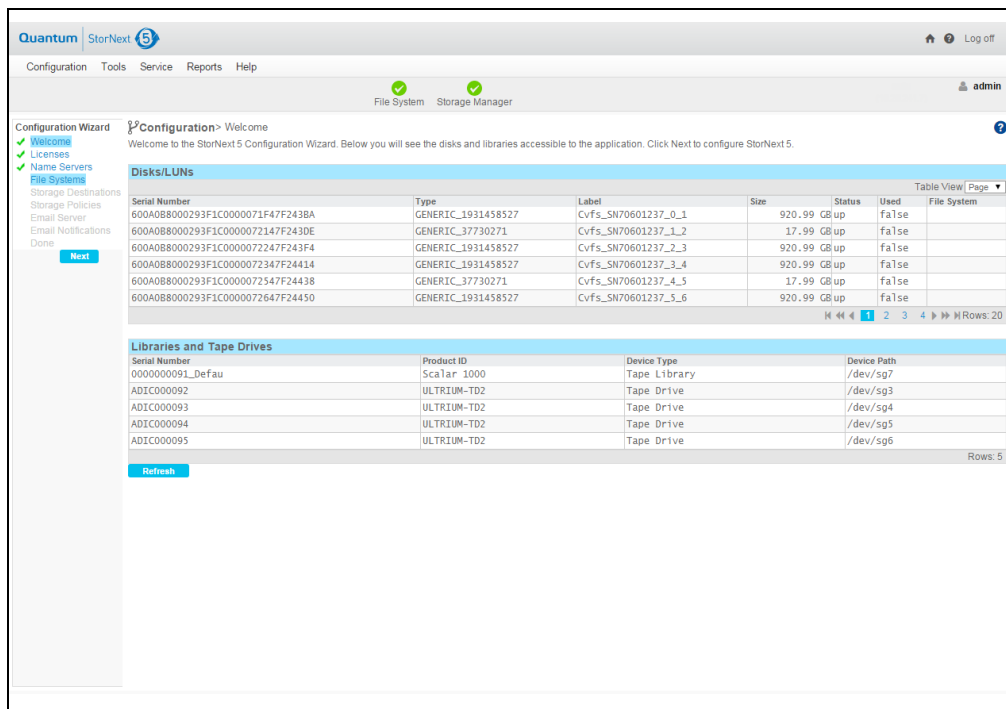
This topic contains some instructions that pertain to high availability (HA) systems, but if you plan to convert to HA you should read [Converting to HA on page 344](#). In particular, be sure to read and follow the [Pre-Conversion Steps on page 347](#).

# Welcome

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

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

Figure 4: Configuration > Configuration Wizard Welcome Page

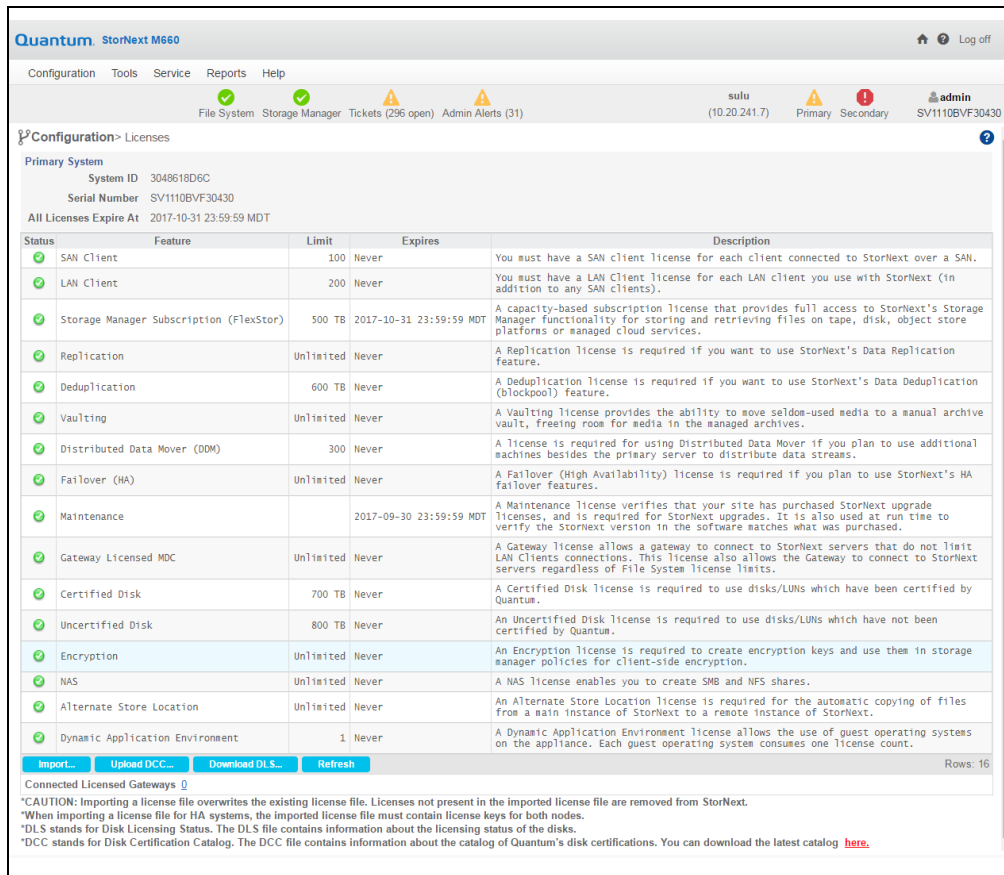


# Licenses

The **Licenses** page allows you to view current licenses and enter new licenses (or update expired licenses) for StorNext features you have purchased. Refer to the *StorNext 5 Licensing Guide* for complete details regarding StorNext license types, license expiration and limits, and other relevant information regarding StorNext licenses.

**Note:** You must have a license to configure or use StorNext features.

Figure 5: Configuration > Licenses Page



If the license.dat file does not contain permanent licenses, StorNext produces an auto-generated license with an expiration date for all StorNext products and features except Deduplication and Archive Conversion. In some cases Quantum may provide evaluation licenses for features. Evaluation licenses also have a fixed expiration date.

Beyond the evaluation period, you must have a permanent license to configure or use StorNext features.

**Caution:** Importing a license file overwrites the existing license file. Licenses not present in the imported license file are removed from StorNext.

**Note:** You cannot mix temporary and permanent licenses. For this reason, once you enter permanent license information ALL of your temporary licenses will be deleted, even if the 30-day period has not expired. For example, if you do not obtain a permanent license for the Distributed Data Mover feature, as soon as you enter permanent licenses for other StorNext products and features, the temporary license for Distributed Data Mover will become invalid. If you want to use the 30-day period to demo StorNext products and features, plan accordingly before entering any permanent licenses.



---

# System

The **System** page allows you to enter or view network configuration information for the primary and secondary nodes. You can also set the system date and time. Your system comes pre-configured for its environment during system installation. IP addresses for the primary and secondary nodes were entered, and the system is configured to operate as High Availability (HA) pair of nodes.

**i Note:** The **System** option is only available on Xcellis, Artico, Pro Foundation, and StorNext Metadata Appliances.

Below is information about the **Configuration > System > Network** page.

## General Settings

The **Configuration > System > Network** page displays pre-configured information about the primary and secondary nodes configured during the initial StorNext installation. The person who initially installed and configured your StorNext Metadata Appliance will have assigned IP addresses for the Primary and Secondary Nodes.

To view or enter network configuration for the primary and secondary nodes in your HA system, perform the following steps.

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

## Virtual Machine Bridging Settings

**i Note:** The **Virtual Machine Bridging Settings** option is **only** available on Xcellis Workflow Director systems with the Dynamic Application Environment license installed.

Modify the virtual machine bridging settings as needed. StorNext supports the recognition of virtual machine allocated NIC ports, as well as the ability to configure the network bridges needed for communication between the node and the virtual machine. In this section you configure the virtual machine bridging settings.

- Virtual machine bridging is required if your virtual machine instance is going to access StorNext File System shares on the metadata network. The virtual machine bridge allows the virtual machine to communicate with the local node from its virtual network interface. There are a maximum of 2 bridges that can be configured in the system; one is expected to be used for metadata network traffic and the other for NAS client traffic. Both are optional in case the virtual machine instance does not require this type of communication.
- When a virtual machine is configured, there is also an option to allocate available dedicated physical network ports for exclusive use by the virtual machine. This is referred to as the “Application” network for dedicated virtual machine application use. Once a port is allocated to a virtual machine, it will no longer be configurable from the StorNext GUI for use by the MDC node since it is reserved for the virtual machine use only. Once the virtual machine is removed from the system, the network port(s) become available for use.

## Bonding Settings

Modify the bonding settings as needed. In this section you specify the bonding mode used by the network interface. Bonds are used by StorNext Metadata Appliances 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 page.

---

**i Note:** If you break the bonds for Bond0 and Bond1, 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 Bond1, 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 Bond0 and Bond1.

### Bonding Mode

The current bonding mode is displayed in this column. There are two modes supported in StorNext Metadata Appliances, **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 StorNext Metadata Appliances. 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.

---

**i 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.

## Bonding rules

- You need two or more ports for a bond.
- 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

- You cannot bond ports with different Ethernet speeds. For example, you cannot bond a 10GbE port with a 1GbE port.
- 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

Enter the appropriate information into the following fields:

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

Parameter	Description
<b>The first column</b>	This lists the current network bond and Ethernet ports available for configuration.
<b>Alias</b>	<p>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.</p> <p><b>i Note:</b> Up to 10 aliases may be defined for every Ethernet port.</p> <p><b>i Note:</b> For systems that have had VLAN tagging manually configured, the following apply:</p> <p>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 <b>bond2.345</b> where <b>345</b> is the assigned VLAN tag. You can view the settings in the GUI, but should avoid changing any VLAN device settings using the GUI.</p>
<b>Primary IP</b>	This is the IP address of the Primary Node.
<b>Secondary IP</b>	This is the IP address of the Secondary Node.
<b>Netmask</b>	This is the netmask for this Bond or Ethernet connection.
<b>Gateway</b>	This is the Gateway used for this Bond or Ethernet connection.

Parameter	Description
<b>External Host</b>	Select an interface from the Configuration Table on the <b>Configuration &gt; System &gt; Network</b> page by clicking the <b>External Host</b> radio-button for that interface. The selected interface IP address will be added to the file <code>/etc/hosts</code> as the external host IP address of the host (the IP address will be used for the FQDN entry for the host in the <code>/etc/hosts</code> file). When the <b>Configuration &gt; System &gt; Network</b> page is loaded, the <b>External Host</b> radio-button is set to the current external host IP.  <b>Note:</b> You can only select an external host whose IP address' network matches that of the default gateway IP address' network.
<b>Jumbo Frame</b>	Use this option for networks which support this type of ethernet payload.

Figure 6: Example of Network Configuration Settings

Configuration								
	Alias	Primary IP	Secondary IP	Netmask	Gateway	External Host	Jumbo Frame	
bond0 (Metadata)	bond0:2	10.60.192.41	10.60.192.42	255.255.240.0	10.60.192.1	↔	<input type="checkbox"/>	
bond1 (SysMgmt)	bond1:2	10.60.226.41	10.60.226.42	255.255.255.0	10.60.226.1	↔	<input type="checkbox"/>	
bond2	bond2:2	10.60.224.41	10.60.224.42	255.255.255.0	10.60.224.1	↔	<input type="checkbox"/>	
eth8	eth8:1	17.218.184.173	17.218.184.174	255.255.252.0	17.218.184.1	↔	<input type="checkbox"/>	
eth9	eth9:1	17.218.188.173	17.218.188.174	255.255.252.0	17.218.188.1	↔	<input type="checkbox"/>	

## 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
<b>The first column</b>	Defines the interface for which a static route is defined. The packets for the destination of the static route are sent through this interface.	
<b>Destination IP</b>	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 <b>Destination IP</b> address and <b>Destination Netmask</b> may be a network address or a host address.  The combination will be a host address if the <b>Destination IP</b> address is a host IP address and the <b>Destination Netmask</b> is 255.255.255.255, and the defined route will be a host route.
<b>Destination Netmask</b>	Enter the netmask address to use as the netmask for the <b>Destination IP</b> address.	The combination will be a network address if the <b>Destination IP</b> address is a host or network IP address and the <b>Destination Netmask</b> is not 255.255.255.255, and the defined route will be a network route.

Parameter	Description	Notes
<b>Destination Gateway</b>	Enter the IP address of the gateway to send the packets to. The IP address reflects the destination defined with the combination of the <b>Destination IP</b> and the <b>Destination Netmask</b> .	
<b>Target Node</b>	Select the node(s) where the static route is configured.  Available node choices: <ul style="list-style-type: none"> <li>• Both</li> <li>• Primary</li> <li>• Secondary</li> </ul>	The <b>Target Node</b> option allows you to define static routes. The default for a new static route entry in the <b>Routing</b> table (host-route, or subnet-route) is <b>Both</b> .  You can also apply the route to only the <b>Primary</b> , or the <b>Secondary</b> target node instead by changing the selection in the <b>Target Node</b> drop-down list. In other words, not only host-routes, but also subnet-routes may be applied to either the <b>Primary</b> or the <b>Secondary</b> target node as well.

Figure 7: Example of Network Routing Configuration Settings

Routing	Destination	Netmask	Gateway	Target Node
eth8:1	17.218.184.172	255.255.255.255	17.218.184.173	Primary
	17.218.184.179	255.255.255.255	17.218.184.173	Primary
	17.218.184.174	255.255.255.255	17.218.184.173	Primary
	17.111.120.5	255.255.255.255	17.218.184.1	Both
	17.111.120.6	255.255.255.255	17.218.184.1	Both
	17.111.120.133	255.255.255.255	17.218.184.1	Both
	17.111.120.134	255.255.255.255	17.218.184.1	Both
	17.218.184.172	255.255.255.255	17.218.184.174	Secondary
	17.218.184.179	255.255.255.255	17.218.184.174	Secondary

## Status

Review your network configuration settings and make changes as necessary.

## Date and Time Settings

To set the system date and time, perform the following steps:

1. On the **Configuration > System** page, click the **Date & Time** tab. The **Configuration > System > Date & Time** page appears.
2. Select 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.

## Set the Date and Time via NTP

1. After you select **Network Time Protocol (NTP) Server (Recommended)**, select one of the following:
  - **Select a Timeserver Pool**
  - **Specify a Timeserver or Pool**
2. If you select **Select a Timeserver Pool**, select the desired server from the dropdown list next to **Timeserver** or **Pool**.
3. If you select **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 your StorNext Metadata Appliance.

## Set the Date and Time Manually


1. After you select **Manual Date & Time Settings**, click **Update Date & Time**.
2. Input the values into the following fields:
  - **New System Date**
  - **New System Time**
3. In the **System Timezone** list, select the area corresponding to the location of your StorNext Metadata Appliance.

## Apply 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**.)

---

 **Caution:** Changing the network configuration requires a system reboot (both nodes) in order for all system services to function correctly. Rebooting the system may take several minutes.

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**.

---

 **Note:** After the network configuration has been saved, please close your browser window and wait for at least 5 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.

---

# Name Servers

The **Name Servers** page allows 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).

---

**i 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).

---

**i 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 (), but this is not a requirement.

If the `fsnameservers` file does not exist, is empty or contains the localhost IP address (127.0.0.1), the file system operates as a local file system requiring both a client and a server. The file system will not communicate with any other StorNext File System product on the network, thus eliminating sharing the FSS over the SAN.

The addresses in the `fsnameservers` file define the metadata networks and therefore the addresses used to access file system services. When an 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.

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.

## View Name Servers

When you access this page by choosing **Name Servers** from the **Configuration menu**, a list of any previously entered IP addresses appears. 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.

## Set the Order of a Name Server


StorNext allows you to specify the order in which name servers are used.

- Select a name server and then click **Move Up** or **Move Down** until the selected server is in the correct order.


## Add a Name Server

1. Enter the IP address in the field to the left of the **Add** button.
2. Click **Add**.
3. When the confirmation message appears, click **Yes** to proceed or **No** to abort.
4. When the new name server appears in the list of available name servers, you can reposition the new name server by using the **Move Up** or **Move Down** buttons.

---

 **Caution:** Changing the name server file affects all servers and clients on your SAN. All servers and clients **MUST** have the same name server file in order to ensure correct SAN operation. After the name server file has been updated on the server, the SAN administrator must copy the name server file to ALL connected clients and then restart StorNext File System services on those clients.

## Delete a Name Server

-  **Note:** Although not required, you might want to disable the name server before deleting it to prevent complications.
- Select the name server you want to delete and then click **Delete**.

## Configure 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 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 systems 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.

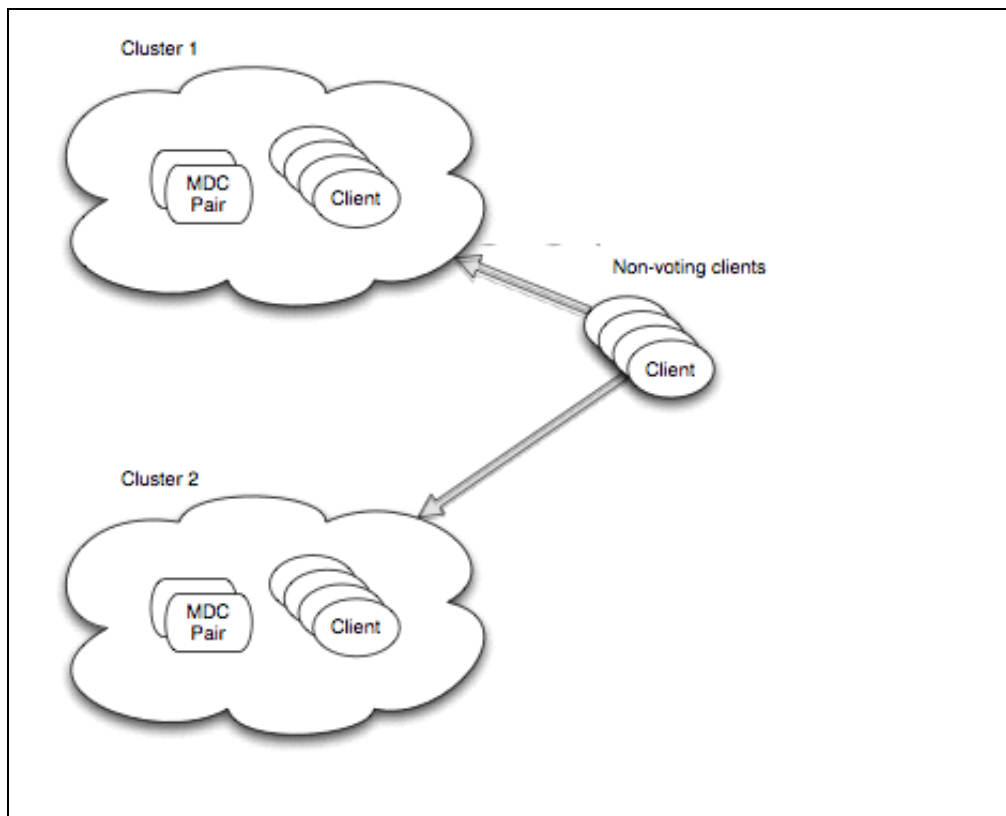
The format for the `fsforeignservers` file is similar to the `fsnameservers` file in that it contains a list of IP addresses or hostnames, preferably IP addresses. One difference is that the addresses in the `fsforeignservers` file are MDC addresses, addresses of hosts that are running the FSM services. 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 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.



## 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 s.
- 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-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 , 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- 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 4 Mbits 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  $500 \times 4$  or 2000 Mbits or 2 Gbits of heartbeat reply messages per second. If the 4 `fsnameservers` addresses belonged to two nameservers, each server would be generating 1 Gbit of heartbeat reply messages per second.

---

**i Note:** During stable periods, the heartbeat rate for non- 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.

---

## File Systems

Whether you access the **Setup > File System** page from the Configuration Wizard or by selecting **File Systems** from the **Configuration** menu, you can view, add, edit, or delete a file system. The procedures are the same regardless of your method of access.

When you reach this step, any previously created file systems are displayed.

---

**i Note:** See [StorNext Limits on page 4](#) for a list of StorNext limitations.

---

**i Note:** For Xcellis, Artico, Pro Foundation, and StorNext Metadata Appliances, the High Availability (HA) shared file-system is already pre-configured from the factory.

On the **Configuration** menu, click **File Systems** to display the **Configuration > File Systems** page and perform the following file system-related tasks:

- [View a File System on page 37](#)
- [Add a File System on page 38](#)
- [Edit a File System on page 45](#)
- [Delete a File System on page 51](#)
- [Perform File System Actions on page 51](#)

## Configuring Spotlight Proxy

Spotlight Proxy enables communication between a StorNext MDC and a Xsan Spotlight server. The Spotlight server enables Spotlight search functionality for Xsan clients accessing StorNext volumes. While Spotlight includes two configurable Spotlight search levels, **ReadWrite** and **FsSearch**, only the **ReadWrite** level works with Spotlight Proxy. Quantum recommends not to configure the **FsSearch** level as it will disable Spotlight Proxy.

Spotlight Search Level	Description
<b>ReadWrite</b>	The <b>ReadWrite</b> level indexes attributes as well as read and index each file in the volume. By default, Spotlight is configured for <b>ReadWrite</b> indexing.
<b>FsSearch</b>	The <b>FsSearch</b> level only indexes file attributes.

**Note:** If the **ReadWrite** search level is configured, every file in the volume is read once by the Spotlight indexing process. For Managed file systems, this includes retrieving and reading all truncated files. Consider the impact of Spotlight indexing and ensure that Spotlight indexing is appropriate for your environment.

**Caution:** Enabling Spotlight Proxy on a managed file system is not advised. Spotlight indexes data as well as metadata and may cause truncated files to be restored from archival storage.

### Enable Spotlight

1. On the **Configuration** menu, click **File Systems**. The **Configuration > File Systems** page appears.
2. Perform one of the following:
  - a. Edit an existing file system. See [Edit a File System on page 45](#).
  - b. Create a new file system. See [Add a File System on page 38](#).
3. Expand the **Advanced Parameters** heading, and then click the **Features** tab.
4. Select the **Spotlight Proxy** check box.
5. Click **Apply** to save the changes, or click **Cancel** to return to the **Configuration > File Systems** page.

### Configure Spotlight on a macOS X version 10.11 (or later) Xsan Client

**Note:** You must configure a Mac client to act as a Spotlight Proxy server.

1. Set with:

```
defaults write /Library/Preferences/com.apple.xsan enableSpotlightServer -  
bool true
```

2. Enable Spotlight for external volumes (SNFS or otherwise):

```
mdutil -i on /Volumes/RH6-DATA
```

---

**i Note:** If the error message, "Error: unknown indexing state" is displayed, then use the following commands to reload the Xsan client.

```
launchctl unload /System/Library/LaunchDaemons/com.apple.xsan.plist  
launchctl load -w /System/Library/LaunchDaemons/com.apple.xsan.plist
```

## Troubleshoot Spotlight

**i Note:** Multiple Xsan clients can be configured as Spotlight proxy servers.

The first proxy server the fsm locates becomes the proxy. If that client leaves the cluster, the fsm sequentially scans the list of clients willing to be the proxy server and requests that each of them try to renegotiate to become the new proxy. The first client to respond affirmatively becomes the new proxy server until it leaves the cluster, even if the previous proxy server rejoins.

### Possible Troubleshooting Items

- To clear the Spotlight state for an external volume, execute the following

```
$ sudo mdutil -E -i off /Volumes/peterson_47_acls_OFF
```

- With sufficient file system load, and/or sufficient load on the Spotlight proxy server, there can be a noticeable delay between something changing in the file system and Spotlight search reflecting that change.

## 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 5 Tuning Guide*.

## View a File System

### View a File System Report, Usage Statistics and Current Option Parameters

1. From the **Configuration > File System** page, select the file system you wish to view.
2. Click **View**.
3. To view the file system's parameters, click **Option Parameters**. (Click again to hide parameter information.)
4. Click **Done** to return to the **Configuration > File System** page.

### File System Report

The file system report provides the following information:

Parameter	Description
<b>File System</b>	The name of the file system whose information is displayed.
<b>Licensed SAN Clients</b>	The number of licensed SAN clients in the file system.
<b>Licensed LAN Clients</b>	The number of licensed LAN clients in the file system.
<b>Licensed Gateways</b>	The number of licensed gateways in the file system.
<b>Mount Point</b>	The file system's mount point name.
<b>SAN Clients</b>	The number of SAN clients in the file system.
<b>LAN Clients</b>	The number of LAN clients in the file system.
<b>Gateways</b>	The number of gateways in the file system.
<b>Status</b>	The current status, indicated by a green checkmark icon (ok) or a red x icon (unknown or down).

## Stripe Groups

Parameter	Description
<b>Stripe Group</b>	A list of stripe groups associated with the file system

## Usage

Parameter	Description
<b>Total Space</b>	The total amount of space in the file system
<b>Total Inodes</b>	The total number of inodes currently in the file system
<b>Used Space</b>	The amount of space currently consumed in the file system
<b>Free Space</b>	The amount of free space currently available in the file system

## Option Parameters

Displays the file system parameters that are defined when you create a new file system (see **Advanced Parameters** in [Add a File System below](#)), or when you edit an existing file system (see **Advanced Parameters** in [Edit a File System on page 45](#)).

## Add a File System

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

### Add a File System

1. Choose **File Systems** from the **Configuration** menu. The **Configuration > File System** page displays all currently configured file systems. (You can also access this page from the StorNext Configuration Wizard by choosing the **File Systems** option.)
2. Click **New** to add a new file system.
3. Enter the following fields:
  - **File System Name:** Enter the name for the new file system.
  - **Mount Point:** Enter the mount point for the new file system, or accept the displayed default mount point.
  - **Storage Manager:** Select this option if you want this file system to be managed by StorNext Storage Manager.

- **Replication/Deduplication:** Select this option if you would like to enable data replication/deduplication on the new file system.
  - **Stripe Group Configuration:** Choose **Generated** or **Manual**. When you choose **Generated**, StorNext creates the file system with typical parameters after you enter basic configuration information. When you choose **Manual**, you are given the opportunity to manually create the file system by specifying all parameters. If you choose **Manual** configuration, proceed to **Manual Configuration**.
4. If you choose **Generated** configuration, click **Continue** to proceed to the second configuration page.
  5. Enter the following fields:
    - **RAID Type:** Select from the dropdown list the RAID type that corresponds to your system configuration. (If you are using a StorNext G300 Gateway Appliance, the default value is **Quantum Disk**.)
    - **Data Disks per LUN:** The number of data disks per LUN.
    - **Segment Size (Bytes):** The amount of data written to one drive in a RAID LUN before writing data to the next drive in RAID LUN. Configure the **Segment Size** using the RAID user interface.
    - **Data Stripe Breadth:** The **Data Stripe Breadth** is the amount of data that StorNext writes to a LUN before switching to the next LUN within a stripe group. For many RAIDs, setting the **Data Stripe Breadth** to the product of the (**Segment Size x Disks per LUN**) results in best performance.
- 
- Note:** Required fields are marked by an asterisk (\*).
6. Select one or more disk to assign to the file system. (Use the checkbox column to select or deselect.)
  7. After selecting one or more disks, click **Meta** to designate the disk(s) as used for metadata, or click **Journal** to use the disk(s) for journaling. A disk can be used for both metadata and journaling.
  8. At the field to the left of the Label button, enter a label name. Click **Label** to apply the label name to the selected labels. (Clicking **Unlabel** removes label names from selected labels.)
  9. 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.)
  10. Click **Continue**.
  11. 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.
  12. 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. (If necessary, click the arrow to the left of these headings to display the tabs and fields.)



1. When you are finished entering **Advanced Parameters** 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**.

## Advanced Parameters > Allocation Tab

The **Allocation** tab contains fields that affect how resources are allocated on your file system.

- **Journal Size:** Defines the size of the file system journal.
- **Strategy:** Defines the method for choosing stripe groups when allocating disk blocks. Options are **Round**, **Fill**, or **Balance**.
- **Reserved Space:** This option enables delayed allocations on clients. Reserved space is a performance feature that allows clients to perform buffered writes on a file without first obtaining real allocations from the metadata controller. The allocations are performed later during periodic cache synchronization.  
  
If the **Reserved Space** option is not enabled, slightly more disk space will be available, but file fragmentation may increase and performance may not be satisfactory.
- **StripeAlign Size:** Defines the minimum allocation size to trigger automatic stripe-aligned allocations.
- **Inode Stripe Width:** If non-zero, causes large files to have their allocations striped across stripe groups in chunks of the specified size.
- **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. At the first field enter the desired chunk size. At the second field specify the chunk unit of measure (B=bytes, KB=kilobytes, MB=megabytes, GB=gigabytes, TB=terabytes). For more information about the **Allocation Session Reservation** feature, refer to the *StorNext File System Tuning Guide*.
- **Affinity Preference:** If checked, permits files of a particular affinity to have their allocations placed on other available stripe groups (with non-exclusive affinities) when the stripe groups of their assigned affinity do not have sufficient space. Otherwise, allocation attempts will fail with an out-of-space error.
- **Auto Affinities:** This table designates the affinity (stripe group[s]) to which allocations will be targeted for all files on this file system whose name has the specified file extension. To add a new entry to the **Auto Affinities** table, type a file extension (omit the "." dot), select an affinity from the **Affinity** menu, and click **Add**. Use the file extension \* (asterisk) to indicate "all other" file extensions that are not explicitly listed. The **Affinity** menu will list only affinities that are currently assigned to a stripe group of this file system. The affinity **NoAffinity** indicates that allocations will be targeted at stripe groups that have no affinity. To delete one or more entries, simply check the row(s) to delete, and click **Delete**. Each unique file extension can be targeted at only one affinity. However, each affinity may serve as the allocation target for more than one file extension. To sort the auto affinities table by file extension or affinity name, click the **File Extension** or **Affinity** column title, respectively. Multiple clicks cause the sort order to alternate between ascending and descending alphabetic order.

## Advanced Parameters > Performance Tab

The **Performance** tab fields allow you to adjust parameters for optimized performance.

- **Buffer Cache Size:** Defines the amount of memory used by the FSM process for caching metadata.
- **Inode Cache Size:** This value defines the number of inodes that can be cached in the SNFS server. The default and minimum setting for the cache size is 16.
- **Use Physical Memory Only:** When this option selected, the file system will use only physical memory, not swapped or paged.
- **High Priority FSM:** Determines whether the FSM process should run with real-time priority.

## Advanced Parameters > Debug Tab

The **Debug** tab fields allow you to enable or disable debugging and influence the debug log.

- **Enable Debugging:** Enables detailed file system debug tracing. When debug tracing is enabled, file system performance could be significantly reduced.
- **Debug Log Settings:** Settings to turn on debug functions for the file system server. The log information may be useful if a problem occurs. A Quantum Technical Assistance Center representative may ask for certain debug options to be activated to analyze a file system or hardware problem.
- **Maximum Log Size:** This value defines the maximum number of bytes (size) to which a StorNext Server log file can grow. When the log file reaches the specified size, it is rolled and a new log is started. In this situation, the two log files could use twice the maximum log size space specified in this field. The range is from 1 to 32 megabytes.
- **Maximum Number of Logs:** This value determines the number of rolled logs kept. Choices range from 4 to 64.
- **OP Hang Limit (Seconds):** Defines the time threshold (in seconds) used by the FSM process to discover hung operations.

## Advanced Parameters > Features Tab

The **Features** tab fields allow you to enable or disable various file system-related features.

- **Security Model:** Determines the scheme for specifying and enforcing security policies. The available options are **legacy**, **unixpermbits**, and **acl**. The default value is **legacy**.
  - If the **Security Model** is **legacy**, then the **Unix Id Mapping** field is grayed out (disabled); however, the **Windows Security** option and the **Enforce ACLs** option remain enabled.
  - If the **Security Model** is **acl**, then the **Unix Id Mapping** field is not grayed out; however, the **Windows Security** and the **Enforce ACLs** are grayed out (disabled).
  - If the **Security Model** is **acl**, the **Unix Id Mapping** field is not allowed to be **none**. You must select a value from the **Unix Id Mapping** list.

- **Unix Id Mapping:** Determines the Unix Id mapping. The available options are **none**, **algorithmic**, and **winbind**. The default value is **none**.
- **Windows Security:** Determines whether Windows ACLs are enabled for the file system.
- **Enforce ACLS:** Determines whether ACLs are enforced on XSan clients.
- **Windows Global ShareMode:** Determines whether Windows Global ShareMode is enabled for the file system.
- **Quotas:** Determines whether the Quota feature is enabled for the file system. 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.

---

**i Note:** It is not possible to use the StorNext GUI to manipulate the actual quotas, so you must use **snquota** to modify the quotas on individual users. However, you cannot enable or disable quotas entirely through that interface. While user and group quotas cannot be manipulated via the StorNext GUI, directory quotas can be set on the **Manage Quotas** page (for additional information, see [Manage Quotas on page 322](#)).

- **Quota Logs Retention Period:** If **Quotas** is enabled, you can configure the length of time (**Days**, **Weeks**, **Years**) to retain the **Quota** logs. In other words, **Quota** logs are not retained after the length of time determined by the **Quota Logs Retention Period** value.
- **Named Streams:** Determines whether a file system includes support for the Xsan Named Streams feature. Accessing files with Named Streams from a non-Xsan client is not supported. Also, **Named Streams** enabled file systems cannot be configured as Storage Manager, Replication or Deduplication enabled file systems. The **Named Streams** feature enables the storing of additional file system metadata and as such the **Named Streams** feature cannot be disabled after it has been applied to a file system.
- **Spotlight Proxy:** Determines if Spotlight proxy is enabled for the file system. For additional information, see [Configuring Spotlight Proxy on page 35](#).
- **Use Active Directory SFU:** Determines if Active Directory is enabled for the file system.
- **File Locks:** Determines whether the FSM tracks and enforces file locks across all clients.
- **FileLock Resync Timeout:** Defines the timeout for clients re-registering file locks following FSM failover.
- **Restore Journal:** The parameter is used to enable or disable restore journal (also known as metadump or metadata dump) creation by the FSM process. A restore journal logs file system operations and is a key piece of restoring a file system after a disaster on non-managed or managed file systems. By default, restore journal creation is disabled on non-managed file systems, and enabled on managed file systems.

## Advanced Parameters > LDAP Tab

The **LDAP** tab fields allow you to enter parameters related to LDAP (Lightweight Directory Access Protocol, an application protocol for querying and modifying directory services running over TCP/IP).

- **Unix File Creation Mode on Windows:** The number of mode bits for UNIX files
- **Unix Directory Creation Mode on Windows:** The number of mode bits for UNIX directories

- **Unix Nobody UID on Windows:** UNIX user ID to use if no other mapping can be found
- **Unix Nobody GID on Windows:** UNIX group ID to use if no other mapping can be found
- **Unix ID Fabrication on Windows:** Allows you to enable or disable using fabricated IDs on a per-file system basis. If enabled, Windows user IDs are mapped using fabricated IDs.

## Advanced Parameters > Special Tab

The **Special** tab fields allow you to enable or disable special file system-related features.

- **Inode Delete Maximum:** Determines throttling factor for reclaiming space from files allocated using PerfectFit mode.
- **Trim on Close:** Determines whether excess allocations are removed from files when they are closed.
  - **Global Super User:** Enable this option (check the box) to allow a user with super-user privileges to assert these privileges on the file system.
    - If the **Global Super User** option is *enabled*, super users have global access rights on the file system. This selection is the same as the `maproot=0` directive in the Network File System (NFS).
    - If the **Global Super User** option is *not enabled*, super users can modify only files they can access, like any other users.
- **File System Capacity Threshold:** Defines the file system fill level (in percent) that triggers a RAS event.
- **Extent Count Threshold:** Defines the number of extents in a file required to trigger a fragmentation RAS event.
- **Remote Notification:** Determines whether to enable partial support for cluster-wide Windows directory event notification.

## Stripe Groups/Disk Management Fields

To modify an existing stripe group, under the **Stripe Groups** heading select the stripe group you want to modify, and then change its properties as desired.

To add a new stripe group to the file system, click **Add** and then enter the remaining fields for the new stripe group.

When you are finished on the **Stripe Group** tab, click **Apply** to save your changes, or **Cancel** to abandon your changes.

- **Stripe Group:** Select the stripe group you want to modify or delete.
- **Add:** Click this button to add a new stripe group.

---

**i Note:** If the following indented fields are not displayed, they appear after you click **Add**. Likewise, after you delete a stripe group these fields may not be displayed.

- **Name:** Enter a name for the new stripe group, or skip this field to accept the displayed name.
- **Breadth:** Specify the stripe group breadth, which is the number of kilobytes (KB) that is read from or written to each disk in the stripe group.
- **Content:** Specify whether the stripe group will be used for metadata, journaling, or user data. You can specify one, two, or all content types.
- **Delete:** Click this button to delete the currently selected stripe group.

---

**i Note:** This particular delete function does not provide a confirmation message, so be absolutely sure you want to delete the selected stripe group before you click **Delete**. The selected stripe group will be deleted immediately after you click **Delete**, and there is no Undo feature.

- **Affinity:** An affinity is a label that allows you to control the physical location of files. Managed file systems are restricted to only using two affinities and are used for Disk-to-Disk relocation Storage Manager policies. Managed file systems are limited to using the same two affinity names on all managed file systems. By default, the affinity names on a managed file systems are “Tier1” and “Tier2”. For unmanaged file systems, using affinities is a two step process: first, each stripe group can be assigned one or more affinities during file system configuration, then associate a directory with the affinity. For example, if you configure stripe group SG2 to have affinity AFF2 and then associate directory special\_files with affinity AFF2, all files put into special\_files can exist only on the disks that make up SG2. If an affinity were not assigned and associated with special\_files, the files put into that directory could exist on any stripe group or any disk in the file system. It makes sense to use affinities in environments where performance is critical. For example, you might want to constrain video files to a stripe group made of fibre channel disks tuned for video playback, but have audio files reside on slower SCSI disks in a different stripe group. A stripe group can have multiple affinities.

You may not remove an affinity from a stripe group if that affinity is not assigned to any remaining stripe group AND the Auto Affinities table includes a file extension that is targeted at that affinity. You must first update the Auto Affinities table to delete that auto affinity entry. See the [Auto Affinities](#) section for auto affinity delete instructions.

If you want to associate an affinity to the new stripe group, select the desired affinities.

- **Exclusive:** When this option is enabled, the selected stripe group is used exclusively for the affinity's files.
- **Access:** Specify the permission level for the stripe group: Full R/W (read/write), Read Only, or Disabled.
- **Quality of Service:** Specify parameters for the Quality of Service (QOS) feature. QOS allows real-time applications to reserve a specified amount of bandwidth on the storage system. (For more information about the QOS feature, see the StorNext Installation Guide.)
- **RealTime I/O/sec:** The amount of I/O per second to reserve for realtime applications.
- **RealTime I/O MB/sec:** The amount of I/O space per second to reserve for realtime applications.
- **Non-RealTime I/O/sec:** The amount of I/O per second to reserve for non-realtime applications.

- **Non-RealTime I/O MB/sec:** The amount of I/O space per second to reserve for non-realtime applications.
- **RealTime Timeout secs:** The timeout interval to reserve for realtime applications.
- **Label, Type, Size, Serial Number, etc.:** In this section select one or more disk to assign to the file system. (Press and hold **Shift** or **Ctrl** to select multiple disks.)
- **Label:** At the field to the left of the **Label** button, enter a label name. Click **Label** to apply the label name to the selected disks.
- **Unlabel:** Clicking **Unlabel** removes label names from selected disks.
- **Assign:** Click **Assign** to assign selected disks to the file system stripe group.
- **Unassign:** Clicking **Unassign** removes previous associations between disks and the stripe group.

## Edit a File System

### Edit a File System

1. From the **Configuration > File Systems** page, select the file system you wish to edit.
2. Click **Edit**.
3. Select one or both of the following fields:
  - **Storage Manager:** Select this option to enable the file system for StorNext Storage Manager.
  - **Replication/Deduplication:** Select this option to enable data replication/deduplication or deselect it to disable the feature.

---

**i Note:** If a SDISK already exists in the unmanaged file system, the SDISK must be deleted before the file system can be enabled for **Storage Manager** and/or **Replication/Deduplication** feature(s).

4. Modify file system information as desired by clicking the tabs and editing or adding information. (See field descriptions under the headings **Advanced Parameters Tab** and **Stripe Group/Disk Management Fields** below.)
5. Click **Finish** to save changes and return to the **Configuration > File Systems** page, or click **Cancel** to exit without saving.

### Manual Configuration

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

1. When you are finished entering **Advanced Parameters** 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**.

## Advanced Parameters > Allocation Tab

The **Allocation** tab contains fields that affect how resources are allocated on your file system.

- **Journal Size:** Defines the size of the file system journal.
- **Strategy:** Defines the method for choosing stripe groups when allocating disk blocks. Options are **Round**, **Fill**, or **Balance**.
- **Reserved Space:** This option enables delayed allocations on clients. Reserved space is a performance feature that allows clients to perform buffered writes on a file without first obtaining real allocations from the metadata controller. The allocations are performed later during periodic cache synchronization.

If the **Reserved Space** option is not enabled, slightly more disk space will be available, but file fragmentation may increase and performance may not be satisfactory.

- **StripeAlign Size:** Defines the minimum allocation size to trigger automatic stripe-aligned allocations.
- **Inode Stripe Width:** If non-zero, causes large files to have their allocations striped across stripe groups in chunks of the specified size.
- **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. At the first field enter the desired chunk size. At the second field specify the chunk unit of measure (B=bytes, KB=kilobytes, MB=megabytes, GB=gigabytes, TB=terabytes). For more information about the **Allocation Session Reservation** feature, refer to the *StorNext File System Tuning Guide*.
- **Affinity Preference:** If checked, permits files of a particular affinity to have their allocations placed on other available stripe groups (with non-exclusive affinities) when the stripe groups of their assigned affinity do not have sufficient space. Otherwise, allocation attempts will fail with an out-of-space error.
- **Auto Affinities:** This table designates the affinity (stripe group[s]) to which allocations will be targeted for all files on this file system whose name has the specified file extension.
  - To add a new entry to the **Auto Affinities** table, type a file extension (omit the "." dot), select an affinity from the **Affinity** menu, and click **Add**. Use the file extension \* (asterisk) to indicate "all other" file extensions that are not explicitly listed.
  - The **Affinity** menu will list only affinities that are currently assigned to a stripe group of this file system. The affinity **NoAffinity** indicates that allocations will be targeted at stripe groups that have no affinity.
  - To delete one or more entries, simply check the row(s) to delete, and click **Delete**.
  - Each unique file extension can be targeted at only one affinity. However, each affinity may serve as the allocation target for more than one file extension.
  - To sort the auto affinities table by file extension or affinity name, click the **File Extension** or **Affinity**

column title, respectively. Multiple clicks cause the sort order to alternate between ascending and descending alphabetic order.

## Advanced Parameters > Performance Tab

The **Performance** tab fields allow you to adjust parameters for optimized performance.

- **Buffer Cache Size:** Defines the amount of memory used by the FSM process for caching metadata.
- **Inode Cache Size:** This value defines the number of inodes that can be cached in the SNFS server. The default and minimum setting for the cache size is 16.
- **Use Physical Memory Only:** When this option selected, the file system will use only physical memory, not swapped or paged.
- **High Priority FSM:** Determines whether the FSM process should run with real-time priority.

## Advanced Parameters > Debug Tab

The **Debug** tab fields allow you to enable or disable debugging and influence the debug log.

- **Enable Debugging:** Enables detailed file system debug tracing. When debug tracing is enabled, file system performance could be significantly reduced.
- **Debug Log Settings:** Settings to turn on debug functions for the file system server. The log information may be useful if a problem occurs. A Quantum Technical Assistance Center representative may ask for certain debug options to be activated to analyze a file system or hardware problem.
- **Maximum Log Size:** This value defines the maximum number of bytes (size) to which a StorNext Server log file can grow. When the log file reaches the specified size, it is rolled and a new log is started. In this situation, the two log files could use twice the maximum log size space specified in this field. The range is from 1 to 32 megabytes.
- **Maximum Number of Logs:** This value determines the number of rolled logs kept. Choices range from 4 to 64.
- **OP Hang Limit (Seconds):** Defines the time threshold (in seconds) used by the FSM process to discover hung operations.

## Advanced Parameters > Features Tab

The **Features** tab fields allow you to enable or disable various file system-related features.

- **Security Model:** Determines the scheme for specifying and enforcing security policies. The available options are **legacy**, **unixpermbits**, and **acl**. The default value is **legacy**.
  - If the **Security Model** is **legacy**, then the **Unix Id Mapping** field is grayed out (disabled); however, the **Windows Security** option and the **Enforce ACLs** option remain enabled.
  - If the **Security Model** is **acl**, then the **Unix Id Mapping** field is not grayed out; however, the **Windows Security** and the **Enforce ACLs** are grayed out (disabled).



- If the **Security Model** is **acl**, the **Unix Id Mapping** field is not allowed to be **none**. You must select a value from the **Unix Id Mapping** list.
- **Unix Id Mapping**: Determines the Unix Id mapping. The available options are **none**, **algorithmic**, and **winbind**. The default value is **none**.
- **Windows Security**: Determines whether Windows ACLs are enabled for the file system.
- **Enforce ACLS**: Determines whether ACLs are enforced on XSan clients.
- **Windows Global ShareMode**: Determines whether Windows Global ShareMode is enabled for the file system.
- **Quotas**: Determines whether the Quota feature is enabled for the file system. 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.

---

**i Note:** It is not possible to use the StorNext GUI to manipulate the actual quotas, so you must use **snquota** to modify the quotas on individual users. However, you cannot enable or disable quotas entirely through that interface. While user and group quotas cannot be manipulated via the StorNext GUI, directory quotas can be set on the **Manage Quotas** page (for additional information, see [Manage Quotas on page 322](#)).

- **Quota Logs Retention Period**: If **Quotas** is enabled, you can configure the length of time (**Days**, **Weeks**, **Years**) to retain the **Quota** logs. In other words, **Quota** logs are not retained after the length of time determined by the **Quota Logs Retention Period** value.
- **Named Streams**: Determines whether a file system includes support for the Xsan Named Streams feature. Accessing files with Named Streams from a non-Xsan client is not supported. Also, **Named Streams** enabled file systems cannot be configured as Storage Manager, Replication or Deduplication enabled file systems. The **Named Streams** feature enables the storing of additional file system metadata and as such the **Named Streams** feature cannot be disabled after it has been applied to a file system.
- **Spotlight Proxy**: Determines if Spotlight proxy is enabled for the file system. For additional information, see [Configuring Spotlight Proxy on page 35](#).
- **Use Active Directory SFU**: Determines if Active Directory is enabled for the file system.
- **File Locks**: Determines whether the FSM tracks and enforces file locks across all clients.
- **FileLock Resync Timeout**: Defines the timeout for clients re-registering file locks following FSM failover.
- **Restore Journal**: The parameter is used to enable or disable restore journal (also known as metadump or metadata dump) creation by the FSM process. A restore journal logs file system operations and is a key piece of restoring a file system after a disaster on non-managed or managed file systems. By default, restore journal creation is disabled on non-managed file systems, and enabled on managed file systems.

## Advanced Parameters > LDAP Tab

The **LDAP** tab fields allow you to enter parameters related to LDAP (Lightweight Directory Access Protocol, an application protocol for querying and modifying directory services running over TCP/IP).

- **Unix File Creation Mode on Windows:** The number of mode bits for UNIX files
- **Unix Directory Creation Mode on Windows:** The number of mode bits for UNIX directories
- **Unix Nobody UID on Windows:** UNIX user ID to use if no other mapping can be found
- **Unix Nobody GID on Windows:** UNIX group ID to use if no other mapping can be found
- **Unix ID Fabrication on Windows:** Allows you to enable or disable using fabricated IDs on a per-file system basis. If enabled, Windows user IDs are mapped using fabricated IDs.

## Advanced Parameters > Special Tab

The **Special** tab fields allow you to enable or disable special file system-related features.

- **Inode Delete Maximum:** Determines throttling factor for reclaiming space from files allocated using PerfectFit mode.
- **Trim on Close:** Determines whether excess allocations are removed from files when they are closed.
  - **Global Super User:** Enable this option (check the box) to allow a user with super-user privileges to assert these privileges on the file system.
    - If the **Global Super User** option is *enabled*, super users have global access rights on the file system. This selection is the same as the `maproot=0` directive in the Network File System (NFS).
    - If the **Global Super User** option is *not enabled*, super users can modify only files they can access, like any other users.
- **File System Capacity Threshold:** Defines the file system fill level (in percent) that triggers a RAS event.
- **Extent Count Threshold:** Defines the number of extents in a file required to trigger a fragmentation RAS event.
- **Remote Notification:** Determines whether to enable partial support for cluster-wide Windows directory event notification.

## Stripe Groups/Disk Management Fields

To modify an existing stripe group, under the **Stripe Groups** heading select the stripe group you want to modify, and then change its properties as desired.

To add a new stripe group to the file system, click **Add** and then enter the remaining fields for the new stripe group.

When you are finished on the **Stripe Group** tab, click **Apply** to save your changes, or **Cancel** to abandon your changes.

- **Stripe Group:** Select the stripe group you want to modify or delete.
- **Add:** Click this button to add a new stripe group.

---

 **Note:** If the following indented fields are not displayed, they appear after you click **Add**. Likewise,

---

**i** after you delete a stripe group these fields may not be displayed.

- **Name:** Enter a name for the new stripe group, or skip this field to accept the displayed name.
- **Breadth:** Specify the stripe group breadth, which is the number of kilobytes (KB) that is read from or written to each disk in the stripe group.
- **Content:** Specify whether the stripe group will be used for metadata, journaling, or user data. You can specify one, two, or all content types.
- **Delete:** Click this button to delete the currently selected stripe group.

---

**i Note:** This particular delete function does not provide a confirmation message, so be absolutely sure you want to delete the selected stripe group before you click **Delete**. The selected stripe group will be deleted immediately after you click **Delete**, and there is no Undo feature.

- **Affinity:** An affinity is a label that allows you to control the physical location of files. Managed file systems are restricted to only using two affinities and are used for Disk-to-Disk relocation Storage Manager policies. Managed file systems are limited to using the same two affinity names on all managed file systems. By default, the affinity names on a managed file systems are “Tier1” and “Tier2”. For unmanaged file systems, using affinities is a two step process: first, each stripe group can be assigned one or more affinities during file system configuration, then associate a directory with the affinity. For example, if you configure stripe group SG2 to have affinity AFF2 and then associate directory special\_files with affinity AFF2, all files put into special\_files can exist only on the disks that make up SG2. If an affinity were not assigned and associated with special\_files, the files put into that directory could exist on any stripe group or any disk in the file system. It makes sense to use affinities in environments where performance is critical. For example, you might want to constrain video files to a stripe group made of fibre channel disks tuned for video playback, but have audio files reside on slower SCSI disks in a different stripe group. A stripe group can have multiple affinities.

You may not remove an affinity from a stripe group if that affinity is not assigned to any remaining stripe group AND the Auto Affinities table includes a file extension that is targeted at that affinity. You must first update the Auto Affinities table to delete that auto affinity entry. See the [Auto Affinities](#) section for auto affinity delete instructions.

If you want to associate an affinity to the new stripe group, select the desired affinities.

- **Exclusive:** When this option is enabled, the selected stripe group is used exclusively for the affinity's files.
- **Access:** Specify the permission level for the stripe group: Full R/W (read/write), Read Only, or Disabled.
- **Quality of Service:** Specify parameters for the Quality of Service (QOS) feature. QOS allows real-time applications to reserve a specified amount of bandwidth on the storage system. (For more information about the QOS feature, see the StorNext Installation Guide.)
- **RealTime I/O/sec:** The amount of I/O per second to reserve for realtime applications.
- **RealTime I/O MB/sec:** The amount of I/O space per second to reserve for realtime applications.
- **Non-RealTime I/O/sec:** The amount of I/O per second to reserve for non-realtime applications.

- **Non-RealTime I/O MB/sec:** The amount of I/O space per second to reserve for non-realtime applications.
- **RealTime Timeout secs:** The timeout interval to reserve for realtime applications.
- **Label, Type, Size, Serial Number, etc.:** In this section select one or more disk to assign to the file system. (Press and hold **Shift** or **Ctrl** to select multiple disks.)
- **Label:** At the field to the left of the **Label** button, enter a label name. Click **Label** to apply the label name to the selected disks.
- **Unlabel:** Clicking **Unlabel** removes label names from selected disks.
- **Assign:** Click **Assign** to assign selected disks to the file system stripe group.
- **Unassign:** Clicking **Unassign** removes previous associations between disks and the stripe group.

## Delete a File System

### Delete a File System

1. On the **Configuration > File Systems** page, select the file system you wish to delete.
2. Click **Delete**.
3. Click **Yes** to confirm the deletion, or **No** to cancel.

## Perform File System Actions

After you have created at least one file system, you can perform multiple file system-related actions:

Action	Description
<b>Stop</b>	Stops the file system.
<b>Start</b>	Starts the file system.
<b>Start and Activate</b>	Both starts and activates the file system in one step, saving you the time of starting and activating separately.
<b>Start and Mount</b>	Both starts and mounts the file system in one step, saving you the time of starting and mounting separately.
<b>Activate</b>	Activates the file system.
<b>Mount</b>	Mounts the file system.
<b>Unmount</b>	Unmounts the file system.

Action	Description
<b>Make</b>	Makes the file system.
<b>Update</b>	Applies configuration changes to the file system. <b>i Note:</b> 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.
<b>Check</b>	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.
<b>Expand</b>	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. <b>i Note:</b> 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. 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. <b>i Note:</b> For Window MDCs and Linux MDCs that do not have the StorNex GUI functionality available, see <b>Using Resource Allocation From the Command Line</b> in the <i>StorNext 5 User's Guide</i> .
<b>Migrate</b>	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.

## Perform an Action on a File System

1. Select the desired file system.
2. Select the desired action from the Select **Action** list.
  - The **Check**, **Expand** and **Migrate** actions take you to a separate page and require additional steps to complete the action. (Additional steps are described in the sections below.)
3. When a message informs you whether the action was completed successfully, click **OK** to continue.

## Check a File System

You should perform a file system check whenever you migrate the file system. This operation could take a significant amount of time depending on the size of the file system, so plan accordingly.

After you select the Check File System action, the **Tools > File Systems > Check File System** page appears. This page displays all file systems available for checking. (You can check only one file system at a time.) To initiate a new check from this page, select the desired file system and then click **Check**.

At the bottom of the page, information about previously run checks is displayed, including the file system names, start and end times, and check statuses.

- When you initiate a new check, you can monitor the status in this same area of the page.
- To view a detailed report about one of these checks, select the desired check and then click **Report**.
- To delete a previously run check, select the check you want to delete and then click **Delete**.
- To delete all previously run checks, click **Delete All**.

When are finished on the **Tools > File Systems > Check File System** page, click **Done**.

## Expand a File System

This function allows you to name and assign disks to the file system. The procedure for naming and assigning disks is identical to when you first created the file system. For more information, see [Add a File System on page 38](#).

## Migrate a File System

This feature enables you to migrate file system data to tertiary storage. Depending on the size of the file system, this process could take some time to complete, so plan accordingly.

After you select this action, the **Tools > File Systems > Migrate** page appears. Follow these steps to begin file system migration:

1. At the top of the page, select the file system whose data you want to migrate.
2. At the bottom of the page, select the stripe group to which you can migrate data.
3. Click **Migrate**.
4. When you are finished on the **Tools > File Systems > Migrate** page, click **Done**.

---

# Storage Destinations Overview

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.

See the following topics for more information:

- [Setting Up Libraries on the next page](#)
- [Setting Up Storage Disks on page 57](#)
- [Setting Up Object Storage Destinations on page 443](#)
- [Setting Up Q-Cloud on page 481](#)

- [Configure Data Replication on page 95](#)
- [Configure Data Deduplication on page 97](#)

## Setting Up Libraries

The **Libraries** tab on the **Configuration > Storage Destinations > Library** page enables you to perform actions pertaining to libraries, including adding, editing, and deleting a library.

### View a Library

1. On the **Configuration** menu, click **Storage Destinations**, and then click the **Library** tab.
2. Select the library whose information you want to view.
3. Click **View**, or select **View** from the actions list. The following library information appears:

Parameter	Description
<b>Name</b>	The name of the library.
<b>Type</b>	The type of library. For example, <b>SCSI</b> , <b>ACSL</b> , etc.
<b>EIF Port Config</b>	The current EIF port configuration. EIF stands for Enterprise Instrumentation Framework, and it helps StorNext process data from applications on the server.
<b>Media Count</b>	The number of media in the library.
<b>Serial Number</b>	The library's serial number.
<b>State</b>	The library's current state. For example, <b>Online</b> or <b>Offline</b> .
<b>HA Failover</b>	Indicates whether HA failover is enabled for the library.
<b>Fill Level</b>	The library's current fill level.
<b>Dual Aisle</b>	Indicated whether the library has a dual aisle configuration.
<b>Drive Count</b>	The number of tape drives in the library.

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

### Add a New Library

1. On the **Configuration** menu, click **Storage Destinations**, and then click the **Library** tab.

2. Click **New**.
3. Enter the following fields at the top of the page:
  - **Name**: Enter the name for the new library
  - **Type**: Select the appropriate library type: ACSLS, DAS, DASDual Aisle, SCSI, or Vault
  - **Remote Server (ACSLs, DAS and DAS Dual Aisle only)**: Enter the IP address or host name of the DAS server.
  - **Remote Client (DAS and DAS Dual Aisle only)**: Enter the client name of the corresponding DAS server host.
  - **Archive**: Select the archive for the new library, or click **Scan** to have StorNext discover available archives for you. (Scanning could take a while to complete depending on your configuration)
4. In the **Drives** section, select a tape drive to add to your new library, or click **Scan** to have StorNext discover available drives for you.
5. In the **Media** section, click **Scan** to have StorNext discover available media for you.
6. Click **Apply**.
7. After a message informs you that the library was successfully created, click **OK**.
8. **(Optional)** Repeat **Step 2** through **Step 7** to add additional libraries.

## Edit a Library


1. On the **Configuration** menu, click **Storage Destinations**, and then click the **Library** tab.
2. Select the library you want to edit.
3. Click **Edit**, or click **Edit** from the actions list. After you select this option StorNext scans the library, which could take some time to complete depending on your configuration.
4. Select the tape drives you want included in your library, or click **All** to include all available tape drives. To exclude all drives, click **None**.
5. Click **Apply** to save your changes, or click **Cancel** to exit without saving..
6. When a 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**.
8. **(Optional)** Click **Remove Drives** to remove the selected drives from the library.
9. **(Optional)** If you have increased (or decreased) the available media slots in a Tape Library, then Storage Manager must be updated to match the settings of the physical tape library. Click **Validate Slot Count** to trigger a background job that will query the physical library and set the Storage Manager slot count to the maximum value allowed by the physical archive for each configured media type. A notification appears instructing you to review a specified ID in the **Reports > Jobs** page.

---

**i Note:** The command `vsarchiveconfig` can be used with the `-m` option to modify the number of slots available for a specific media. See the `vsarchiveconfig` manpage in the *StorNext MAN Pages Reference Guide* for more detailed information.




## Delete a Library

 **WARNING:** Before you can use this function, all media must first be deleted from the library. All the media in the library and associated data will be deleted from the system. The data will be lost and the operation cannot be undone.

1. On the **Configuration** menu, click **Storage Destinations**, and then click the **Library** tab.
2. Select the library you want to delete.
3. Click **Delete**, or select **Delete** from the actions 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**.

## Perform Other Library Actions

Towards the middle of the **Configuration > Storage Destinations > Library** page is a 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** list:

 **Note:** If you do not select a library, the actions listed below run on all configured libraries in the list. If a library is selected, the actions listed below run on the selected library:

- **Drives Validation Report**
- **Drives Validation Update**
- **Drive Replacement**
- **Drive Device Path Audit**

Parameter	Description
<b>Audit</b>	Select this option to perform an inventory of the media 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.
<b>Remap-Audit</b>	This option synchronizes the StorNext databases with the library databases.

Parameter	Description
<b>Validate Slot Count</b>	<p>If you have increased (or decreased) the available media slots in a Tape Library, then Storage Manager must be updated to match the settings of the physical tape library.</p> <p>This operation will trigger a background job that will query the physical library and set the Storage Manager slot count to the maximum value allowed by the physical archive for each configured media type.</p> <p>A notification appears instructing you to review a specified ID in the <b>Reports &gt; Jobs</b> page.</p> <p><b>i Note:</b> The command <code>vsarchiveconfig</code> can be used with the <code>-m</code> option to modify the number of slots available for a specific media. See the <code>vsarchiveconfig</code> manpage in the <i>StorNext MAN Pages Reference Guide</i> for more detailed information.</p>
<b>Online</b>	Select this option to set the library online.
<b>Offline</b>	Select this option to take the library offline.
<b>Drives Online</b>	Select this option to place the drives in the library online.
<b>Drives Offline</b>	Select this option to take the drives in the library offline.
<b>Drives Validation Report</b>	Select this option to execute a drive validation report. A notification appears on the banner instructing you to review a specified <b>ID</b> in the <b>Reports &gt; Jobs</b> page.
<b>Drives Validation Update</b>	Select this option to execute a drive validation update. A notification appears on the banner instructing you to review a specified <b>ID</b> in the <b>Reports &gt; Jobs</b> page.
<b>Drive Replacement</b>	Select this option to update the drive serial number mappings. This option is also available when you click the <b>Tools</b> menu, click <b>Storage Manager</b> , and then click <b>Drive Replacement</b> .
<b>Drive Device Path Audit</b>	Select this option to verify that the device path seen by the operating system matches the Storage Manager device path for a configured tape drive. If there is a mismatch, Storage Manager is updated.
<b>Add Media Bulkload</b>	Select this option to add media to the library via the bulk loading method.
<b>Add Media Mailbox</b>	Select this option to add media to the library through the library's mailbox.

## Setting Up Storage Disks

The **Storage Disk** option enables you to view, add, edit or delete storage disks.

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.

---

**i Note:** Before you create a storage disk, the disks you plan to use must reside in an existing, mounted file system.

## Storage Disk Usage Recommendations

Use Case	Recommendation
File System configuration includes storage disks.	Avoid using that file system for any data other than storage disk stored data. <b>i Note:</b> If your file system includes storage disks and you accidentally fill it with unrelated user data (for instance, non-storage disk data) call the Quantum Technical Assistance Center and request the procedure to clean up and transcribe data.
Configuration includes Shared File Systems	Use complete and physically dedicated file systems (snfs, local, NFS, or other,) for storage disk data. <b>i Note:</b> Avoid shared file systems or file systems with linked directories.

---

**! 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.

## View Storage Disks

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Storage Disk** tab. Information for any previously configured storage disks is shown. For each configured storage disk, the page displays 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.
3. Select the storage disk whose information you want to view.
4. Click **View**.
5. When you are finished viewing library information, click **Done**.

## Add a New Storage Disk

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Storage Disk** tab.
3. Click **New**.
4. Enter the following fields:
  - **Storage Disk Name:** The name of the storage disk you are creating.
  - **Number of Streams:** The number of I/O streams that can concurrently write to the disk. The default is 4 streams.
  - **Copy # for policy classes:** The copy number (1-4) assigned to the storage disk.
  - **Mount Point:** The file system mount point for the storage disk. Select an existing mount point from the dropdown list.
  - **Directory:** The directory selected for file storage. Select an existing directory from the dropdown list, or create a new directory by typing the new directory name and then clicking Create Directory.
5. Click **Apply** to save your changes, or **Cancel** to exit without saving.
6. Repeat **Step 2** through **Step 4** to add additional storage disks.

## Edit a Storage Disk

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Storage Disk** tab.
3. Select the storage disk whose information you want to edit.
4. Click **Edit**.
5. Modify any of the fields you entered when creating the storage disk.
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 storage disk was successfully modified, click **OK**.

## Delete a Storage Disk

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Storage Disk** tab.
3. Select the storage disk 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**.

## Set a Storage Disk Online or Offline

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Storage Disk** tab.
3. Select the storage disk.
4. In the **Select Action** list, click one of the following:
  - **Online** to set the storage disk online.
  - **Offline** to set the storage disk offline.

## Setting Up Object Storage Destinations

Just as with tape and storage disk, you can configure object storage destinations using the **Storage Destinations > Object Storage** option under the **Configuration** menu. See [Configuring Object Storage on page 438](#).

Depending on the provider and media type you are configuring, follow the appropriate procedure below to configure your Object Storage Destination:

- [Setting up Lattus Object Storage Destinations on a StorNext Configuration below](#)
- [Setting up S3COMPAT Object Storage Destinations on a StorNext Configuration on page 67](#)
- [Setting up AWS Object Storage Destinations on a StorNext configuration on page 72](#)
- [Setting up Azure Object Storage Destinations on a StorNext Configuration on page 85](#)

For information on troubleshooting Object Storage and Cloud errors, see [Debugging StorNext for Object Storage Systems and Cloud Providers on page 638](#).

## Setting up Lattus Object Storage Destinations on a StorNext Configuration

### View Lattus Object Storage Destinations

Follow this procedure to view a list of currently configured Lattus Object Storage destinations.

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab. Information for any previously configured Lattus Object Storage destination is shown as entries that have **Quantum Lattus** listed as the **Provider**. For each configured destination, the page displays the **Name**, **Provider**, **Appliance 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 destination information, click **Done**.

## Add a New Lattus Object Storage Destination

Follow this procedure to add a new Lattus Object Storage destination.

**i 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 291](#). To manage SSL certificates, click **Object Storage Certificates** on the **Tools** menu. For additional information, see [Basic Secure Sockets Layer \(SSL\) Guidelines on page 497](#) in the *StorNext 5 User's Guide*, as it outlines some standard information about using private and public certificates.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Click **New...** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue...**
4. Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new Lattus Object Storage destination
<b>Provider</b>	Select <b>Quantum Lattus</b> from the <b>Provider</b> list.
<b>Manager Host</b>	Enter the host address for the Lattus Object Storage manager host.
<b>Manager Port</b>	Enter a decimal integer to specify the port number of the Lattus Object Storage manager GUI interface. The default port number is <b>80</b> .
<b>Manager Protocol</b>	Select the <b>http</b> or <b>https</b> protocol.  <b>i Note:</b> If you plan to use HTTPS, you must <b>create</b> or <b>import</b> a security certificate prior to creating a <b>Lattus Object Storage Destination</b> . This applies to each Lattus controller configured for HTTPS. If you only intend to use HTTP, certificates are not needed. To <b>create</b> or <b>import</b> a Lattus Object Storage security certificate, see <a href="#">Lattus Certificates on page 291</a> . To manage SSL certificates, click <b>Object Storage Certificates</b> on the <b>Tools</b> menu. For additional information, see <a href="#">Basic Secure Sockets Layer (SSL) Guidelines on page 497</a> in the <i>StorNext 5 User's Guide</i> , as it outlines some standard information about using private and public certificates.
<b>Authentication</b>	Select if authentication is required for this configuration.

Parameter	Description
<b>User Name</b>	Select a global user name to be used for namespace permission for this configuration. This parameter is mandatory if <b>Authentication</b> is set to “ <b>Enabled</b> ”.
<b>Password</b>	Select a global password to be used for namespace permissions for this configuration. This parameter is mandatory if <b>Authentication</b> is set to “ <b>Enabled</b> ”.

5. In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

6. In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the “ <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> ” error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.
<b>Media Type</b> ( <i>New in StorNext 5 release 5.4</i> )	This parameter specifies the object storage media type assigned to an I/O Path that is associated with a specific Object Storage API. The available values for provider <b>Quantum Lattus</b> are <b>AXR</b> and <b>S3</b> .
<b>URL Style</b>	There are two ways to format the URL: <ul style="list-style-type: none"><li>• PATH</li><li>• VHOST</li></ul> This parameter defines which style of URL to use.

Parameter	Description
<b>Object Access Protocol</b>	<p>Select the protocol to be used for Object Storage object access. By default, the protocol is set to <b>http</b>.</p> <hr/> <p><b>i</b> <b>Note:</b> If you plan to use HTTPS, you must <b>create</b> or <b>import</b> a security certificate prior to creating a <b>Lattus Object Storage Destination</b>. This applies to each Lattus controller configured for HTTPS. If you only intend to use HTTP, certificates are not needed. To <b>create</b> or <b>import</b> a Lattus Object Storage security certificate, see <a href="#">Lattus Certificates on page 291</a>. To manage SSL certificates, click <b>Object Storage Certificates</b> on the <b>Tools</b> menu. For additional information, see <a href="#">Basic Secure Sockets Layer (SSL) Guidelines on page 497</a> in the <i>StorNext 5 User's Guide</i>, as it outlines some standard information about using private and public certificates.</p>
<b>Host [Port]</b>	<p>Enter the connection endpoint address that contains the 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.</p>

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 a user name and password are required, either use the credential specified for the manager host or check the **Use different credentials** box and enter the username 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** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Select ( <b>Scan</b> mode) or enter ( <b>Manual</b> mode) the appropriate container for this configuration.
<b>Media ID</b>	Enter the StorNext Media ID associated with the selected container. The <b>Media ID</b> must be unique.
<b>Media Type</b> <i>(New in StorNext 5 release 5.4)</i>	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. The available values for provider <b>Quantum Lattus</b> are <b>AXR</b> and <b>S3</b> .
<b>Storage Class</b>	This parameter should be left at <b>none</b> as it is not applicable to Lattus media.
<b>Signing Type</b>	For Lattus S3 media, use the default value of <b>V2</b> . This parameter is not applicable to AXR media.



Parameter	Description
<b>Authentication Type</b>	This parameter specifies the authentication type for the container being configured. An authentication type is required for Lattus S3 media but not for AXR. The available values for provider <b>Quantum Lattus</b> are <b>NONE</b> and <b>STANDARD</b> . The <b>STANDARD</b> type authenticates with a user name and password for Object Storage access.
<b>User Name</b>	Enter a user name to be used to access this container. This parameter is mandatory if Authentication is set to " <b>Enabled</b> ". This selection overrides the global permissions settings.
<b>Password</b>	Enter a password to be used to access the container. This parameter is mandatory if Authentication is set to " <b>Enabled</b> ". This selection overrides the global permissions settings.
<b>Copy Number</b>	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.
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details

**i 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. **(Optional)** Repeat **Step 7** to add additional containers to the same Lattus Object Storage Destination.
9. Click **Apply** to save your changes, or **Cancel** to exit without saving.
10. **(Optional)** Repeat **Step 3** through **Step 9** to add additional Lattus Object Storage destinations.

**i Note:** The containers on Lattus-M share the same I/O paths and storage capacity. There is no advantage to be gained by defining multiple containers for the same policy class and copy number. Storage Manager selects the first available container that meets the policy class criteria for the store operation.

## Edit a Lattus Object Storage Destination

Follow this procedure to edit an existing Lattus Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **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.

**i 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**.

## Delete a Lattus Object Storage Destination

Follow this procedure to delete an existing Lattus Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **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 Object Storage destination was successfully deleted, click **OK**.

## Perform 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, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the Lattus Object Storage destination.
4. Click **Launch Manager**. A new browser window appears, and displays the Object Storage Manager GUI application login page. 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.

---

**i Note:** If you are using Safari as your browser, you may have to enable pop-ups. See [Enable Pop-ups in Safari below](#).

## 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.

## 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 ([Change the Current State of Object Storage Destinations, Controllers, and I/O Paths on page 421](#)). 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.

## Converting an AXR Namespace to an S3 Bucket

With Lattus 3.5.1, it is possible to convert AXR namespaces to S3 buckets and make them accessible via the S3 interface. StorNext 5 provides the capability to convert the media type from AXR to S3.

## Convert One or More AXR Namespaces to S3 Buckets

1. Using the StorNext MDC (CLI only).

- 
- i Note:** Ensure there is no store/retrieve operation occurring on the same AXR namespace which you plan to convert. If there is store/retrieve operation occurring on a particular namespace, wait for the operation to complete.
    - a. Execute the command **fsubjcfg** and retrieve the **media ID** of the AXR namespace to be converted.
    - b. Stop TSM.
  - 2. Refer to the documentation in the section titled **Converting an AXR Namespace to an S3 Bucket** located within the *Lattus Service Reference Guide* (Part Number 6-67798-07). Log into the Lattus controller to perform the conversion.
  - 3. After the Lattus conversion is complete from **Step 2**, execute the following commands using the StorNext MDC CLI:
    - a. Add a S3 connection endpoint via the StorNext GUI or the **fsubjcfg** command line. At least one S3 iopath is needed to access these S3 buckets under the same appliance:

```
#/usr/adic/TSM/exec/fsobjcfg -a -o iopath_alias -i connection_endpoint -e  
http -t S3 -n controller_node_alias
```

- b. Change the AXR namespace name and media type:

```
# fsobjcfg -m -b NameSpace -t S3 -U <S3_bucket_username> -P <S3_bucket_  
password> -X -f <Media-ID>
```

using the NameSpace value from **Step 2** and the Media-ID value from **Step 1a**.

- c. Modify the file `/usr/adic/TSM/config/filesize.config` and change **LATTUS** to **S3** for the **media ID** corresponding to the converted namespace. Alternatively, you can achieve the same result using the StorNext GUI by changing the fspolicy's steering information from Lattus to S3.

---

**i Note:** To use the StorNext GUI, TSM should be available.

- d. Start TSM.
- e. Verify the store and retrieve are working as expected with converted media type.
- f. Use the command `fsfileinfo -u` to verify the object URL reflects (S3) for every file stored before or after the conversion.

## Setting up S3COMPAT Object Storage Destinations on a StorNext Configuration

Storage Manager supports a rich set of features for AWS S3 compatible cloud storage providers including:

- HTTP and HTTPS access to AWS S3 compatible buckets
- AWS Signature Version 2 (V2) and Version 4 (V4)
- AWS Standard authentication that makes use of either your AWS Identity and Access Management (IAM) Access Key Id and Secret Access Key or your user name and password

The following S3 compatible providers are supported by Storage Manager for the S3COMPAT media type:

- **IBM Cleversafe**
- **StorageGRID Webscale**
- **Scality RING**

When you configure the Object Storage Destination for one of these providers, select the provider that matches your actual Object Storage system.

## View S3COMPAT Object Storage Destinations

Follow this procedure to view a list of currently configured S3COMPAT Object Storage destinations.

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab. Providers using the **S3COMPAT** media type are listed along with other object storage providers in the **Object Storage** tab. For each configured destination, the screen displays the **Name**, **Provider**, **Appliance 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 **S3COMPAT** Object Storage destination whose information you want to view.
4. Click **View...**
5. When you are finished viewing the destination information, click **Done**.

## Add a New S3COMPAT Object Storage Destination

Follow this procedure to add a new **S3COMPAT** Object Storage destination.

1. Ensure that the S3 buckets that you are configuring with Storage Manager, have been created on your Object Storage system and that you know the names, connection endpoint, Access Key Id, and Secret Access Key to your buckets.

**i Note:** Note: If you plan to use **HTTPS**, you may have to create or import a security certificate prior to creating a **S3COMPAT** Object Storage destination. Follow the documentation from your S3 Object Storage system's vendor to set this up on your Object Storage system. Also, see [HTTPS Default CA ROOT Certificate File or Path on page 292](#) for additional information on how to configure your customized CA PEM files on your StorNext system.

2. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
3. Click the **Object Storage** tab.
4. Click **New...** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue...**
5. Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new S3COMPAT Object Storage destination
<b>Provider</b>	Select one of the <b>S3COMPAT</b> providers from the drop-down list that matches your Object Storage system. <ul style="list-style-type: none"><li>• <b>IBM Cleversafe</b></li><li>• <b>StorageGRID Webscale</b></li><li>• <b>Scality RING</b></li></ul>
<b>Manager Host</b>	Enter the host address for the <b>S3COMPAT</b> Object Storage manager host.

Parameter	Description
<b>Manager Port</b>	Enter a decimal integer to specify the port number of the <b>S3COMPAT</b> Object Storage Host's port. The default port number is <b>80</b> .
<b>Manager Protocol</b>	Select the <b>HTTP</b> or <b>HTTPS</b> protocol.  <b>Note:</b> Note: If you plan to use <b>HTTPS</b> , you may have to create or import a security certificate prior to creating a <b>S3COMPAT</b> Object Storage destination. Follow the documentation from your S3 Object Storage system's vendor to set this up on your Object Storage system. Also, see <a href="#">HTTPS Default CA ROOT Certificate File or Path on page 292</a> for additional information on how to configure your customized CA PEM files on your StorNext system.
<b>Authentication</b>	This parameter <b>does not</b> apply to the <b>S3COMPAT</b> Object Storage Destination. Leave it at the default <b>Disabled</b> .
<b>User Name</b>	This parameter <b>does not</b> apply to the <b>S3COMPAT</b> Object Storage Destination. Leave it at the default blank.
<b>Password</b>	This parameter <b>does not</b> apply to the <b>S3COMPAT</b> Object Storage Destination. Leave it at the default blank.

6. In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

7. In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

**Note:** Using the Host name to configure the **I/O Path** for the provider **Scality RING**.

If you are not using an IP address as an endpoint to Scality, the default Host name endpoint for Scality is to emulate Amazon AWS S3 (for example, **s3.amazonaws.com**). This configuration is provided in the `config.json` file on the Scality host. Using any other host name as an endpoint (for example, **my-scality.host.com**) does not work even if the name resolves to the correct IP address. Scality rejects the request with the error message `HTTP/1.1 400 Bad Request`. If you want to use the hostname configured by default on Scality, you can configure your server to use **s3.amazonaws.com** as the endpoint. However, ensure that the name resolves to the IP address of your Scality host.

If your Scality host's IP address is `10.65.191.2`, you can resolve **s3.amazonaws.com** by having the following entry in your `/etc/hosts` file:

```
10.65.191.2 s3.amazonaws.com
```

If you want to use a DNS name that is not the default on Scality, modify the `config.json` file on the scality host:

```
# cat config.json | grep my-scality  
"localregion": ["my-scality.host.com"]
```

And have an entry in your `/etc/hosts` file that resolves it correctly:

```
10.65.191.2 my-scality.host.com
```

Quantum recommends you consult your Scality vendor for further guidance.

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the " <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> " error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.
<b>Media Type</b> ( <i>New in StorNext 5 release 5.4</i> )	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. Select <b>S3COMPAT</b> from the drop-down list.
<b>URL Style</b>	There are two ways to format the URL: <ul style="list-style-type: none"><li>• PATH</li><li>• VHOST</li></ul> This parameter defines which style of URL to use.
<b>Object Access Protocol</b>	Select the protocol to be used for <b>S3COMPAT</b> Object Storage object access. By default, the protocol is set to <b>http</b> .
<b>Host [:Port]</b>	Enter the connection endpoint address that contains the 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 (for additional information, see <a href="#">Using the Host name to configure the I/O Path for the provider Scality RING. on the previous page</a> )

8. 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**, check the **Use different credentials** box and enter the username (the **Access Key ID**) and password (the **Secret Access Key**). You are then presented with a pre-populated list of available containers when you add a container. 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** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Select ( <b>Scan</b> mode), or enter ( <b>Manual</b> mode) the name of your S3 bucket.
<b>Media ID</b>	Enter the StorNext <b>Media ID</b> associated with the selected container. The <b>Media ID</b> must be unique.
<b>Media Type</b>	Select <b>S3COMPAT</b> from the drop-down list.
<b>Storage Class</b>	This parameter should be left at <b>none</b> as it is not applicable to <b>S3COMPAT</b> media
<b>Signing Type</b>	<p>This parameter specifies the signing type for the requests sent to the <b>S3COMPAT</b> Object Storage server. Available values include <b>V2</b> and <b>V4</b>. In order to use <b>V4</b>, the server must support <b>V4</b> signing for both AWS full payload and chunked uploading.</p> <p><b>i Note:</b> When configuring the signing type for containers from the provider <b>Scality RING</b>, set this parameter at <b>V2</b>, until Scality supports <b>V4</b> chunked uploading.</p>
<b>Authentication Type</b>	This parameter specifies the authentication type for the container being configured. An authentication type is required for all <b>S3COMPAT</b> media. Use the default value of <b>STANDARD</b> , which authenticates with an <b>Access Key ID</b> and <b>Secret Access Key</b> .
<b>User Name</b>	Enter the <b>Access Key Id</b> for this container.
<b>Password</b>	Enter the <b>Secret Access Key</b> for this container.
<b>Copy Number</b>	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.
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details.

**i Note:** If no data has been written to a controller, I/O path, or container, click **Delete** to remove the item, and then click **Apply** to save the changes.



9. **(Optional)** Repeat **Step 8** to add more containers to the same **S3COMPAT** Object Storage Destination.
10. Click **Apply** to save your changes, or **Cancel** to exit without saving.
11. **(Optional)** Repeat **Step 4** through **Step 10** to add additional **S3COMPAT** Object Storage Destinations.

## Edit a S3COMPAT Object Storage Destination

Follow this procedure to edit an existing **S3COMPAT** Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the **S3COMPAT** 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.

---

**i 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 **S3COMPAT** Object Storage destination was successfully modified, click **OK**.

## Delete a S3COMPAT Object Storage Destination

Follow this procedure to delete an existing **S3COMPAT** Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the **S3COMPAT** 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 Object Storage destination was successfully deleted, click **OK**.

## Setting up AWS Object Storage Destinations on a StorNext configuration

Storage Manager supports a rich set of features for Amazon Web Service (AWS) Simple Storage Service (S3) Cloud Storage including:

- HTTP and HTTPS access to AWS S3 buckets, though HTTPS is the recommended protocol for access AWS S3 buckets
- AWS Signature Version 2 (V2) and Version 4 (V4)
- Multiple AWS authentication types including:
  - AWS Standard authentication that makes use of your AWS Identity and Access Management (IAM) Access Key Id and Secret Access Key, for AWS public cloud and GovCloud
  - AWS Security Token Service (STS) authentication for AWS public cloud and GovCloud
  - AWS Commercial Cloud Service (C2S) Access Portal (CAP) authentication for AWS FedCloud
- Multiple AWS storage classes including **standard**, **standard\_ia**, and **glacier**

## View AWS Object Storage Destinations

Follow this procedure to view a list of currently configured **AWS** Object Storage destinations.

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab. Information for any previously configured **AWS** Object Storage destinations is shown as entries that have **AWS** listed as the **Provider**. For each configured destination, the screen displays the **Name**, **Provider**, **Appliance 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 **AWS** Object Storage destination whose information you want to view.
4. Click **View...**
5. When you are finished viewing the destination information, click **Done**.

## Add a New AWS Object Storage Destination that uses STANDARD Authentication

Prior to configuring the AWS Object Storage Destinations with Storage Manager for AWS STANDARD authentication, complete the following two steps, which apply to both AWS public cloud and GovCloud:

1. Have your AWS IAM Access Key Id and Secret Access Key ready.
2. Ensure that the S3 buckets that you are configuring with Storage Manager have been created with AWS and that you know the names and the AWS region endpoint of your buckets.

See <http://docs.aws.amazon.com/general/latest/gr/rande.html> for details on AWS Regions and Endpoints for AWS public cloud.

See <http://docs.aws.amazon.com/govcloud-us/latest/UserGuide/using-govcloud-endpoints.html> for details on AWS GovCloud (US) Endpoints.

Follow this procedure to add a new **AWS** Object Storage destination that uses STANDARD authentication.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Click **New....** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue....**
4. Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new <b>AWS</b> Object Storage destination
<b>Provider</b>	Select <b>AWS</b> from the <b>Provider</b> list.
<b>Manager Host</b>	Enter <b>s3.amazonaws.com</b> for the <b>AWS</b> Object Storage Manager Host's address.
<b>Manager Port</b>	Enter <b>443</b> for the <b>AWS</b> Object Storage Manager Host's port.
<b>Manager Protocol</b>	Select <b>HTTPS</b> .
<b>Authentication Type</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default <b>Disabled</b> .
<b>User Name</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.
<b>Password</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.

5. In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

6. In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the " <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> " error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.

Parameter	Description
<b>Media Type</b> <i>(New in StorNext 5 release 5.4)</i>	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. Select <b>AWS</b> from the drop-down list.
<b>URL Style</b>	Select <b>VHOST</b> from the drop-down list.
<b>Object Access Protocol</b>	This specifies the network protocol to be used for the host. Select <b>HTTPS</b> from the drop-down list.
<b>Host[:Port]</b>	Enter the AWS region endpoint for your S3 bucket. For example: <ul style="list-style-type: none"><li>• <b>s3-us-west-2.amazonaws.com</b> for an S3 bucket created in the US West (Oregon) region</li><li>• <b>s3.amazonaws.com</b> for an S3 bucket created in the US East (N. Virginia) region</li><li>• <b>s3-us-gov-west-1.amazonaws.com</b> for an S3 bucket created in AWS GovCloud</li></ul>

7. In the **Containers** section, perform the following:
  - a. Leave the **Container Selection** at the default **Manual**.
  - b. Click **Add** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Enter the name of your S3 bucket.
<b>Media ID</b>	Enter the StorNext <b>Media ID</b> associated with the selected container. The <b>Media ID</b> must be unique.
<b>Media Type</b>	Select <b>AWS</b> from the drop-down list.
<b>Storage Class</b>	This parameter specifies the storage class of your S3 bucket's lifecycle property. Make sure that this value always matches that configured for the S3 bucket itself. Available values include <b>standard</b> , <b>standard_ia</b> , and <b>glacier</b> . Most configurations use the default <b>standard</b> storage class.
<b>Signing Type</b>	This parameter specifies the signing type for the requests sent to the AWS Object Storage server. Available values include <b>V2</b> and <b>V4</b> (default).

Parameter	Description
<b>Authentication Type</b>	<p>This parameter specifies the authentication type for the container being configured. Available values include:</p> <ul style="list-style-type: none"><li>• <b>STANDARD</b></li><li>• <b>STS_PUBLIC</b></li><li>• <b>STS_GOVCLOUD</b></li><li>• <b>CAP</b></li></ul> <p>Select the default <b>STANDARD</b> since you are configuring the container to use <b>STANDARD</b> authentication.</p>
<b>User Name</b>	Enter the <b>Access Key Id</b> for this container.
<b>Password</b>	Enter the <b>Secret Access Key</b> for this container.
<b>Copy Number</b>	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.
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details.

- i Note:** If no data has been written to a controller, I/O path, or container, click **Delete** to remove the item, and then click **Apply** to save the changes.
8. **(Optional)** Repeat **Step 7** to add more containers to the same **AWS** Object Storage Destination.
  9. Click **Apply** to save your changes, or **Cancel** to exit without saving.
  10. **(Optional)** Repeat **Step 3** through **Step 9** to add additional **AWS** Object Storage Destinations.

## Add a New AWS Object Storage Destination that uses STS Authentication

Prior to configuring the AWS Object Storage Destinations with Storage Manager for AWS STS authentication, complete the following three steps, which apply to both AWS public cloud and GovCloud:

1. Have your AWS IAM Access Key Id and Secret Access Key ready.
2. Ensure that the S3 buckets, to be configured in Storage Manager, exist and you know both the name and the AWS region for each.

See <http://docs.aws.amazon.com/general/latest/gr/rande.html> for details on AWS Regions and Endpoints for AWS public cloud.

See <http://docs.aws.amazon.com/govcloud-us/latest/UserGuide/using-govcloud-endpoints.html> for details on AWS GovCloud (US) Endpoints.

3. Ensure that the IAM roles and their managed policies have been defined for your S3 buckets with AWS and that you know the roles' Amazon Resource Name (ARN). [Table 1 below](#) provides an example of a role's managed policy that Storage Manager recommends.

**Table 1:** Example of a role's managed policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListAllMyBuckets",
      "Resource": "arn:aws:s3:::*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": "arn:aws:s3:::joe-bucket"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:PutObject",
        "s3:DeleteObject",
        "s3:AbortMultipartUpload",
        "s3:RestoreObject"
      ],
      "Resource": "arn:aws:s3:::joe-bucket/*"
    }
  ]
}
```

**Note:** The Action "s3:RestoreObject" is only required if you are using Glacier Storage Class for your S3 bucket's life-cycle property.

Follow this procedure to add a new **AWS** Object Storage destination that uses **STS** authentication.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.

- Click **New...** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue...**
- Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new <b>AWS</b> Object Storage destination
<b>Provider</b>	Select <b>AWS</b> from the <b>Provider</b> list.
<b>Manager Host</b>	Enter <b>s3.amazonaws.com</b> for the <b>AWS</b> Object Storage Manager Host's address.
<b>Manager Port</b>	Enter <b>443</b> for the <b>AWS</b> Object Storage Manager Host's port.
<b>Manager Protocol</b>	Select <b>HTTPS</b> .
<b>Authentication Type</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default <b>Disabled</b> .
<b>User Name</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.
<b>Password</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.

- In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

- In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the " <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> " error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.
<b>Media Type</b> ( <i>New in StorNext 5 release 5.4</i> )	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. Select <b>AWS</b> from the drop-down list.
<b>URL Style</b>	Select <b>VHOST</b> from the drop-down list.

Parameter	Description
<b>Object Access Protocol</b>	This specifies the network protocol to be used for the host. Select <b>HTTPS</b> from the drop-down list.
<b>Host[:Port]</b>	Enter the AWS region endpoint for your S3 bucket. For example: <ul style="list-style-type: none"><li>• <b>s3-us-west-2.amazonaws.com</b> for an S3 bucket created in the US West (Oregon) region</li><li>• <b>s3.amazonaws.com</b> for an S3 bucket created in the US East (N. Virginia) region</li><li>• <b>s3-us-gov-west-1.amazonaws.com</b> for an S3 bucket created in AWS GovCloud</li></ul>

7. In the **Containers** section, perform the following:
  - a. Leave the **Container Selection** at the default **Manual**.
  - b. Click **Add** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Enter the name of your S3 bucket.
<b>Media ID</b>	Enter the StorNext <b>Media ID</b> associated with the selected container. The <b>Media ID</b> must be unique.
<b>Media Type</b>	Select <b>AWS</b> from the drop-down list.
<b>Storage Class</b>	This parameter specifies the storage class of your S3 bucket's life-cycle property. Make sure that this value always matches that configured for the S3 bucket itself. Available values include <b>standard</b> , <b>standard_ia</b> , and <b>glacier</b> . Most configurations use the default <b>standard</b> storage class.
<b>Signing Type</b>	This parameter specifies the signing type for the requests sent to the AWS Object Storage server. Available values include <b>V2</b> and <b>V4</b> (default).
<b>Authentication Type</b>	This parameter specifies the authentication type for the container being configured. Available values include: <ul style="list-style-type: none"><li>• <b>STANDARD</b></li><li>• <b>STS_PUBLIC</b></li><li>• <b>STS_GOVCLOUD</b></li><li>• <b>CAP</b></li></ul> Since you are configuring your bucket to use <b>STS</b> authentication, select <b>STS_PUBLIC</b> if your bucket is in AWS public cloud. Select <b>STS_GOVCLOUD</b> if your bucket is in AWS GovCloud.
<b>Role</b>	This parameter specifies the Amazon Resource Name (ARN) of the IAM role to assume for obtaining temporary credentials. Enter the IAM role's ARN defined to access this container.



Parameter	Description
<b>Role Duration</b>	This parameter specifies the duration, in seconds, of the role session or temporary credentials. The value must be in the range <b>900</b> to <b>3600</b> . A default value of <b>3600</b> seconds is used if a role duration is not specified.
<b>Authentication Endpoint</b>	This parameter specifies an alternate authentication endpoint , which is used to override the default STS server for the AWS public or GovCloud region. Otherwise, leave it at the default blank.
<b>User Name</b>	Enter the <b>Access Key Id</b> for this container.
<b>Password</b>	Enter the <b>Secret Access Key</b> for this container.
<b>Copy Number</b>	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.
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details.

**i Note:** If no data has been written to a controller, I/O path, or container, click **Delete** to remove the item, and then click **Apply** to save the changes.

8. **(Optional)** Repeat **Step 7** to add more containers to the same **AWS** Object Storage Destination.
9. Click **Apply** to save your changes, or **Cancel** to exit without saving.
10. **(Optional)** Repeat **Step 3** through **Step 9** to add additional **AWS** Object Storage Destinations.

## Add a New AWS Object Storage Destination that uses CAP Authentication

If you use CAP authentication, complete the following six steps prior to configuring an object storage destination:

1. Ensure that the S3 buckets, to be used by Storage Manager, have been created with AWS FedCloud (C2S account) and that you know the name and the AWS region endpoint of each bucket.
2. Ensure that the IAM roles and their managed policies have been defined for your S3 buckets with AWS FedCloud (C2S account). See [Table 1 on page 77](#) for an example of a managed policy which includes the required capabilities of the role.
3. Ensure that you have the following information:
  - a. The IAM role associated with your C2S account.
  - b. The agency associated with your C2S account.
  - c. The mission associated with your C2S account.

- d. Your CAP server connection endpoint.
  - e. Your customized certificate authority (CA) file in PEM format.
  - f. Your X.509 client certificate in PEM format, and any applicable private key file or passphrase.
4. Add the following Storage Manager sysparms to the file `/usr/adic/TSM/config/fs_sysparm_override` on the StorNext system to enable Storage Manager to communicate with the CAP server:
- a. **FS\_OBJSTORAGE\_C2S\_CAP\_HOSTPORT** identifies the connection endpoint for the CAP server and can be configured as follows:

```
FS_OBJSTORAGE_C2S_CAP_HOSTPORT=cap-portal:port;
```

- b. **FS\_OBJSTORAGE\_CAPATH** identifies the directory in which the issuer's certificate authority (CA) can be found if it is not already included in the operating system's default trusted root certificate file.

---

**i Note:** The certificate should be in PEM format.

For example, the certificate can be copied to `/usr/cvfs/config/ssl` and configured as follows:

```
FS_OBJSTORAGE_CAPATH=/usr/cvfs/config/ssl;
```

---

**i Note:** If your customized CA PEM file contains more than one certificate, we recommend that you append the content of your customized CA PEM file to your operating system's default CA bundle and to **NOT** use sysparm **FS\_OBJSTORAGE\_CAPATH** to set the location of your customized CA PEM file. Otherwise, you could split your CA PEM file into multiple CA PEM files, each of which contains a single CA certificate, and use sysparm **FS\_OBJSTORAGE\_CAPATH** to set the location of your newly split single certificate CA PEM files. See [HTTPS Default CA ROOT Certificate File or Path on page 292](#) for additional information on how to configure your customized CA PEM files.

- c. **FS\_OBJSTORAGE\_CLIENTCERT** identifies the location of the X.509 client certificate installed on the system for the CAP server to authenticate.

---

**i Note:** The certificate should be in PEM format.

For example, the client certificate can be copied to `/usr/cvfs/config/ssl/client-cert-filepath`, and configured as follows:

```
FS_OBJSTORAGE_CLIENTCERT=/usr/cvfs/config/ssl/client-cert-filepath;
```

- d. **FS\_OBJSTORAGE\_CLIENTKEY** which sets the location of the client private key if the client private key

is kept separately from (for example, not included in) the client certificate file. This parameter can be configured as follows:

```
FS_OBJSTORAGE_CLIENTKEY=/usr/cvfs/config/ssl/client-key-filename;
```

- e. **FS\_OBJSTORAGE\_CLIENTKEY\_PASS** which specifies the passphrase used to protect the client private key and can be configured as follows:

```
FS_OBJSTORAGE_CLIENTKEY_PASS=passphrase;
```

- 5. Execute the following command to generate the hash for your certificates:

```
/opt/quantum/openssl/bin/c_rehash /usr/cvfs/config/ssl
```

- 6. Restart TSM to allow the system parameter changes to take effect.

**i Note:** You can also use the GUI's [System Parameters on page 437](#) page to set the above system parameters.

Follow this procedure to add a new **AWS** Object Storage destination that uses **CAP** authentication.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Click **New....** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue....**
4. Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new <b>AWS</b> Object Storage destination
<b>Provider</b>	Select <b>AWS</b> from the <b>Provider</b> list.
<b>Manager Host</b>	Enter <b>s3.amazonaws.com</b> for the <b>AWS</b> Object Storage Manager Host's address.
<b>Manager Port</b>	Enter <b>443</b> for the <b>AWS</b> Object Storage Manager Host's port.
<b>Manager Protocol</b>	Select <b>HTTPS</b> .
<b>Authentication Type</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default <b>Disabled</b> .

Parameter	Description
<b>User Name</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.
<b>Password</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.

5. In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

6. In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the " <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> " error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.
<b>Media Type</b> <i>(New in StorNext 5 release 5.4)</i>	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. Select <b>AWS</b> from the drop-down list.
<b>URL Style</b>	Select <b>VHOST</b> from the drop-down list.
<b>Object Access Protocol</b>	This specifies the network protocol to be used for the host. Select <b>HTTPS</b> from the drop-down list.
<b>Host[:Port]</b>	Enter the AWS region endpoint for your S3 bucket.

7. In the **Containers** section, perform the following:

- a. Leave the **Container Selection** at the default **Manual**.
- b. Click **Add** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Enter the name of your S3 bucket.
<b>Media ID</b>	Enter the StorNext <b>Media ID</b> associated with the selected container. The <b>Media ID</b> must be unique.

Parameter	Description
<b>Media Type</b>	Select <b>AWS</b> from the drop-down list.
<b>Storage Class</b>	This parameter specifies the storage class of your S3 bucket's lifecycle property. Make sure that this value always matches that configured for the S3 bucket itself. Available values include <b>standard</b> , <b>standard_ia</b> , and <b>glacier</b> . Most configurations use the default <b>standard</b> storage class.
<b>Signing Type</b>	This parameter specifies the signing type for the requests sent to the AWS Object Storage server. Available values include <b>V2</b> and <b>V4</b> (default).
<b>Authentication Type</b>	<p>This parameter specifies the authentication type for the container being configured. Available values include:</p> <ul style="list-style-type: none"><li>• <b>STANDARD</b></li><li>• <b>STS_PUBLIC</b></li><li>• <b>STS_GOVCLOUD</b></li><li>• <b>CAP</b></li></ul> <p>Select <b>CAP</b> since you are configuring the container to use <b>CAP</b> authentication.</p>
<b>Role</b>	This parameter specifies the IAM role associated with the target C2S account for obtaining temporary credentials. Enter the IAM role associated with your C2S account to access this container.
<b>Role Duration</b>	This parameter specifies the duration, in seconds, of the role session or temporary credentials. The value must be in the range <b>900</b> to <b>3600</b> . A default value of <b>3600</b> seconds is used if a role duration is not specified.
<b>CAP Agency</b>	This parameter specifies the CAP agency associated with the target C2S account for obtaining temporary credentials. Enter the agency associated with your C2S account for this container.
<b>CAP Mission</b>	This parameter specifies the CAP mission associated with the target C2S account for obtaining temporary credentials. Enter the mission associated with your C2S account for this container.
<b>Copy Number</b>	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.
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details.

- 
- i Note:** If no data has been written to a controller, I/O path, or container, click **Delete** to remove the item, and then click **Apply** to save the changes.
  - 8. **(Optional)** Repeat **Step 7** to add more containers to the same AWS Object Storage Destination.
  - 9. Click **Apply** to save your changes, or **Cancel** to exit without saving.
  - 10. **(Optional)** Repeat **Step 3** through **Step 9** to add additional **AWS** Object Storage Destinations.

## Edit an AWS Object Storage Destination

Follow this procedure to edit an existing **AWS** Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the **AWS** 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.

---

**i 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 **AWS** Object Storage destination was successfully modified, click **OK**.

## Delete an AWS Object Storage Destination

Follow this procedure to delete an existing **AWS** Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the **AWS** 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 Object Storage destination was successfully deleted, click **OK**.

## Setting up Azure Object Storage Destinations on a StorNext Configuration

Storage Manager has Object Storage support for Microsoft Azure Cloud Services that includes:

- HTTP and HTTPS access to Microsoft Azure containers, though HTTPS is the recommended protocol for access Azure containers
- Append blob storage service

## View Azure Object Storage Destinations

Follow this procedure to view a list of currently configured Azure Object Storage destinations.

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab. Information for any previously configured **Azure** Object Storage destinations is shown as entries that have **Microsoft Azure** listed as the **Provider**. For each configured destination, the screen displays the **Name**, **Provider**, **Appliance 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 **Azure** Object Storage destination whose information you want to view.
4. Click **View...**
5. When you are finished viewing the destination information, click **Done**.

## Add a New Azure Object Storage Destination

Follow this procedure to add a new **Azure** Object Storage destination.

1. Ensure that the **Azure** containers that you are configuring with Storage Manager, have been created with Microsoft Azure Cloud Services; and that you know the **Storage Account Name**, **Storage Access Key**, and names of your **Azure** containers.
2. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
3. Click the **Object Storage** tab.
4. Click **New...** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue...**
5. Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new <b>Azure</b> Object Storage destination
<b>Provider</b>	Select <b>Microsoft Azure</b> from the <b>Provider</b> list.
<b>Manager Host</b>	Enter <b>portal.azure.com</b> for the <b>Azure</b> Object Storage Manager Host's address.
<b>Manager Port</b>	Enter <b>443</b> for the <b>Azure</b> Object Storage Manager Host's port.
<b>Manager Protocol</b>	Select <b>HTTPS</b> .

Parameter	Description
<b>Authentication Type</b>	This parameter <b>does not</b> apply to <b>Azure</b> Object Storage Destination. Leave it at the default <b>Disabled</b> .
<b>User Name</b>	This parameter <b>does not</b> apply to <b>Azure</b> Object Storage Destination. Leave it at the default blank.
<b>Password</b>	This parameter <b>does not</b> apply to <b>Azure</b> Object Storage Destination. Leave it at the default blank.

6. In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

7. In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the " <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> " error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.
<b>Media Type</b> <i>(New in StorNext 5 release 5.4)</i>	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. Select <b>Azure</b> from the drop-down list.
<b>URL Style</b>	Leave the URL style at the default <b>PATH</b> .
<b>Object Access Protocol</b>	This specifies the network protocol to be used for the host. Select <b>HTTPS</b> from the drop-down list.
<b>Host[:Port]</b>	Enter <b>blob.core.windows.net</b> for the I/O path's connection endpoint to your Azure containers.

8. 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.
  - i. If you select **Scan**, check the **Use different credentials** box and enter the username (the **Azure Storage Account Name**) and a password (the **Azure Storage Access Key**). The



page displays a pre-populated list of available containers when you add container.

ii. If you select **Manual**, the page displays a text box to manually enter the name of the container.

b. Click **Add** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Select ( <b>Scan</b> mode) or enter ( <b>Manual</b> mode) the name of your <b>Azure</b> container.
<b>Media ID</b>	Enter the StorNext <b>Media ID</b> associated with the selected container. The <b>Media ID</b> must be unique.
<b>Media Type</b>	Select <b>Azure</b> from the drop-down list.
<b>Storage Class</b>	This parameter specifies the storage class for the <b>Azure</b> Object Storage media. Leave it at the default <b>azure_append_blob</b> .
<b>Signing Type</b>	This parameter specifies the signing type for the requests sent to the <b>Azure</b> Object Storage server. Leave it at the default <b>azure</b> .
<b>Authentication Type</b>	This parameter specifies the authentication type for the container being configured. Use the default value of <b>STANDARD</b> , which authenticates with the Storage Account Name and Storage Access Key.
<b>Account Name</b>	Enter the <b>Azure Storage Account Name</b> for this container.
<b>Key</b>	Enter the <b>Azure Storage Access Key</b> for this container.
<b>Copy Number</b>	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.
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details.

**i** **Note:** If no data has been written to a controller, I/O path, or container, click **Delete** to remove the item, and then click **Apply** to save the changes.

9. **(Optional)** Repeat **Step 8** to add more containers to the same **Azure** Object Storage Destinations.
10. Click **Apply** to save your changes, or **Cancel** to exit without saving.
11. **(Optional)** Repeat **Step 4** through **Step 10** to add additional **Azure** Object Storage Destinations.

## Edit an Azure Object Storage Destination

Follow this procedure to edit an existing **Azure** Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the **Azure** 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.

**i 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 **Azure** Object Storage destination was successfully modified, click **OK**.

## Delete an Azure Object Storage Destination

Follow this procedure to delete an existing **Azure** Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the **Azure** 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 Object Storage destination was successfully deleted, click **OK**.

## Change the Current State of Object Storage Destinations, Controllers, and I/O Paths

You can also change the current state of existing Object Storage destinations, controllers, and I/O paths. To change the state, select the Object Storage destination, and then choose one of these options from the **Select Action** drop-down list:

Parameter	Description
<b>Online</b>	Select this option to set the Object Storage destination online.
<b>Offline</b>	Select this option to take the Object Storage destination offline.
<b>Controllers Online</b>	Select this option to set the controllers online.
<b>Controllers Offline</b>	Select this option to take the controllers offline.

Parameter	Description
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.


## Setting Up Q-Cloud

The **Q-Cloud** tab on the **Configuration > Storage Destinations** page enables you to perform actions pertaining to Q-Cloud resources. The **Q-Cloud** page displays a table showing the Q-Cloud resources that have been registered/activated for use in StorNext.

For information on troubleshooting Object Storage and Cloud errors, see [Debugging StorNext for Object Storage Systems and Cloud Providers on page 638](#).

## Q-Cloud Retrieval

- File retrieval from Q-Cloud Vault can take up to 5 hours.
- Retrieval requests to StorNext for files in Q-Cloud Vault are issued immediately, similar to other devices supported by Storage Manager.
- If an application requests multiple files (issues these requests one file at a time) and then waits for each retrieval to complete before issuing the next file to be retrieved, each file can take up to 5 hours. Since the requests are sequential (it waits for the previous retrieve to complete for two files), file retrieval can take up to 10 hours to complete. However, if the retrievals are submitted to StorNext concurrently, it can take up to 5 hours for two files to complete.
- While each file retrieval waits for completion (up to 5 hours), the retrieval process consumes and retains a certain amount of system resources for that length of time until completion.

 **Caution:** A large number of file retrieval requests issued concurrently can severely impact system performance.

## Q-Cloud and Partial File Retrieval (PFR)

- The **Partial File Retrieval** feature is supported with Q-Cloud Archive targets, with the exception of configurations where client-side encryption or compression is used.
- The **Partial File Retrieval** feature is not supported with Q-Cloud Vault targets.

For additional information on **Partial File Retrieval**, see the *StorNext Partial File Retrieval User's Guide*.

## Firewall Rules and IP Ranges Specific to Q-Cloud

If the StorNext configuration has restricted network access because of firewall rules, the firewall configuration may need to be updated to use Q-Cloud.

### Configure the Firewall if the Software Supports Named Entries

1. Open HTTPS access (port **443**) to `s3-us-west-2.amazonaws.com` (note, this is an example for the **US-West-2** region). The appropriate region should be used in its place.
2. For access to the Q-Cloud Controller, open HTTPS access (port **443**) to `api-qcloud.quantum.com`. This allows you to configure and validate your Q-Cloud installation.

### Configure the Firewall if the Software Does Not Support Named Entries

1. Locate the IP addresses for the appropriate AWS region in <https://ip-ranges.amazonaws.com/ip-ranges.json>. Allow access to these addresses. The list may change several times per week.


**i Note:** Amazon rotates IP addresses through a range in order to provide access to the AWS S3 service. Each region of AWS has a specific set of IP ranges that are described at <https://aws.amazon.com/blogs/aws/aws-ip-ranges-json/>, and listed at <https://ip-ranges.amazonaws.com/ip-ranges.json>.

2. Determine the IP address of `api-qcloud.quantum.com`, and open HTTPS access to this address.

Alternatively, open all outgoing HTTPS traffic from the StorNext installation to the Internet.

## Parameters and Buttons on the the Q-Cloud Page

Parameter/Button	Description
<b>Media ID</b>	Displays the <b>Media ID</b> representing a purchased resource. It relates the Q-Cloud bucket to a StorNext media.
<b>Product Key</b>	Displays the <b>Product Key</b> associated with the StorNext media.
<b>Provider</b>	Displays the underlying storage provider.
<b>State</b>	Displays the operational state of the storage media.
<b>Max Streams</b>	The <b>Max Streams</b> value defines the number of concurrent I/O operations that can write concurrently. By default, the maximum number of streams is set to <b>48</b> , or you can select a different value from the <b>Max Streams</b> drop-down list when you click <b>Preferences....</b> See <a href="#">Configure a Media ID Preference on page 94</a> .
<b>Copy Number</b>	Displays the copy number assigned to the Q-Cloud media. If files do not exist on the media, you can select copy number <b>1</b> , <b>2</b> , <b>3</b> , or <b>4</b> from the list when you click <b>Preferences....</b> See <a href="#">Configure a Media ID Preference on page 94</a> .

Parameter/Button	Description
<b>Policy Class</b>	Displays the name of the policy.
<b>Evaluation Key</b>	Displays <b>Yes</b> or <b>No</b> . If <b>Yes</b> , the product key is an evaluation key. If <b>No</b> , the product key is not an evaluation key.
<b>Expiration Date</b>	Displays the expiration date of the selected product key.
<b>Overused</b>	This parameter applies to evaluation keys and indicates if you have exceeded the limits of the evaluation. For example, too much data stored in Q-Cloud.
<b>Region</b>	Displays the region the Q-Cloud media belongs to.
<b>File Count</b>	Displays the number of files currently stored on the media.
<b>Manage Keys...</b>	Click to display the <b>Manage Keys</b> page where you can add or remove additional <b>Product Keys</b> . For additional procedures, see <a href="#">Add a Product Key on the next page</a> and <a href="#">Delete a Product Key on the next page</a> below.
<b>Preferences...</b>	Click to configure your preferences for the currently selected product key. See <a href="#">Configure a Media ID Preference on page 94</a> .
<b>Check Connectivity</b>	Click to check the connectivity to the Q-Cloud media. If connectivity fails, contact Quantum Technical Support.
<b>Refresh</b>	Click to update the <b>Q-Cloud</b> page with current information.
<b>Select Action (Online)</b>	Select <b>Online</b> manually change the media state to online.
<b>Select Action (Offline)</b>	Select <b>Offline</b> to manually change the media state to offline.
<b>Select Action (Remove All Q-Cloud Storage)</b>	Select <b>Remove All Q-Cloud Storage</b> to purge and factory reset the system.  <b>Note:</b> The option is only available if the Q-Cloud media do not contain files.

## Parameters and Buttons on the Manage Keys Page

On the **Configurations > Storage Destinations > Q-Cloud** page, click **Manage Keys...** to display the **Manage Keys** page.

For additional procedures, see [Add a Product Key on the next page](#) and [Delete a Product Key on the next page](#) below.

Parameter/Button	Description
<b>Access ID</b>	Enter the customer access ID provided by Quantum.

Parameter/Button	Description
<b>Product Keys</b>	Enter the customer product key provided by Quantum.
<b>Delete</b>	Click to remove the selected <b>Product Key</b> from the list.  You must enter an <b>Access ID</b> in order to delete <b>Product Keys</b> .
<b>Add</b>	Click to add a new <b>Product Key</b> to the list after you enter it into the <b>Product Keys</b> field. You must enter an <b>Access ID</b> in order to add <b>Product Keys</b> .
<b>Apply</b>	Click to commit the edited list of <b>Product Keys</b> and/or <b>Access ID</b> .
<b>Reset</b>	Click to restore the current list.
<b>Cancel</b>	Click to cancel the editing session and return to the <b>Q-Cloud</b> page.

## Add a Product Key

1. In the **Access ID** field, enter the customer access ID provided to you by Quantum. Upon entering the **Access ID** once and successfully committing a configuration, the **Access ID** field remains populated when you return to the **Manage Keys** page.
2. In the **Product Keys** field (adjacent to the **Add** button), enter the customer product key provided to you by Quantum, select a **Policy Class** to assign the Q-Cloud media directly to it, and then click **Add**. The **Product Keys** list (above the field) displays the value just added. You must enter an **Access ID** in order to add **Product Keys**. Repeat this step to add additional product keys.

**i Note:** If a **Policy Class** is not selected, the media is assigned to **System Blank** (in other words, the default blank pool).

3. Click **Apply**. The **Product Key** appears on the main **Q-Cloud** page. Upon clicking **Apply**, the actual process to commit the configuration is run as a background task. The GUI returns to the **Q-Cloud** page immediately if there are no locally detected errors. In some cases, the background process may require some time to execute and the **Q-Cloud** page will not immediately reflect your changes. As with other similar processes, navigate to the [Jobs on page 365](#) page to verify the outcome of the background task. In the case of a delayed success, click **Refresh** on the **Q-Cloud** page to display the current state of the system. In the case of a failure, the **Q-Cloud** page may not update; navigate to the [Jobs on page 365](#) page to determine the cause of the error.
4. **(Optional)** Click **Reset** to restore the current list.
5. **(Optional)** Click **Cancel** to cancel the editing session and return to the **Q-Cloud** page.

## Delete a Product Key

1. In the **Access ID** field, enter the customer access ID provided to you by Quantum.
2. In the **Product Keys** list, click a product key, and then click **Delete** to remove the selected product key from the **Product Keys** list. You must enter an **Access ID** in order to delete **Product Keys**. Repeat

this step to delete additional product keys.

3. Click **Apply**. The **Product Key** is removed on the **Q-Cloud** page. Upon clicking **Apply**, the actual process to commit the configuration is run as a background task. The GUI returns to the **Q-Cloud** page immediately if there are no locally detected errors. In some cases, the background process may require some time to execute and the **Q-Cloud** page will not immediately reflect your changes. As with other similar processes, navigate to the [Jobs on page 365](#) page to verify the outcome of the background task. In the case of a delayed success, click **Refresh** on the **Q-Cloud** page to display the current state of the system. In the case of a failure, the **Q-Cloud** page may not update; navigate to the [Jobs on page 365](#) page to determine the cause of the error.
4. **(Optional)** Click **Reset** to restore the current list.
5. **(Optional)** Click **Cancel** to cancel the editing session and return to the **Q-Cloud** page.

## Configure a Media ID Preference

1. On the **Configurations > Storage Destinations > Q-Cloud** page, click a radio button to select a media ID you want to configure (corresponding to a media ID located under the **Media ID** column).
2. Click **Preferences...** to display the **Preferences** page. The table below provides a description for each parameter/button on the **Preferences** page.

Parameter/Button	Description
<b>Product Key</b>	Displays the product key corresponding to the <b>Media ID</b> you selected on the Q-Cloud page from Step 1.
<b>Max Streams</b>	The <b>Max Streams</b> value defines the number of concurrent I/O operations that can write concurrently. By default, the maximum number of streams is set to <b>48</b> , or you can select a different value from the <b>Max Streams</b> drop-down list.
<b>Copy Number</b>	Displays the copy number assigned to the <b>Q-Cloud</b> media. Select copy number <b>1, 2, 3,</b> or <b>4</b> from the list if files do not exist on the media.
<b>Update</b>	Click to update your preferences for the currently selected <b>Product Key</b> .
<b>Cancel</b>	Click to cancel the editing session and return to the <b>Q-Cloud</b> page.

3. In the **Max Streams** drop-down list, select the number of streams. For **Copy Number**, the only option available is 1.
4. Click **Update** to proceed and display the confirmation dialog requesting you confirm your preferences. Perform one of the following:
  - Click **Yes** to confirm your selected preferences. A confirmation dialog appears; click **OK**. The **Q-Cloud** page is displayed with your updated preferences.
  - Click **No** to disregard your selected preferences and return to previous page.
5. **(Optional)** Click **Cancel** to cancel the editing session and return to the **Q-Cloud** page.

## Replace a Product Key

Follow the procedure below to replace your Q-Cloud Product Key. For example, if your Q-Cloud Product Key is compromised and you are reissued a new Product Key from Quantum.

1. On the **Configurations > Storage Destinations > Q-Cloud** page, click a radio button to select a media ID you want to configure (corresponding to a media ID located under the **Media ID** column).
2. Click **Preferences...** to display the **Preferences** page (see [Configure a Media ID Preference on the previous page](#)).
3. Note the **Product Key** corresponding the **Media ID** you selected in **Step 1**.
4. Click **Cancel** to return to the **Q-Cloud** page.
5. Click **Manage Keys.....**
6. In the **Product Keys** list, click the obsolete product key (from **Step 3**) and then click **Delete**.
7. In the **Product Keys** field (adjacent to the **Add** button), enter the **new** product key provided to you by Quantum, and then click **Add**. The **Product Keys** list (above the field) displays the value just added.
8. Click **Apply**. The new **Product Key** appears on the **Q-Cloud** page.

---

**i Note:** If you click **Apply** after deleting the obsolete product key (**Step 6**) without adding the **new** product key, you will receive a **putConfig** error message.

## Enable Encryption on the Q-Cloud Data

1. On the **Configuration** menu, click **Storage Policies**.
2. Click the radio button corresponding to the Q-Cloud policy class you want to enable encryption on, and then click **Edit....**
3. Click the **General** tab.
4. Click the **Encryption (Q-Cloud only)** list, and then select **Server AES256**.
5. Click **Apply** to save your changes to the policy class.
6. **(Optional)** Click **Cancel** to discard your changes and return to the **Storage Manager Policies** page.

## Backup and Restore Using Q-Cloud

Contact Quantum Technical Support for assistance with disaster recovery to and from Q-Cloud.

## Configure Data Replication

The **Configuration** menu's **Storage Destinations > Replication Targets** option enables you to view or edit currently configured data replication targets, and to add a new host or mount point for additional targets.

For more information about the replication feature, see [Replication and Deduplication on page 213](#) and [Additional Replication and Deduplication Information on page 521](#).



## View Replication Targets

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Replication Targets** tab.
3. Click the plus sign (+) beside the **Replication Targets** heading in the box titled **Replication Target Configuration**.
4. Click the plus sign beside the replication target to view the target's mount point.


## Edit a Replication Target

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Replication Targets** tab.
3. If necessary, click the plus sign (+) beside the **Replication Targets** heading in the box titled **Replication Target Configuration**.
4. Select the replication target you want to edit.
5. Click **Edit**.
6. At the **Host Name** or **IP Address** field, modify either the host name or IP address for the replication target.
7. Click **Update** to save your changes, or **Cancel** to abort.

## Delete a Replication Target

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Replication Targets** tab.
3. If necessary, click the plus sign (+) beside the **Replication Targets** heading in the box titled **Replication Target Configuration**.
4. Select the replication target you want to delete.
5. 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**.

## Add a New Host

1. Click the **Replication Targets** tab. The **Configuration > Storage Destinations > Replication Targets** page appears.
2. Click **New...**

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

## Add a New Mount Point

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Replication Targets** tab.
3. Select the replication target (host) to which you would like to add a mount point. If necessary, click the plus sign (+) beside the **Replication Targets** heading in the box titled **Replication Target Configuration**.
4. Click **Add Mount Point**.
5. Click **Scan Host** to identify available mount points on the selected host.
6. Enter the path name for the mount point you are adding. The first part of the path may be displayed by default.
7. Click **Add** to save the change, or **Cancel** to abort without saving.
  - Repeat **Step 3** through **Step 6** to add additional mount points.
8. When you are finished adding mount points, click **Apply**.
9. After a message informs you that changes were successfully incorporated, click **OK**.

## Configure Data Deduplication

The **Deduplication** tab enables you to create a blockpool on a specified file system. To create the blockpool, select the desired file system from the dropdown list next to the **Blockpool File System** label, and then click **Apply**.

---

**i 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.

# Storage Policies

There are two types of StorNext storage policies:

- Storage Manager
- Replication/Deduplication

Replication/Deduplication storage policies control the way StorNext's replication and deduplication features behave. For more information about the replication and deduplication features, see [Replication and Deduplication on page 213](#) and [Additional Replication and Deduplication Information on page 521](#).

## What is the functionality of a Storage Manager storage policy?

- A Storage Manager storage policy defines how the Storage Manager feature manages files in a directory and its 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
- A Storage Manager storage policy is used to make copies of files to tape or to storage disk, and also to relocate files between stripe groups within a file system.
- A Storage Manager storage policy does not control replication or deduplication.

## Examples of Storage Manager 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.

---

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

## 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. In this case, the policy could be configured 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`.

A Replication/Deduplication storage policy defines the parameters governing the data replication process, including inbound and outbound replication parameters, and enabling or disabling data deduplication and truncation.

## Adding a Storage Manager Policy


### Add a New Storage Manager Storage Policy

1. If you have not already done so, choose **Storage Policies** from the **Configuration** menu.
2. Click **New**.
3. Enter the following fields:
  - **Policy Class:** The name of the new policy you are creating. 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 fields on 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 page 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.

The **General** tab contains the following fields:

Parameter	Description
<b>File Age Before Copy is Made</b>	<p>This value determines the minimum number of minutes, hours or days a file must reside unmodified on disk before it is considered a candidate for copying to storage media. Enter the number in the first field, and the unit of measure in the second pull-down field. The minimum value is 1 minute. The <b>default</b> value is 5 minutes.</p> <p> <b>Note:</b> A minimum of one hundred files is required to trigger the copy after 1 minute.</p>
<b>File Age Before Truncation</b>	<p>This value determines the minimum number of minutes, hours or days a file must reside on a disk unaccessed before it is considered a candidate for truncation from disk. Truncation removes the disk blocks of a stored file, but not the file itself. Enter the number in the first field, and the unit of measure in the second pull-down field. The minimum value is 5 minutes.</p>
<b>Default Media Type</b>	<p>Select the appropriate media type that applies to devices associated with this policy. When creating a policy, the <b>Default Media Type</b> you select applies to all defined copies in the policy unless overridden in the policy. For additional information, see <a href="#">The Steering Tab on page 102</a>.</p>
<b>Truncate File Immediately After Store</b>	<p>Enable this option (check this box) to truncate files immediately after they are stored.</p>
<b>Clean Database When File Removed</b>	<p>If this option is enabled (the box is checked), StorNext cleans or consolidates the database after a file is removed.</p>
<b>Generate Checksum</b>	<p>If this option is enabled, (the box is checked), checksums are generated and retained in the database for files stored by the corresponding policy.</p>
<b>Validate Checksum</b>	<p>If this option is enabled (the box is checked), checksums are compared to retained values for the files retrieved by the corresponding policy. The Checksum feature consumes additional space in the StorNext database whether it is enabled or not. When disabled, this feature consumes approximately 2 bytes per stored file; when enabled, this feature consumes approximately 18 bytes per stored file. The database stores data in files on the host computer, so the increase in database size translates to a corresponding increase in disk space requirements. The exact amount of space consumed (whether the feature is enabled or disabled) may vary.</p>
<b>Alternate Store Location</b>	<p>Provides an automatic system for copying files from a main instance of StorNext to a remote instance of StorNext at the same time as copies are made to tertiary storage at the main site.</p>
<b>Maximum Inactive Versions</b>	<p>The maximum number of inactive versions of a file StorNext keeps track of for recovery purposes.</p>

Parameter	Description
<b>Tape Drive Pool</b>	Associates the drive pool to use with the policy. If you specify a drive pool, the drive pool name must be defined before any data operation can occur.
<b>Media Email Warning Limit</b>	The warning limit for the number of media that can be allocated in the policy class. If you specify an amount at this field, when the storage disk reaches the amount you specify you will receive an email message warning that the device is nearing capacity. The maximum number for the limit is 65,000.
<b>Media Limit</b>	The maximum number of media that are allowed in this policy class. When this limit is reached you will receive an email message warning. Files can still be stored in the policy class as long as there is room on the media that have already been used to store files in that policy class. The maximum number for the limit is 65,000.
<b>Media Cleanup</b>	This option determines the action StorNext applies to storage media follows after media are cleaned. Select <b>return to scratch pool</b> or <b>keep in policy</b> .
<b>Stub Files</b>	Select this option to enable the Stub File feature. When this feature is enabled, third-party applications can gather information about a file by reading a portion of the file (called a stub) rather than reading the entire file. When you enable stub file support for a storage policy, you must specify the size of the stub file (in kilobytes). When stub file support is enabled, the beginning portion of the file (up to the size you specified) remains on disk after data blocks are freed during policy management or space management.
<b>Stub File Size (KB)</b>	When the Stub File feature is enabled, specify the target size for the stub file in kilobytes.
<b>Delay File Storage</b>	The store candidates in the policy class have to add up to Minimum File Size before any of the files will be stored. When this option is enabled, files are not stored immediately.
<b>Minimum File Size (1 to 999 in MB or GB)</b>	This value determines the minimum size (in megabytes or gigabytes) all valid store candidates in the policy class combined must reach before they are stored. This field is enabled only if the <b>Delay File Storage</b> option is also enabled.
<b>Maximum File Age (Hours)</b>	This option works in conjunction with the <b>Delay File Storage</b> option so that files will never be stored because the <b>Minimum File Size</b> limit has not been reached. As soon as any file in the policy class reaches the Maximum File Age, all of the files will be stored. This value determines the time after which any valid store candidate in the policy is stored.
<b>Retrieve to Affinity</b>	This field enables you to indicate that you want StorNext to retrieve files to an affinity other than the primary affinity. To retrieve to an alternative affinity, select the desired affinity from the drop down list. To retrieve to the primary affinity, select <b>None</b> , which is the default value.

Parameter	Description
<b>Compression (Q-Cloud only)</b>	This option allows you to enable compression on the data corresponding to the policy. Note, compression is only available for Q-Cloud-based policies. For additional information, see <a href="#">Client Side Compression and Encryption on page 114</a> .
<b>Encryption (Q-Cloud only):</b>	This option allows you to enable encryption on the data corresponding to the policy. Note, encryption is only available for Q-Cloud-based policies. For additional information, see <a href="#">Client Side Compression and Encryption on page 114</a> .
<b>Encryption Master Key (Q-Cloud only)</b>	The option allows you to select an existing master key. If you intend to create a new master key, you can do so by selecting the <b>Add New...</b> option from the drop-down list. This parameter is enabled only if <b>Encryption (Q-Cloud only)</b> is set to <b>Client AES 256</b> .

## 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.

The **Relocation** tab contains the following fields:

Parameter	Description
<b>Disk-to-Disk Relocation</b>	To enable this feature, select from the dropdown list the destination disk to which you want to move data. Select <b>None</b> , the default value, to disable this feature.
<b>File Age Before Relocation</b>	When Disk-to-Disk relocation is enabled, specify the number of seconds, minutes, hours or days a file must reside on a disk unaccessed before file relocation commences. Enter the number in the first field, and the unit of measure in the second pull-down field. The minimum value is 5 minutes.

## 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. In addition, you can manage the number of tape drives to use per store policy by using the **Drive Limit** feature.

The **Steering** tab contains the following fields:

Parameter	Description
<b>Media Type</b>	For each copy (1 - 4), select from the <b>Media Type</b> list the media type for devices in the drive pool. This can be an actual device type such as LTO, Lattus, S3COMPAT, Q-Cloud, or a Storage Disk. If you select <b>None</b> for a given copy, the <b>Media Format</b> , <b>Tape Drive Pool</b> , and <b>Drive Limit</b> lists for that specific copy are grayed out, and you are not able to configure the parameters for that specific copy.
<b>Media Format</b>	For each copy (1 - 4), select from the <b>Media Format</b> list the media format for LTO devices in the drive pool. Select <b>ANTF</b> or <b>LTFS</b> . <b>ANTF</b> is the Quantum internal tape format. <b>LTFS</b> is the Linear Tape File System specification tape format. <b>ANTF</b> will format media with a single partition containing <b>ANTF</b> volume labels. This data partition will store StorNext file data. <b>LTFS</b> will format media with two partitions containing <b>LTFS</b> volume labels. The index partition will store <b>LTFS</b> metadata and the data partition will store StorNext file data.
<b>Tape Drive Pool</b>	For each copy (1 - 4), select from the <b>Tape Drive Pool</b> list the specific drive pool to which you want to write that copy's data. A <b>Tape Drive Pool</b> must be defined before it can be selected. For more information, see refer to <a href="#">Drive Pools on page 159</a> .
<b>Drive Limit</b>	For each copy (1 - 4) in the <b>Drive Limit</b> list, select the number of tape drives to use per policy and copy. The <b>Drive Limit</b> feature allows you to manage the number of drives used per store policy. The default is <b>None</b> ; no limit will be set and all drives will be used on stores. The <b>Drive Limit</b> feature is not supported for Storage Disks or Object Storage.

## 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 180](#))

The **Schedule** tab contains the following field:

- **Store Files Automatically:** Select this option to enable the **Store Files Automatically** feature.

## The Associated Directories Tab

The **Associated Directories** tab enables you to view or delete existing relation points (directories) in the file system for the policy, and to add new relation points.

- **Note:** Review the content for the CLI command `fsaddrrelation`, in latest version of the *MAN Pages Reference Guide* (available online at <http://www.quantum.com/sndocs>) for information on limitations and constraints.

The **Associated Directories** tab contains the following fields:

- **File System:** Select from the dropdown list the file system for which you want to view, delete or add relation points.




- **Directory:** Select the directory you want to add, or type the name of a new directory in the field to the left of the **Create Directory** button. Click **Add** to add that directory as a new relation point.
- **Associated Directories:** This area displays any existing associated directories.

## Deleting an Associated Directory

To delete a directory listed under the **Associated Directories** heading, select the desired relation point and then click **Delete**. The directory must be empty before you can delete it.

---

 **Caution:** This particular delete function does not provide a confirmation message, so be absolutely sure you want to delete the selected relation point before you click **Delete**. The selected relation point will be deleted immediately after you click **Delete**, and there is no Undo feature.

## Viewing, Running, Editing, Deleting or Testing a Storage Policy

On the **Setup > Storage Policy** page you can view, edit or delete existing storage policies (in addition to creating new policies as described in [Adding a Storage Manager Policy on page 99](#) and [Adding a Replication or Deduplication Policy on the next page](#)).

### View Storage Policy Details for a Storage Manager or Replication Policy

1. On the **Configuration > Storage Policies** page, select the storage policy you wish to view.
2. Click **View**.
3. Click **Done** to return to the **Configuration > Storage Policies** page.

### Run a Storage Policy

1. On the **Configuration** menu, click **Storage Policies**.
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, on the **Reports** menu, click **Jobs**.


### Edit a Storage 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 [Adding a Storage Manager Policy on page 99](#).

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, there is no **Blackout** tab. For more information about fields on these tabs, see [Adding a Storage Manager Policy on page 99](#).

1. From the **Configuration > Storage Policies** page, 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.
4. Click **Apply** to save changes and return to the **Configuration > Storage Policies** page, or **Cancel** to abort.


---

 **WARNING:** Adding or removing steering copies from an active policy is not retroactive to existing, stored files. Only new or modified files are stored to the updated steering destinations. To re-process existing stored files to reflect the updated steering destinations, contact Quantum Technical Support.

## Delete a Storage Policy

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

---

 **WARNING:** Adding or removing steering copies from an active policy is not retroactive to existing, stored files. Only new or modified files are stored to the updated steering destinations. To re-process existing stored files to reflect the updated steering destinations, contact Quantum Technical Support.

## Test a Replication Storage Policy

1. From the **Configuration > Storage Policies** page, select the replication policy for you wish to do a test run.
2. Click **Test**.
3. A message informs you that a job has been submitted to run the policy. Click **OK** to continue.
4. If desired, check job status by clicking **Jobs** from the **Reports** menu. For more information, see [Jobs on page 365](#).

# Adding a Replication or Deduplication Policy

## Add a Replication Storage Policy

Before you create a replication storage policy you should have already created, at a minimum, a source and

target file system. Also, you will not be able to create the policy unless you have specified the blockpool. For more information, see [Name Servers on page 30](#).

1. If you have not already done so, choose **Storage Policies** from the **Configuration** menu.
2. Click **New**.
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
    - The **Replication/Deduplication** button remains disabled (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 background job can take multiple minutes depending on your licensed deduplication capacity. This button will become 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 from the dropdown list the file system you intend to use as your source
4. Click **Configure**.
5. Enter information on the **Deduplication**, **Outbound Replication**, **Inbound Replication** and **Source Directories** tabs. (See the sections following for more information about fields on these tabs.)
6. When you are finished entering information about the new policy, click **Apply**, or click **Cancel** to exit without saving.
7. After a message informs you that the policy was created successfully, click **OK**.

## Inheritance and Overriding

Many fields contain a button that toggles between **Inherit** and **Override**. When the button is **Inherit**, the value for that parameter is "inherited" from the global replication policy that is automatically created when you create a replication file system. To change a parameter's value, click the **Inherit** button so that it changes to **Override**, and then enter the new value in the field that appears.

## The Deduplication Tab

The **Deduplication** tab enables you to enter parameters for data deduplication and contains the following fields and buttons:

Parameter	Description
<b>Deduplication</b>	Click the button beside this field to turn deduplication on or off.

Parameter	Description
<b>Address for Replication and Deduplication</b>	Enter the address to use for replication to and from this host. In HA configurations this address must match the vIP address. In non-HA configurations the address may be left as localhost.  <b>Note:</b> This field appears only when editing the Replication / Deduplication policy named "Global." To change the virtual IP address from "localhost" you must edit the "Global" policy. Additionally, you must edit the "Global" policy on every file system that has a replication/deduplication policy.
<b>Minimum File Idle Time Before Deduplication</b>	Specify the interval of time for a file to remain idle before deduplication begins. This field uses the format 00:00:00:00, which refers to days:hours:minutes:seconds.
<b>Minimum File Size to Deduplicate</b>	Specify the minimum size a file must be in order to be eligible for deduplication.
<b>Filenames Excluded from Deduplication</b>	Specify any filenames you want excluded from the deduplication process.
<b>Relative Deduplication Priority</b>	If desired, indicate relative deduplication priority. Preference is given to files from policies with a lower priority number when StorNext considers deduplication candidates.
<b>Metadata Content Filtering</b>	When metadata content filtering is enabled, the Deduplication engine will attempt to interpret the content of various backup file formats such as tar and netbackup to extract the component files and apply deduplication algorithms to them. Only use filtering if you are storing this type of file in the policy.

The **Deduplication** tab also contains fields which enable you to enter parameters related to data truncation as it applies to deduplication.

**Note:** If a Deduplication-enabled directory is also a Storage Manager relation point, both StorNext and Storage Manager truncation concepts apply. In this case Storage Manager truncation performs any required truncation, but Storage Manager will not truncate a file unless the StorNext truncation parameters have also been satisfied.

The **Truncation** portion of the **Deduplication** tab contains the following fields:

Parameter	Description
<b>Truncation</b>	Click the button beside this field to turn truncation on or off.

Parameter	Description
<b>Minimum File Idle Time before Truncation</b>	Specify the interval of time for a file to remain idle before truncation begins.
<b>Minimum File Size to Truncate</b>	Specify the minimum size a file must be in order to be eligible for truncation.
<b>Files Excluded from Truncation</b>	Specify any filenames you want excluded from the truncation process.
<b>Truncation Low Water Mark *</b>	Watermarks help you determine disk space thresholds for your file system. These thresholds determine the point at which StorNext applies or stops applying truncation. At this field, enter the percentage of occupied disk space a file system must reach before StorNext stops applying truncation.
<b>Truncation High Water Mark *</b>	Enter the percentage of occupied disk space a file system must reach before StorNext applies or starts truncation.
<b>Stub File Size</b>	The stub file is the readable portion of the file that remains after truncation. Enter the desired target size to allocate for the truncated file stub.

*\* These fields do not appear when adding a replication storage policy, but are editable when editing the Global replication policy.*

## The Outbound Replication Tab

The **Outbound Replication** tab enables you to enter parameters related to outgoing replicated data copied to a destination target and contains the following parameters. For all parameters you can either accept the default values by checking the box to the right of the parameter, or uncheck the box to manually enter the value.

**i Note:** This tab displays replication target machines and files systems that have already been configured on the **Storage Destinations > Replication Targets** page. You cannot enter replication targets on the Outbound Replication tab; you can only select previously configured replication targets. For more information about creating replication targets, see [Configure Data Replication on page 95](#).

Parameter	Description
<b>Outbound Replication</b>	Click the button beside this field to turn outbound replication on or off. In this context "On" means that directories associated with this policy will be replicated.

Parameter	Description
<b>Pathname on Target</b>	This is the pathname on the target to which replicated data is copied.
<b>Copies to Keep on Target</b>	This is the number of replicated copies created on the target system.
<b>Replicate Deduplicated Content</b>	This parameter determines whether deduplicated data is included in the replication process.
<b>Use Compression on Network</b>	This parameter determines whether data compression is used prior to and during replication.
<b>Use Encryption on Network</b>	This parameter determines whether data encryption is used prior to and during replication.
<b>Generate Completion Reports</b>	This parameter determines whether a replication report is generated.
<b>Replicate Inline Size</b>	If you specify a size, files smaller than the size you specified will have their contents replicated inline with metadata and will not be transferred separately.
<b>'Name' Passed to Target</b>	When entered, a string used to expand the %N parameter in the realization string on the target. StorNext defaults to the source directory path.
<b>Filenames Excluded from Replication</b>	Specify any filenames you want excluded from the replication process. This field works the same way as a UNIX shell which lets you pattern match names. For example, entering *.0 core would exclude all .o files and also files named "core." You could also skip all core files by entering rep_skip=core*.
<b>Relative Replication Priority</b>	When you enter a priority, preference is given to replicating contents of files from policies with lower priority.

## Add a New Replication Schedule

When deciding whether to create a replication schedule, consider the following:

- This scheduling function schedules the “namespace realization” phase of replication, not replicated data movement. (Namespace realization refers to the process of creating the directory structure for replicated data.)
- If no replication schedule is specified, namespace replications are triggered manually (on demand).

1. Under **Replication Schedules**, click **New**.
2. Under **Replication Targets**, select at least one replication target to which the new schedule should be applied.
3. Enter the parameters for the new replication schedule. Your entries in the **Weekday**, **Day**, **Hour**, and **Minute** columns determine precisely when replication occurs. You must either select at least one entry in each heading column, or else click **All**.

To quickly select recurrence, you can click one of the six small boxes under the **Use** button:

- **Y**: Yearly
- **M**: Monthly
- **W**: Weekly
- **D**: Daily
- **H**: Hourly
- **B**: Business Days (Monday through Friday every week)

---

**i Note:** Do not check **All** under the **Minute** heading. Doing so can trigger a replication each minute, which is a lot of unnecessary overhead.

4. **(Optional)** In the **CRON Spec** field, input a CRON-like schedule in CRON form by entering an allowable string, and then click **Use**.
  - **@yearly**: Run once a year
  - **@annually**: Same as **@yearly**
  - **@monthly**: Run once a month
  - **@weekly**: Run once a week
  - **@daily**: Run once a day,
  - **@midnight**: Same as **@daily**
  - **@hourly**: Run once an hour

After you enter the string, the CRON format will appear at the field, if applicable. For example, if you enter **@daily, 0 0 \* \* \*** appears.

5. **(Optional)** For the **Specials** heading, check **@Reboot** to run the schedule once, at startup.
6. Click **Continue**.
7. Click **Apply** to save the new schedule, or click **Cancel** to exit without saving.
8. When the confirmation message appears, click **Yes** to proceed or **No** to abort.
9. When a message informs you that the replication schedule was successfully created, click **OK** to continue.

## Edit a Replication Schedule

On the right side of the page is a list of previously configured replication schedules.

1. Select the replication schedule you want to modify.
2. Click **Edit**. A page similar to the one on which you added the schedule appears.
3. Modify entries under the **Weekday**, **Day**, **Hour**, and **Minute** columns as desired. (Alternatively, you can enter a CRON-like schedule in CRON format at the **CRON Spec** field, and then click **Use**.)
4. Click **Continue**.
5. When you are finished making changes, click **Apply** to save your modifications, or click **Cancel** to exit without saving.
6. When a message informs you that the replication schedule was successfully modified, click **OK** to continue.

## Delete a Replication Schedule

1. Select the replication schedule you want to delete.
2. Click **Delete**.
3. When the confirmation message appears, click **Yes** to proceed or **No** to abort.
4. When a message informs you that the replication schedule was successfully deleted, click **OK** to continue.

## The Inbound Replication Tab

The **Inbound Replication** tab enables you to enter parameters related to incoming replicated data sent from the source blockpool server and contains the following fields:

Parameter	Description
Inbound Replication	Turn inbound replication on or off, or disable the process. Quantum recommends not enabling inbound replication when creating new policies, but rather enabling it in the "Target" policy.
Pathname on Target	Specify the name of the directory on the target system to which replicated data is saved.
Copies to Keep on Target	Specify the number of copies of replicated data to store on the source system. Quantum recommends not specifying this number on the target policy, but specifying it in the source policy.

**i Note:** See [Additional Replication and Deduplication Information on page 521](#) for additional information about setting up replication and recommended settings for these parameters. If these values are entered for both the source and target policies, the values entered in the target policy are used.



## The Source Directories Tab

The **Source Directories** tab enables you to select or create the file system directory used for the policy and contains the following fields and buttons:

Parameter	Description
<b>Directory</b>	Select from the pulldown list an existing directory you want to use for the policy.
<b>Create Directory Button</b>	To create a new directory, enter a directory location at the field to the left of this button, and then click Create Directory to create the specified directory.
<b>Add Button</b>	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.

To delete a directory on this tab, select the desired directory from the **Source Directories** list and then click **Delete**. Confirm the deletion by answering **Yes** to the confirmation prompt.

## The Blackout Tab

The **Blackout** tab allows you to schedule times and/or days when you do not want deduplication or replication to run. Specifically, the blackout period specifies periods during which the background data transfers for replication (replication phase 1) do not run. A blackout overrides the replication schedule when there is overlap or conflict. The process for scheduling blackouts is the same for deduplication or replication.

**i Note:** The Blackout tab is not shown when you add a replication storage policy, but it appears when you edit the Global replication policy.

1. Select **Deduplication Blackout Window** and/or **Replication Blackout Window** to display scheduling options.
2. Specify the weekday(s), month(s), day(s), hour(s) and minute(s) when you do not want deduplication and/or replication to run. (For each of these items you can click **All**.) Alternatively, you can enter a CRON-like schedule in CRON format at the CRON Spec field, and then click **Use**.
3. When you are finished, click **Apply** to save your changes or **Cancel** to abort.

## Client-side Encryption

The **Tools > Storage Manager > Client-side Encryption** page lists the master keys that can be used for client side encryption.

## Information on the Client-side Encryption Page

Parameter	Description
<b>Master Key</b>	Displays the name of the master key.
<b>Creation Time</b>	Displays the time when the selected master key was created.
<b>Modified Time</b>	Displays the time when the selected master key was last modified.
<b>Number of Instances</b>	Displays when the selected master key is edited, or a new instance of the master key is created. This column displays the number of instances that exist for the selected master key.
<b>New...</b>	Click to create a new master key.
<b>Edit...</b>	Click to edit an existing master key.
<b>Refresh</b>	Click to refresh the data on this page.
<b>Select Action</b>	Click to display a list of available drop-down options. The available options are: <ul style="list-style-type: none"><li>• <b>Change Data Protection Key:</b> This action will internally change the data protection key.</li></ul>

## Create a Master Key

1. Click **New...** The **Tools > Storage Manager > Client-side Encryption > New** page appears.
2. In the **Master Key Store ID** field, input the master key store identification number. This is a mandatory field limited to 15 characters. Only alphanumeric characters and the special characters - (dash), \_ (underscore), and . (period) are allowed.

**Note:** This field only appears when you create master keys for the first time.

3. In the **Master Key Name** field, input the master key name. This is a mandatory field limited to 31 characters. Only alphanumeric characters and the special characters \_ (underscore), - (dash), . (period), and ! (exclamation mark) are allowed.
4. In the **Master Key Passphrase** field, input the Master Key passphrase. This is a mandatory field limited to 127 characters. This passphrase must be remembered or be kept in safe place. It is required for updating the passphrase.
5. **(Optional)** Click the **Show Passphrase** box to display the passphrase text. Since the passphrase text is secret text, it is masked by default.
6. Click **Apply** to create the new master key, or click **Cancel** to discard your changes and exit without creating the master key.

## Edit a Master Key

1. On the **Tools > Storage Manager > Client-side Encryption** page, select a **Master Key**.
2. Click **Edit...**. The **Tools > Storage Manager > Client-side Encryption > <your key>** page appears.

---

**i Note:** The **Master Key Name** field displays the master key name. The name cannot be edited.

3. In the **Current Master Key Passphrase** field, input the **Current Master Key** passphrase. This is a mandatory field.
4. In the **New Master Key Passphrase** field, input the **New Master Key** passphrase. This is a mandatory field limited to 127 characters. This passphrase must be remembered or be kept in safe place. It is required for updating the passphrase.
5. **(Optional)** Click the **Show Passphrase** box to display the passphrase text. Since passphrase text is secret text, it is masked by default.
6. Click **Apply** to edit the master key, or click **Cancel** to discard your changes and exit without editing the master key.

For additional information, see [Client Side Compression and Encryption below](#).

## Client Side Compression and Encryption

Beginning with StorNext 5 release 5.3, support for Encryption and Compression with Q-Cloud policies has been added. Both client side compression and client side encryption are enforced and configured as a Policy class attribute. To enable client side encryption for a policy class, a master key must be selected. If a master key does not exist, create a master key first. Master keys are created and managed by the command **fskey**. Refer to the *StorNext 5 MAN Pages Reference Guide* for the various policy class commands.

For example:

- **fsaddclass**
- **fsmodclass**
- **fsrclass**
- **fsclassinfo**
- **fskey**

## The fskey Command

The command **fskey** adds, modifies and reports master keys used in the client-side encryption feature in the Quantum storage system. The command can also be used to generate a new data protection key associated with a specific master key. A data protection key (DPK) is used to encrypt data content before it is uploaded to an Object Storage when the client-side encryption is enabled, while master keys are used to wrap (encrypt) data protection keys.

A master key's content is derived from a user-supplied passphrase. Each master key has a unique name. This unique key name can be assigned to a particular policy if the client-side encryption feature is enabled for this policy. The key content of a master key can be changed by providing a new passphrase. In this case, a new master key instance is created. The old instance is then removed after all data protection keys wrapped by the old instance are rewrapped by the new instance.

For additional information, see [Client-side Encryption on page 112](#).

## The `qcloud_audit` Utility

The `qcloud_audit` utility generates a peer device key file containing CSV list of files stored to Q-Cloud in the given output directory. Execute the command `qcloud_audit -h` to display a list of arguments and the usage description of the tool.

The CSV list contains the following for each file stored to Q-Cloud:

Parameter	Description
<b>Path</b>	Displays the path from the relation point to the file.
<b>Name</b>	Displays the name of the file.
<b>Owner</b>	Displays the ID of the file owner.
<b>Entype</b>	Displays the encryption type: <ul style="list-style-type: none"><li>• 0 for none</li><li>• 1 for server</li><li>• 2 for client</li></ul>
<b>Namespace</b>	Displays the name of the bucket where the file is stored.
<b>Objid</b>	Displays the object ID of the file.
<b>Modtime</b>	Displays the time-stamp of the file's last modification in the form: mm-dd-yyyy:hh:mm:ss
<b>Addtime</b>	Displays the time-stamp when the Q-Cloud copy batch was completed in the form: mm-dd-yyyy:hh:mm:ss

For example, execute the following command:

```
qcloud_audit -o /var/tmp
```

The command generates a CSV file titled "**Qcloud\_1.audit**" (assume the device key is 1). The content of the file contains information similar to the following:

```
path1, file1, 123, 0, bucket1, 000001, 07-09-2015:19:24:13, 07-09-2015:19:24:54  
path2, file2, 456, 1, bucket1, 000002, 07-09-2015:19:24:13, 07-09-2015:19:24:54  
path3, file3, 789, 2, bucket1, 000003, 07-09-2015:19:24:13, 07-09-2015:19:24:54
```

---

**i Note:** The time required to complete the command is directly proportional to the number of files stored to Q-Cloud.

**i Note:** The output of the Q-cloud audit log can reach up to 6GB per million copies stored. Ensure your system contains sufficient storage space before running the audit process.

## The Compression and Encryption Usage Report

Beginning with StorNext 5 release 5.3, with compression and encryption for Q-Cloud devices, you can request to view compression and encryption usage information. The compression and encryption usage information is reported by the command **fsobjinfo**.

### The fsobjinfo Command

Execute the command **fsobjinfo** to generate the compression and encryption usage report. The command **fsobjinfo** produces a summary usage report for object store media. Object store usage is summarized based on object store media ID and policy class ID. Reported usage can be limited to the optionally specified set of policy class IDs or object store media IDs.

### The qcloud\_migrate.pl Command

In order to generate accurate reports, existing Q-Cloud Archive usage must be accounted for. The usage information in the `filecomp` table must be populated in the `filecomp_obj` and `classobj_info` tables. The command **qcloud\_migrate.pl** provides the capability; execute the command **qcloud\_migrate.pl** after you upgrade your system to StorNext 5 release 5.3 (or later).

The full path of the command is: `/usr/adic/TSM/util/install/qcloud_migrate.pl`

---

**i Note:** Your system will operate normally without execution of the **qcloud\_migrate.pl** command.

### Considerations for the qcloud\_migrate.pl Command

- You can execute the command anytime after an upgrade to StorNext 5 release 5.3 (or later).
- Execute the command **qcloud\_migrate.pl** only once.
- The compression and encryption usage report may not be accurate if your system contains existing Q-Cloud Archive devices.
- If your system does not contain existing Q-Cloud Archive devices, do not execute the command **qcloud\_migrate.pl**.

## The Compression and Encryption Usage Report

Beginning with StorNext 5 release 5.3, with compression and encryption for Q-Cloud devices, you can request to view compression and encryption usage information. The compression and encryption usage information is reported by the command **fsobjinfo**.

### The **fsobjinfo** Command

Execute the command **fsobjinfo** to generate the compression and encryption usage report. The command **fsobjinfo** produces a summary usage report for object store media. Object store usage is summarized based on object store media ID and policy class ID. Reported usage can be limited to the optionally specified set of policy class IDs or object store media IDs.

### The **qcloud\_migrate.pl** Command

In order to generate accurate reports, existing Q-Cloud Archive usage must be accounted for. The usage information in the `filecomp` table must be populated in the `filecomp_obj` and `classobj_info` tables. The command **qcloud\_migrate.pl** provides the capability; execute the command **qcloud\_migrate.pl** after you upgrade your system to StorNext 5 release 5.3 (or later).

The full path of the command is: `/usr/adic/TSM/util/install/qcloud_migrate.pl`

---

**i Note:** Your system will operate normally without execution of the **qcloud\_migrate.pl** command.

### Considerations for the **qcloud\_migrate.pl** Command

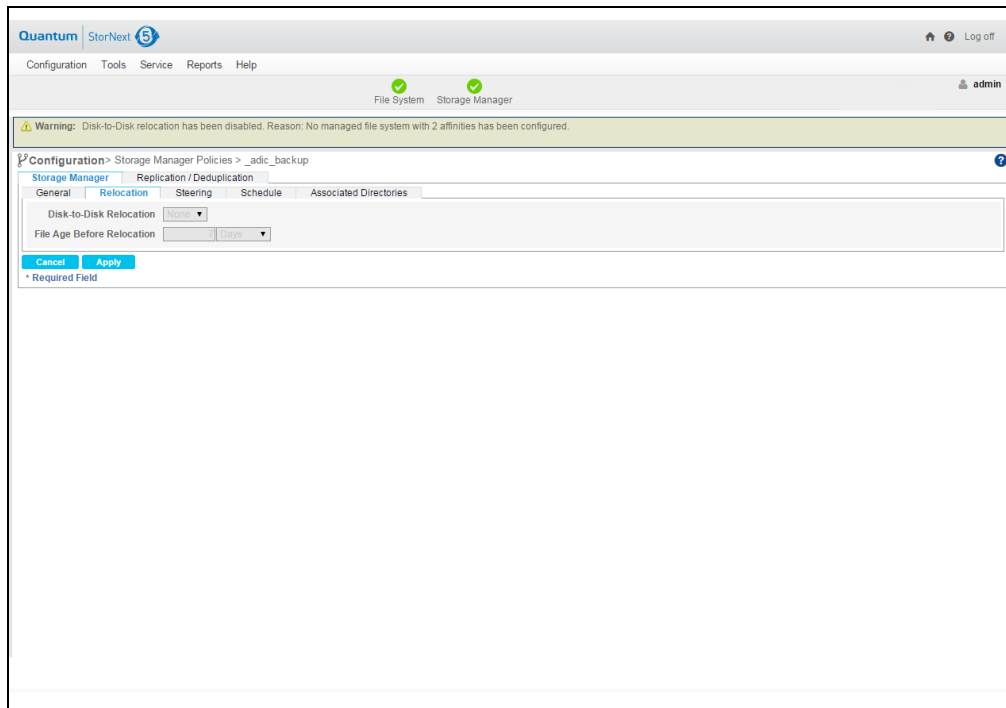
- You can execute the command anytime after an upgrade to StorNext 5 release 5.3 (or later).
- Execute the command **qcloud\_migrate.pl** only once.
- The compression and encryption usage report may not be accurate if your system contains existing Q-Cloud Archive devices.
- If your system does not contain existing Q-Cloud Archive devices, do not execute the command **qcloud\_migrate.pl**.

### 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 8:** 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.

## Linear Tape File System (LTFS) Media Format

Beginning with StorNext 5, users with LTO5/LTO6 tape media can choose to store archive copies to the traditional native StorNext tape format (ANTF) and/or to the new Open LTFS tape format.

**i Note:** LTO7 tape media may be used with StorNext 5 release 5.4.0 (and later).

The new LTFS tape format offers most of the same StorNext 5 functionality as the existing ANTF format with the additional benefit of being portable. That is, tape media written in LTFS format by StorNext 5 can be vaulted and removed from the StorNext Storage Manager system and mounted elsewhere using the freely available Open LTFS package, allowing access to the data without requiring any additional software from Quantum.

The StorNext 5 implementation for LTFS is based on the Open LTFS Format Specification v2.0.0 and is compatible with implementations from other vendors at or above this version. The initial StorNext 5 implementation only supports LTFS on LTO5 and LTO6 tape media.

---

**i Note:** StorNext 5 release 5.4.0 (and later) uses Open LTFS Format Specification v2.2.1 and supports LTO7 tape media.

## Supported Features

StorNext 5 for LTFS includes support for the following features:

- Policy-class-copy steering by media format type for LTO tape media
- Automatic media selection and formatting
- Distributed data-mover support
- File versioning on LTFS media
- File recovery from LTFS media
- LTFS media defragmentation
- Tape-to-tape copy from ANTF or LTFS to ANTF or LTFS formats
- LTO7 media with StorNext 5 release 5.4.0 (or later)

## Features Not Supported

Features NOT supported in StorNext 5 for LTFS:

- Segmented files
- LTFS volume spanning
- StorNext backup to LTFS
- Import of native LTFS media into StorNext
- Export/import of StorNext managed LTFS media
- LTO7 media prior to StorNext 5 release 5.4.0

## Accessing StorNext-generated LTFS Media

Access of StorNext-generated LTFS media outside of StorNext (for example, by an Open LTFS package,) should be considered read-only.

A StorNext-generated LTFS medium may be rendered unusable within StorNext 5 if it is modified outside of StorNext 5. If the medium is required for data retrieval of truncated files within the StorNext file system, access to those files may not be possible if the medium is modified outside of StorNext.

## Configuration Considerations

Now that two different media formats (ANTF/LTFS) are supported for LTO media in StorNext 5, certain items should be considered for configuration and usage so that StorNext users can make the best choices for their environment.



Following is a list of these items to consider:

- StorNext Storage Manager LTFS support is based on LTFS Format Specification v2.0.0. The LTFS Format Specification does not support backwards compatibility with earlier versions of the specification.
- The LTFS Format Specifications v2.0.0 and v2.2.1 are compatible with each other. If a tape is formatted with version v2.0.0 and is written by StorNext 5 release 5.4.0, the tape format will be upgraded to v2.2.1. Similarly, a tape formatted with version v2.2.1 will be downgraded to a v2.0.0 if it is written by StorNext 5 prior to release 5.4.0.
- Neither LTFS Format v2.0.0 or v2.2.1 provide a means for maintaining file ownership and permissions when storing files to LTFS media. Instead, this is managed at LTFS file system mount time using options passed to the Open LTFS package commands. When used within StorNext Storage manager, access is restricted to user root and is not accessible outside of the StorNext Storage Manager processes. When mounted outside of StorNext Storage Manager, care should be taken to ensure proper access permissions are used when the LTFS volume is mounted.
- Storage Manager does not support segmented files LTFS. This includes files being stored as well as files being copied from tape to tape (fsfilecopy/fsmedcopy). As a means to manage this, the configuration parameter `DEF_MED_SPC_LTO_LTFS` in file `/usr/adic/TSM/config/fs_sysparm` specifies the maximum file size that can be stored to LTFS media.
- StorNext 5 supports a maximum tape block size of 1MB for LTFS. As a means to manage this, the configuration parameter `FS_LTO_LTFS_BLOCK_FACTOR` in file `/usr/adic/TSM/config/fs_sysparm` is used to calculate the LTFS tape block size and limit it to 1MB for LTFS formatted media.
- LTFS allocates at least one tape block for each file, and no two files share the same tape block. Therefore, every file consumes an amount of space on tape that is a multiple of the tape block size. As a consequence, very small files do not make efficient use of tape capacity, especially when large tape block sizes are configured.
- StorNext import of native LTFS media is not supported (that is, media that has been formatted and written to outside of StorNext Storage Manager).
- Also, StorNext import of StorNext LTFS media is not supported between different installations of Storage Manager.
- StorNext Backup(`_adic_backup policy class`) is not supported for LTFS.
- Since LTFS is a file system on a tape, additional overhead is incurred when Storage Manager processes have to mount and unmount the LTFS file system. Because of this, ANTF will outperform LTFS in situations where small numbers of files and/or small amounts of data need to be stored, retrieved, or copied. Therefore, StorNext users will NOT want to configure LTFS for the primary copy used for file retrieval (for example, copy 1).
- In order to minimize this additional overhead, system parameters should be configured to process high file counts and high byte counts for store and retrieve requests. The configuration parameter `MAX_FILES_PER_CLUSTER` in file `/usr/adic/TSM/config/fs_sysparm` should remain at the default of 3000 and not be reduced. The undocumented configuration parameter `FS_CLUSTER_LIMIT_LTO` in file `/usr/adic/TSM/config/fs_sysparm` should be increased to 1,000,000,000,000 (1TB) for LTO5 media, 2,000,000,000,000 (2TB) for LTO6 media, and 6,000,000,000,000 (6TB) for LTO7 media. The `FS_CLUSTER_LIMIT_LTO` parameter should be set to the lowest recommended LTO generation value in environments where multiple LTO generations are being used with LTFS.
- The amount of meta data stored in the LTFS index partition grows as files are added to an LTFS volume. This meta data is always loaded into the `fs_fmover` process memory when the LTFS volume is mounted,

and continues to grow as files are added to the LTFS volume. Therefore, it is expected behavior that the `fs_fmover` process memory requirements are larger when the source and/or destination media format type is LTFS as compared to when the source and destination media format types are both ANTF.

- The LTFS standard prohibits the use of the following characters in either directory or file names:

: /

Attempting to store files to LTFS that have prohibited characters in either directory or file names in the path will fail and cause an admin alert to be generated. These files will then be removed from the store policy candidate list by having the `FS_NO_STORE` flag set in the file attributes. In order to successfully store these files, they must be renamed such that they no longer contain any prohibited characters in any of the directory or file names in the path. Once renamed, the `FS_NO_STORE` attribute flag must then be cleared using the `fschfiat -sp` command or the **Tools > File and Directory Actions > Modify File Attributes** GUI page in order for policy to place it back on the store candidate list. Future file modifications will be ignored and no additional copies will be made regardless of media type or format until the `FS_NO_STORE` attribute flag had been cleared. It is also strongly recommended to avoid using certain characters for interoperability between operating systems. The following characters should NOT be used in either directory or file names if the files are to be stored to LTFS media:

\* ? < > " | \

While not prohibited on Linux or OS X, these characters cannot be used on Windows because of additional restrictions for characters used in directory and file names.

In the **Steering** tab in the Storage Manager Policies configuration page, you can select from the Media Format list, either ANTF or LTFS formats for the **LTO (tape) Media Type**.

---

**i Note:** This only applies to tape formats; there are no Media Format selections available for either Lattus or Storage Disks.

LTFS appears in the **Configuration > Storage Manager Policies > Storage Manager** tab (in the **Steering** column), and in the **View** page (in the **Media Format** column).

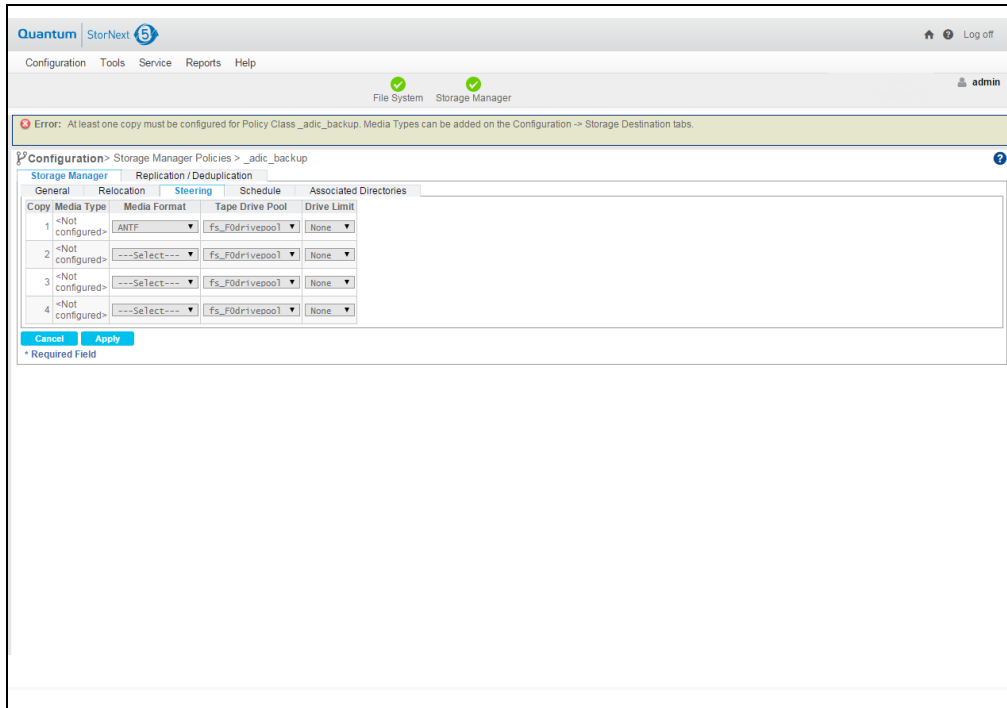
The LTFS media format also appears in the **Reports > Media** page (in the **Media Format** column).

Additional checks have been included in the StorNext Linux-based installation to ensure that some additional RPM packages required to support the LTFS media format are installed prior to installing StorNext. For LTFS requirements, see the *StorNext Compatibility Guide*, the *StorNext Installation Guide*, and the *StorNext Upgrade Guide*.

## 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 9 on the next page](#)). 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 9: Storage Policies > New > Steering Tab



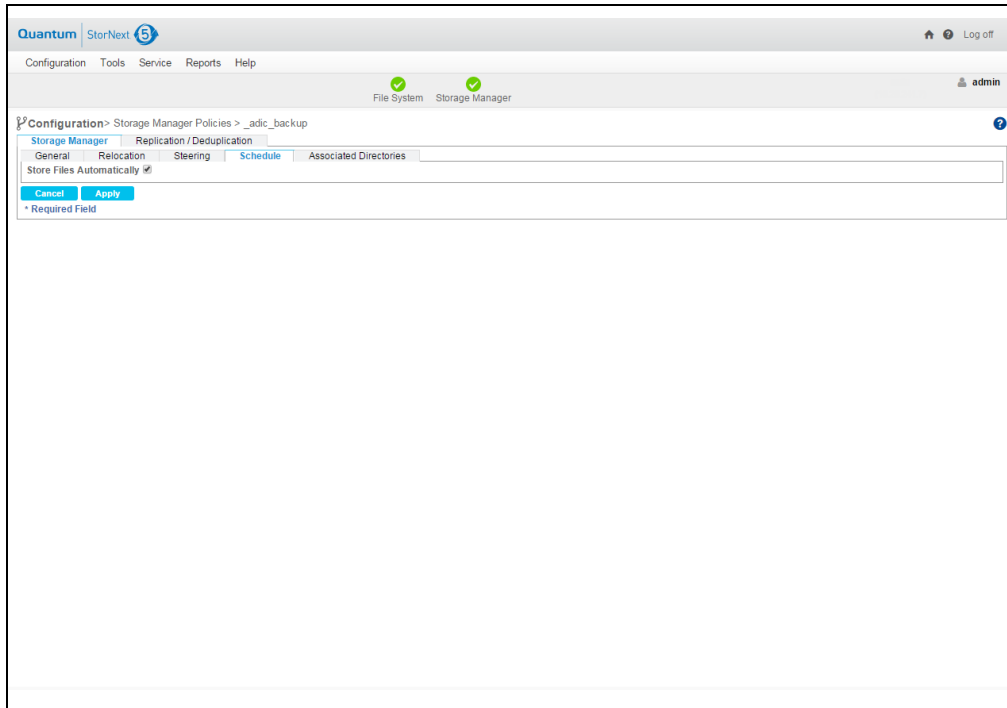
For instructions on what to enter on this screen, see the online help.

## 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 180](#).

**Figure 10:** 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 11:** Storage Policies > New > Associated Directories Tab



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

---

## Email Server

The **Email Server** option allows you to specify the email server used for processing StorNext notification email messages. On this page 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.

**i 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.

### Configure the Email Server

1. Select **Email Server** from the **Configuration** menu.
2. Complete the following fields related to your email system configuration on the **Configuration > Email Server** page. Required fields are marked with an asterisk (\*).

Parameter	Description
<b>SMTP Server</b>	Enter the identification for the server that stores and processes your email account information. This might be a valid server name or an IP address.
<b>Verify SMTP Connectivity</b>	Enable this option if you want StorNext to verify that the SMTP server you entered is valid and reachable.
<b>Authentication</b>	If your email provider requires a password upon sign on, select the <b>PASSWORD</b> option. Otherwise, select <b>NONE</b> . When you select the <b>PASSWORD</b> option, the following fields appear: <ul style="list-style-type: none"><li>• <b>Account:</b> Enter a valid email account for outgoing email messages.</li><li>• <b>Password:</b> Enter the email account's sign-on password, if required.</li></ul>
<b>Sender Address</b>	Enter the email address for the entity responsible for sending alert messages to designated recipients.
<b>Send Test Email</b>	Enter an email address to which you can send test messages in order to confirm successful configuration.

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

---

## Email Notifications

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

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

### Add an Email Recipient

1. When the Configuration Wizard is displayed, select **Email Notifications** on the left side of the page. Alternatively, select **Email Notifications** from the **Configuration** menu. The **Configuration > Email Notification** page appears.
2. On the **Configuration > Email Notifications** page, click **New**. The **Configuration > Email Notifications > New** page appears.
3. Enter the following fields for the new email recipient:

- **Contact Name:** Enter the name of the person who should receive email notifications.
  - **Admin Alerts:** Select this option to send an email message to the recipient whenever an admin alert is generated. For more information about Admin Alerts, see [Admin Alerts on page 337](#).
  - **StorNext Backups:** Select this option to send an email message to the recipient after a backup has occurred on your system.
  - **Service Tickets:** Specify the way you want StorNext to handle email notifications whenever a service ticket for your system is generated. Choose one of the following:
    - **Disabled:** No email notification messages are sent for service tickets.
    - **Alert Level High:** An email notification message is sent only when the alert level is **High**.
    - **Alert Level Middle:** An email notification message is sent when the alert level is **Middle** and higher.
    - **Alert Level Low:** An email notification message is sent when the alert level is **Low** and higher.
  - **Policy Class:** Select this option to receive email about policy class. You must specify the pertaining policy class.
4. Click **Apply** to save your changes, or click **Cancel** to exit without saving. You can also clear all fields and start over by clicking **Reset**.
  5. When the confirmation message appears, click **Yes** to proceed or **No** to abort.
  6. When a message informs you that the email notification recipient was successfully added, click **OK** to return to the **Configuration > Email Notifications** page.

## View Email Recipient Information

To view details for an existing email recipient:

1. If you have not already done so, select **Email Notifications** from the **Configuration** menu.
2. On the **Configuration > Email Notifications** page, 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** page.

## Edit an Email Recipient

1. If you have not already done so, select **Email Notifications** from the **Configuration** menu.
2. On the **Configuration > Email Notifications** page, select the recipient whose information you want to edit and then click **Edit**.
3. Modify any of the following fields:
  - **Admin Alerts:** Select this option to send an email message to the recipient whenever an admin alert is generated. For more information about Admin Alerts, see [Admin Alerts on page 337](#).

- **StorNext Backups:** Select this option to send an email message to the recipient after a backup has occurred on your system.
  - **Service Tickets:** Specify the way you want StorNext to handle email notifications whenever a service ticket for your system is generated. Choose one of the following:
    - **Disabled:** No email notification messages are sent for service tickets.
    - **Alert Level High:** An email notification message is sent only when the alert level is **High**.
    - **Alert Level Middle:** An email notification message is sent when the alert level is **Middle** and higher.
    - **Alert Level Low:** An email notification message is sent when the alert level is **Low** and higher.
  - **Policy Class:** Select this option to receive email about policy class. You must specify the pertaining policy class.
4. When you are finished making modifications, click **Apply** to save your changes and return to the **Configuration > Email Notifications** page, or click **Cancel** to exit without saving.

## Delete an Email Recipient

To delete a previously entered email recipient:

1. If you have not already done so, select **Email Notifications** from the **Configuration** menu.
2. On the **Configuration > Email Notifications** page, 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** to return to the **Configuration > Email Notifications** page.

## Notify Quantum When Service Tickets are Generated

Select this option to automatically send the Quantum Technical Assistance Center a message whenever a service ticket is generated.

---

# Done

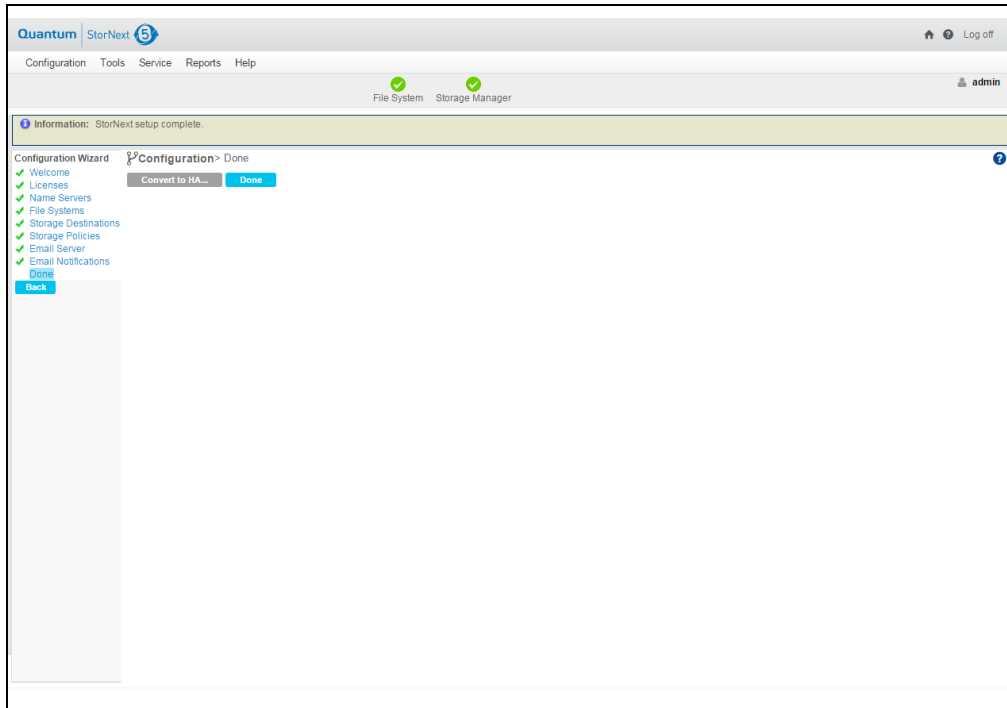
The last step in the Configuration Wizard is to click **Done** to indicate that you have completed all configuration steps.

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 344](#).



Chapter 3: The Configuration Wizard  
Done

Figure 12: Configuration > Configuration Wizard Done Page





# Chapter 4: File System Tasks

---

In addition to the basic file system tasks described for the **Configuration Wizard**, the **File Systems** menu contains additional options that enable you to perform the following file system-related tasks:

Task	Description
<b>Label Disks</b>	Apply EFI or VTOC label names for disk devices in your StorNext libraries
<b>Check File System</b>	Run a check on StorNext files systems prior to expanding or migrating the file system
<b>Affinities</b>	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
<b>Migrate Data</b>	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
<b>Truncation Parameters</b>	Enter truncation parameters for your file systems in order to free up file storage that isn't being actively used
<b>Manage Quotas</b>	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 [Rename a Standalone \(unmanaged\) StorNext File System on page 153](#).

This chapter contains the following topics:

Label Disks .....	130
Understanding Resource Allocation .....	132

Using Resource Allocation From the Command Line .....	133
Check File System .....	137
Affinities .....	139
Migrate Data .....	144
Truncation Parameters .....	147
Manage Quotas .....	148
Rename a Standalone (unmanaged) StorNext File System .....	153

---

## 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.

---

**i 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.

## Label a Device

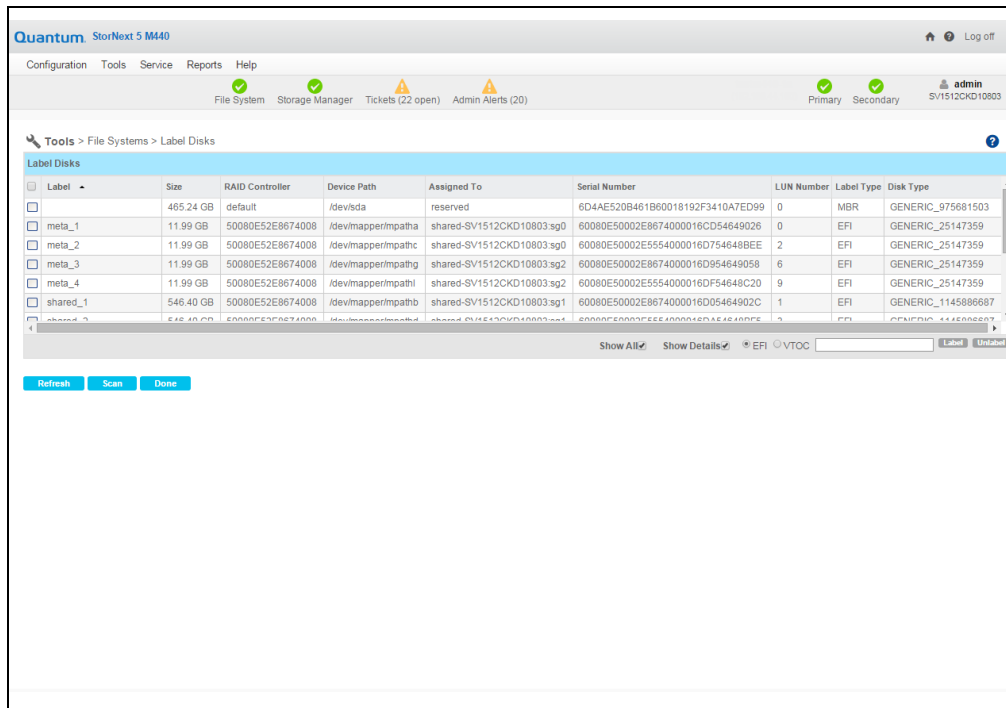
Follow this procedure to label any new or unused devices, or relabel a device that has been unlabeled.

---

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

1. On the **Tools** menu, click **File Systems**, and then click **Label Disks**. The **Tools > Label Disks** page appears.

Figure 13: Label Disks Page



2. (Optional) 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. Optionally, 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. To cancel the operation, click **Cancel**.

**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 [Label a Device on page 130](#).

---

**i 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. On the **Tools** menu, click **File Systems**, and then click **Label Disks**.
2. Select the disk devices from which you want to remove labels. Optionally, 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). To cancel the operation, click **Cancel**.

---

**⚠ 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
- 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.

## 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).

## Expansion and Movement Steps

Below are high-level 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 137](#).
2. Expand the file system. See [File Systems on page 34](#).
3. Move data stripe groups or metadata/journal stripe groups. See [Migrate Data on page 144](#).
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 on the next page](#)
- [Adding and Moving a Data Stripe Group on the next page](#)
- [Moving a Metadata/Journal Stripe Group on page 135](#)

 **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.

---

**⚠ Caution:** Do not use this procedure if the file system was created with the **Windows Simple Config Tool**.

**i Note:** Changes to the configuration file must be manually copied to the peer MDC if you have a pair of Windows MDCs or if your Linux MDCs are not using a dedicated shared HA file system, also known as RPM-only StorNext install.

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 above](#).
4. Add the new stripe groups to the file system.
5. Stop the File System Manager (FSM).
6. Run the `cvupdatefs` command.
7. Restart the FSM.

## Adding and Moving a Data Stripe Group

New functionality has been added to the `snfsdefrag` utility to support operations on multiple stripe groups.

---

**i 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 on the previous page](#).
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 command, **sndiskmove**, only works on disks or LUNs of the same size. For instance, if the destination LUN is smaller than the source LUN, then the command fails. However, if the destination LUN is larger than the source, then the additional capacity is ignored.

---

**⚠ 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 on page 134](#).
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)

---

**i 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.

---

**i 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.

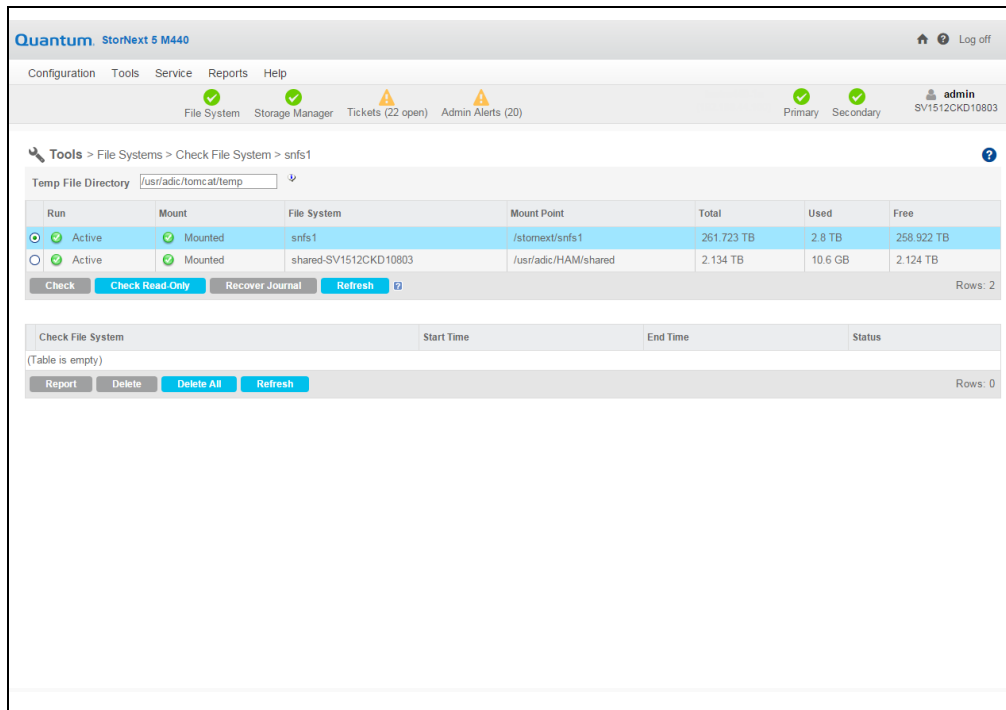
Use the following procedure to perform a file system check.

---

**i 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. On the **Tools** menu, click **File Systems**, and then click **Check File System**. The **Tools > Check > [file system name]** page appears.

Figure 14: Check File System Page

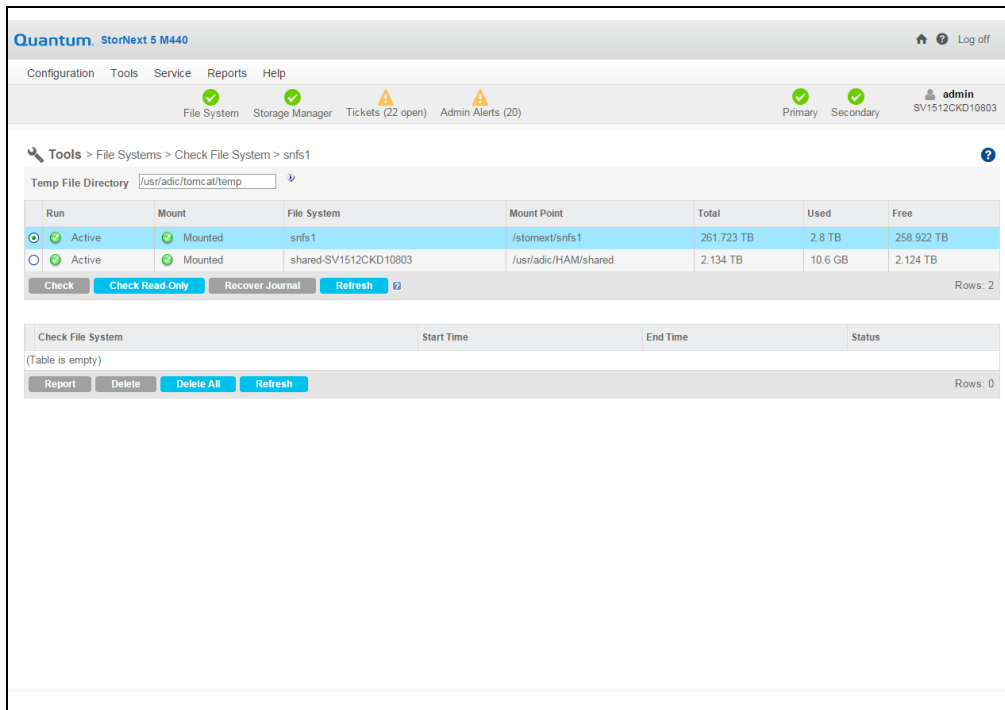


2. At the **Temp File Directory** field, enter a new directory if the specified directory does not have enough space to perform the check. The checking process on large file systems can take hundreds of megabytes or more of local system disk space for working files.
3. Select the file system you want 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 15:** 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:

File	Example
<code>/usr/cvfs/data/&lt;fsname&gt;/trace/cvfsck-&lt;timestamp&gt;</code>	<code>/usr/cvfs/data/snfs1/trace/cvfsck-02_22_2010-12_15_19</code>
<code>/usr/adic/gui/logs/jobs/CHECK_FS-&lt;timestamp&gt;-&lt;jobid&gt;</code>	<code>/usr/adic/gui/logs/jobs/CHECK_FS-20100222_121519-77</code>

## Affinities

This section describes StorNext’s “stripe group affinity” feature, and also provides some common use cases.

Term	Definition
<b>Stripe Group</b>	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.
<b>Affinity</b>	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.
<b>Exclusivity</b>	Defines a stripe group that has one or more affinities and the <code>exclusive</code> attribute set to <code>true</code> , 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.

## Auto Affinities

**Auto Affinities** designate the affinity (stripe group[s]) to which allocations will be targeted for all files on the file system whose name has the specified file extension.

See the *StorNext Online Help* for more details on how to configure **Auto Affinities** using the GUI.

## Affinity Preference

**Affinity Preference** allows files of a particular affinity to have their allocations placed on other available stripe groups (with non-exclusive affinities) when the stripe groups of their assigned affinity do not have sufficient space. Otherwise, allocation attempts will fail with an out-of-space error.

See the *StorNext Online Help* for more details on how to configure **Affinity Preference** using the GUI.

## Allocation Strategy

StorNext has multiple allocation strategies which can be set at the file system level. These strategies control where a new file's first blocks will be allocated. Affinities modify this behavior in two ways:

- A file with an affinity will be allocated only on a stripe group with matching affinity.
- A stripe group with an affinity and the exclusive attribute will be used only for allocations by files with matching affinity.

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

## Example Use Cases

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

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

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

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

To reserve high-speed disk for critical files:

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

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

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

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

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

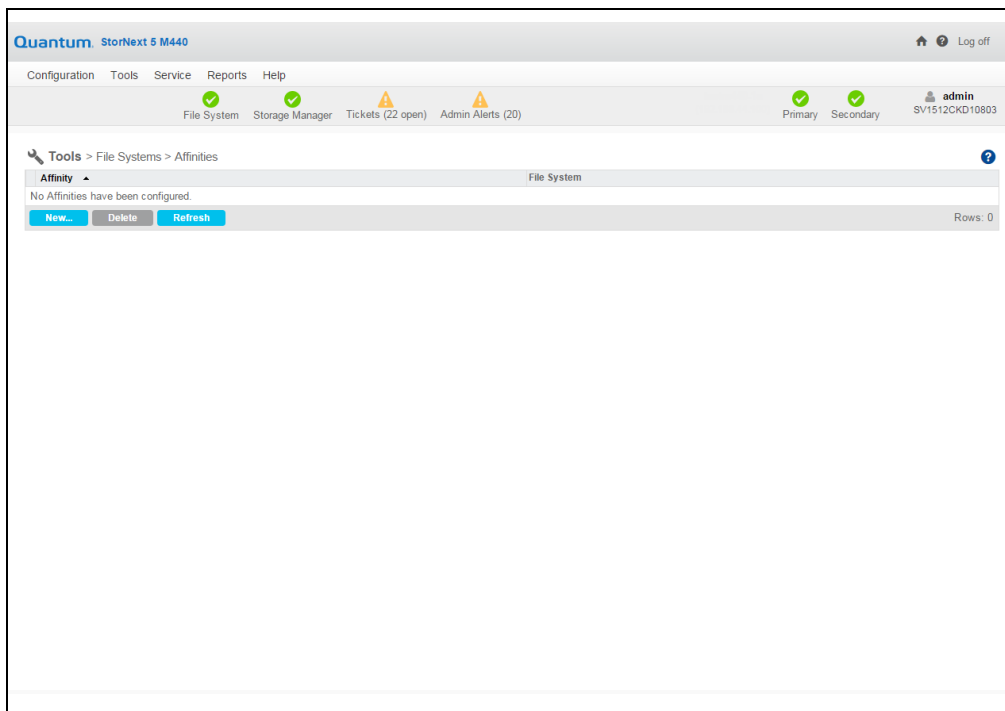
This will not work if critical files may be created later, unless there is a process to move them to the FAST stripe group, or an affinity is set on the critical files by inheritance but removed after their first allocation (to allow them to grow onto non-FAST groups).

## Adding a New Affinity

Follow this procedure to add affinities:

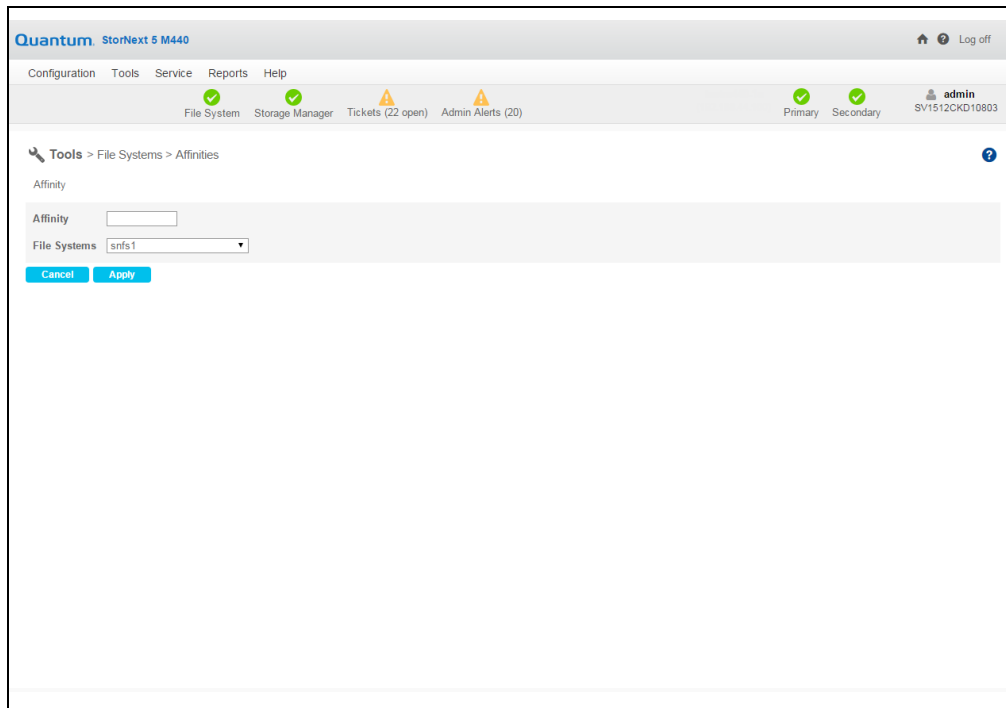
1. On the **Tools** menu, click **File Systems**, and then click **Affinities**. The **Tools > File Systems > Affinities** page appears.

**Figure 16:** Affinities page



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

Figure 17: New Affinity Page



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. On the Tools menu, click File Systems, and then click **Affinities**. The **Tools > File Systems > Affinities** page 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.

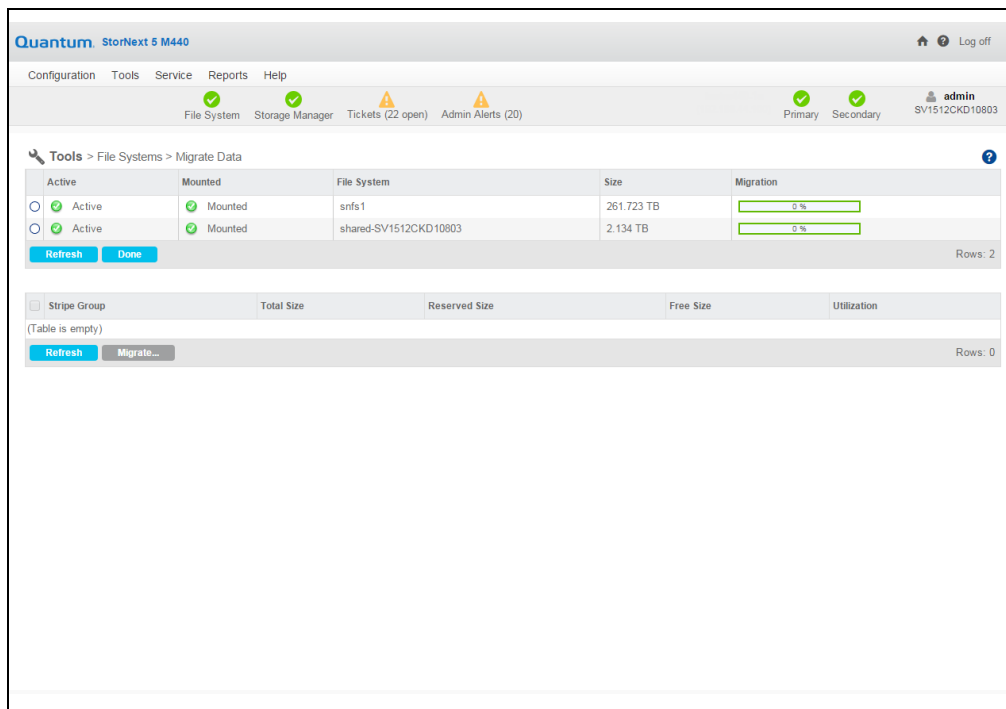


# Migrate Data

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 18 below](#)).

Figure 18: Migrate Data page



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

- To migrate **Metadata/Journal Data** stripe groups, see [Migrate Metadata and Journal Data on the next page](#).
- To migrate **User Data** stripe groups, see [Migrate User Data on page 146](#).

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.

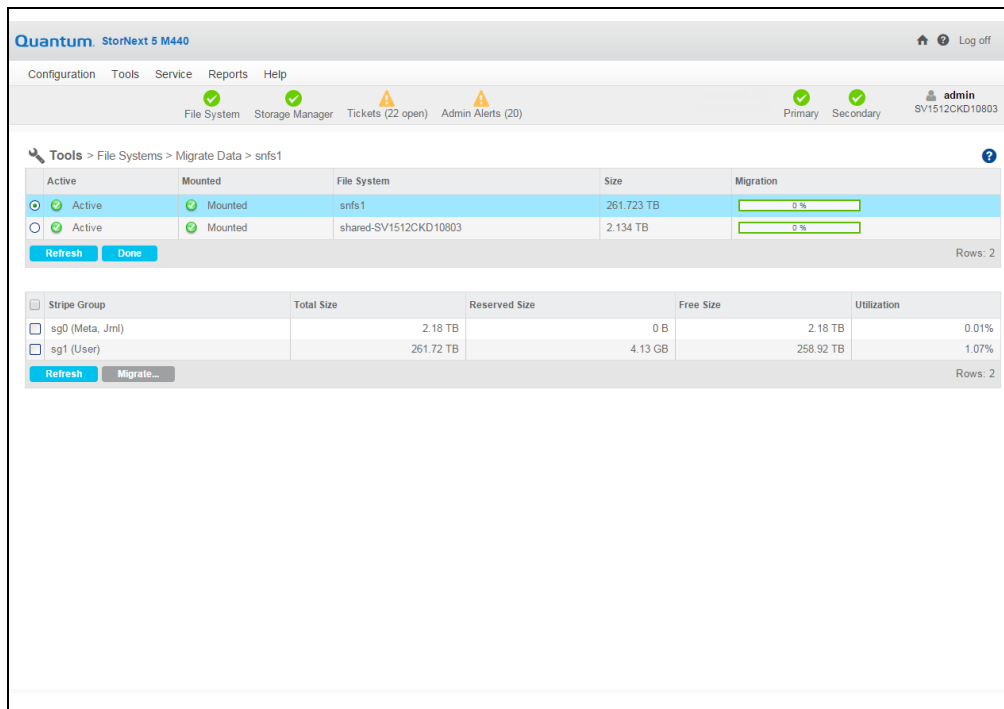
**Note:** When migrating a metadata stripe group to new disks, the overall metadata size will remain the same despite the size of the target disks. In order to add more metadata capacity, add new metadata stripe groups (refer to [Adding a Stripe Group Without Moving on page 134](#)).

## Migrate Metadata and Journal Data

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

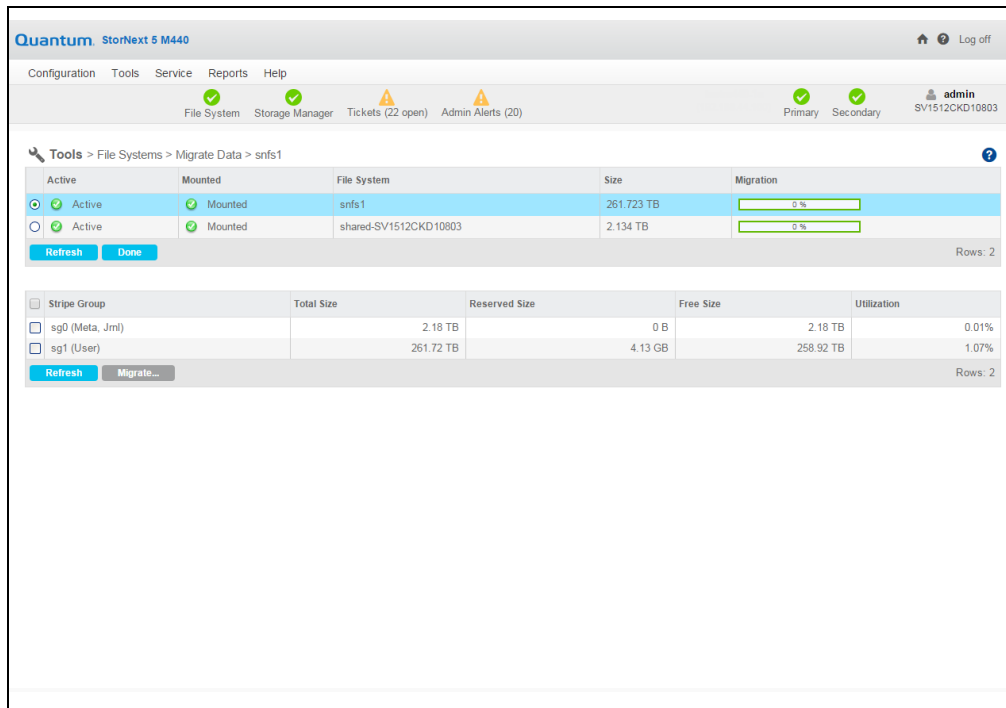
1. On the **Tools** menu, click **File Systems**, and then click **Migrate Data**. The **Tools > File Systems > Migrate Data** page appears (refer to [Figure 19 below](#)).
2. Click a **File System**.
3. Click a **Metadata/Journal Stripe Group** in the **File System** to migrate (refer to [Figure 19 below](#)).

Figure 19: Metadata/Journal Data migration page



4. Click **Migrate....** A new page appears (refer to [Figure 20 on the next page](#)).

Figure 20: Metadata/Journal Data migration (new page)



5. Click a disk from the **Choose Disk to Migrate From** table.
6. Click a disk from the **Choose Disk to Migrate To** table.
7. Click **Migrate**. The following occurs:
  - The file system stops and is unmounted.
  - The GUI executes the command `sndiskmove`.

**Caution:** The command, `sndiskmove`, only works on disks or LUNs of the same size. For instance, if the destination LUN is smaller than the source LUN, then the command fails. However, if the destination LUN is larger than the source, then the additional capacity is ignored.

- The source disk/LUN is relabeled to `$LABEL .old`.
  - The destination disk/LUN is relabeled to `$LABEL`.
8. Click **Refresh** to manually update the status.
  9. When this migration task is completed, start and mount the file system using the GUI.

## Migrate User Data

To migrate stripe groups in a managed file system, you must use the command line interface. For additional information, contact Quantum Technical Support.

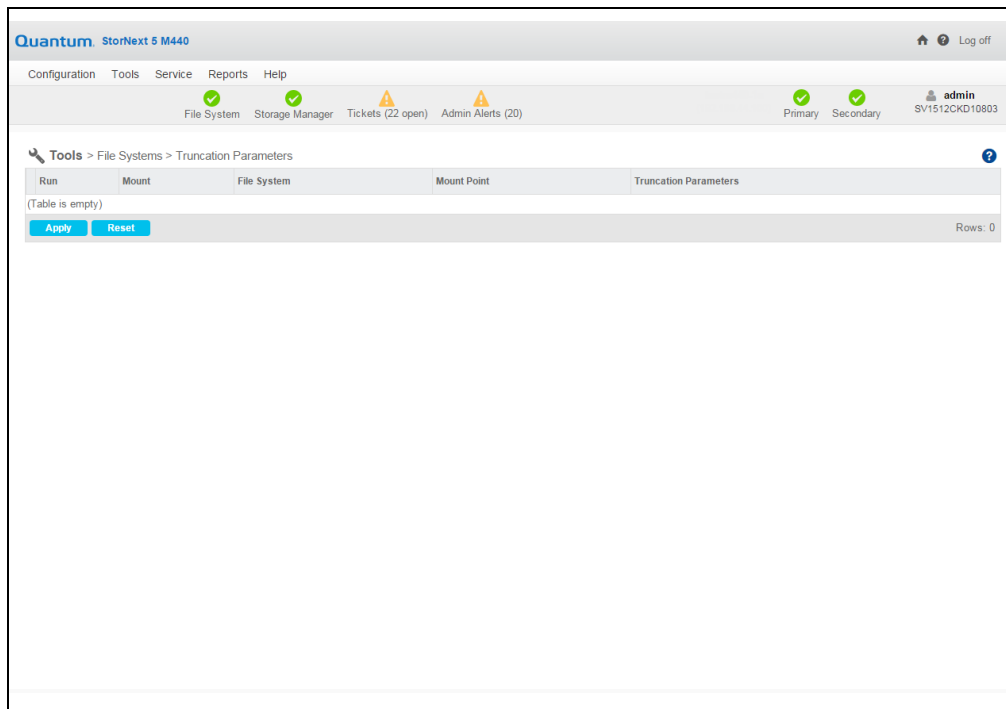
# Truncation Parameters

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

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

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

Figure 21: Truncation Parameters page




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 **StorNext 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

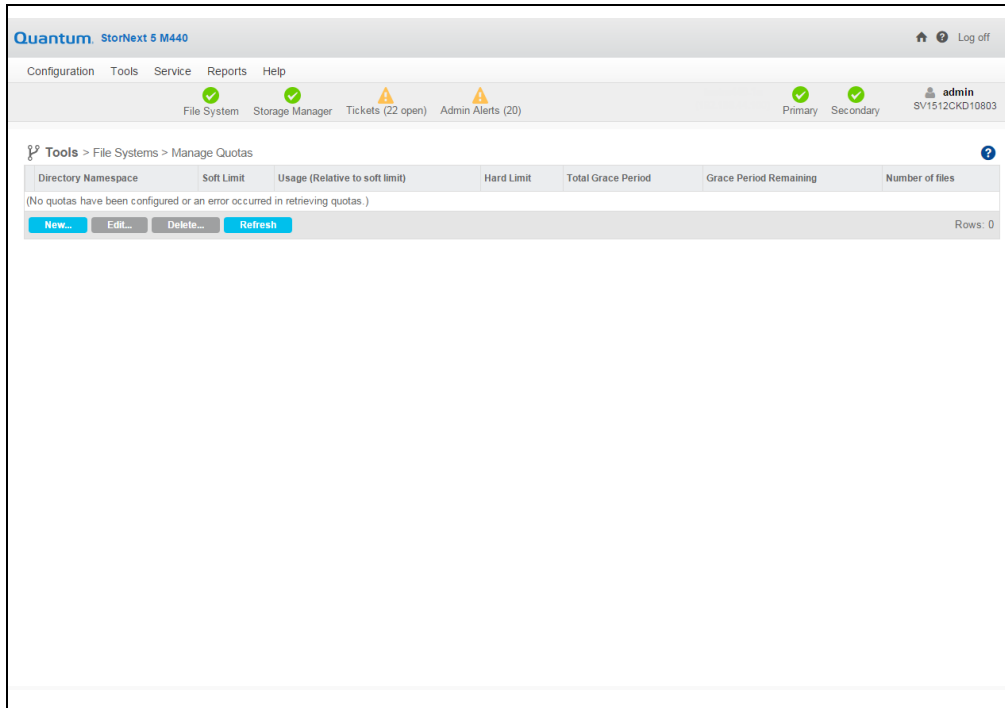
The quota system provides a means for limiting the amount of disk storage consumed on a per user or per group basis across an entire file system or within a designated directory hierarchy. Quota limits apply to the space consumed by disk-block allocations for a user or group, which is not equal to the 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.

---

 **WARNING:** If you enable quotas on an existing file system, the file system becomes unavailable while the changes are applied.

The Manage Quotas page enables you to **view**, **edit**, **delete** or configure **new** quota values as it pertains to StorNext file system. See [Figure 22 on the next page](#) and [Table 2 on the next page](#).

**Figure 22:** Manage Quotas Page



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

**Table 2:** Manage Quotas Table

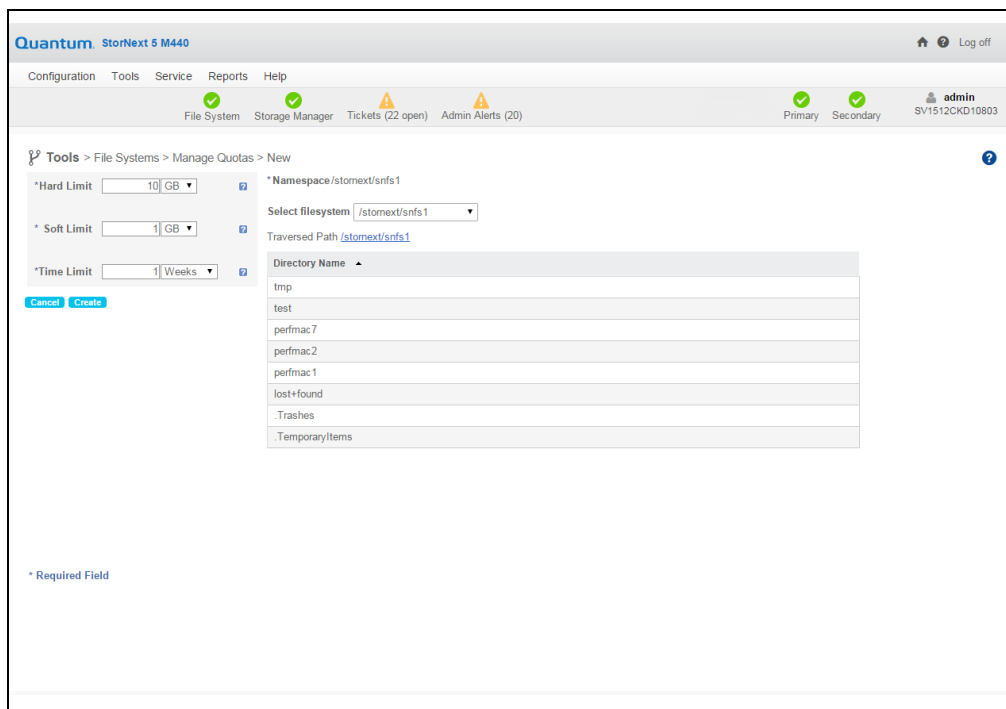
Quota Property	Description
<b>Directory Namespace</b>	Specifies the file system and directory name.
<b>Soft Limit</b>	Specifies the soft limit quota value configured on the system. The <b>Soft Limit</b> is the maximum amount of available usage. You are warned upon reaching the <b>Soft Limit</b> quota value.
<b>Usage (Relative to Soft Limit)</b>	Specifies the percent (%) usage relative to the <b>Soft Limit</b> quota value on the system.
<b>Hard Limit</b>	Specifies the hard limit quota value configured on the system. The <b>Hard Limit</b> quota value is the absolute amount of available usage. You cannot go beyond the <b>Hard Limit</b> quota value.

Quota Property	Description
<b>Total Grace Period</b>	Specifies the total grace period configured on the system. The <b>Total Grace Period</b> is used when you have exceeded the <b>Soft Limit</b> quota value, but are still under the <b>Hard Limit</b> quota value. As soon as the <b>Soft Limit</b> quota value has been exceeded, you have the configured <b>Total Grace Period</b> amount of time to free up space to return your usage under the <b>Soft Limit</b> quota value.
<b>Grace Period Remaining</b>	Specifies the grace period remaining on the system. The <b>Grace Period Remaining</b> is the amount of time remaining to free up space to return your usage under the <b>Soft Limit</b> quota value.
<b>Number of Files</b>	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 23 below](#)).

**Figure 23:** Edit Quotas Page



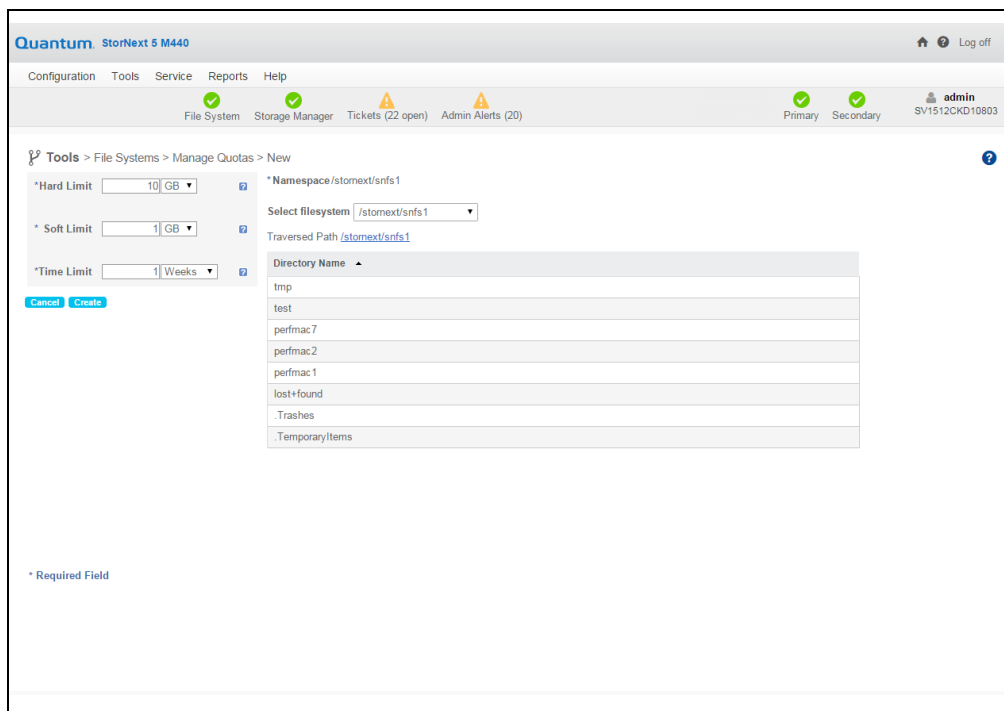
3. Configure the following quota value properties:
  - a. In the **Hard Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down. Refer to [Quota Limits on page 152](#) for additional information on limits.

- 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 the next page](#) 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 24 below](#)).

**Figure 24:** New Quotas Page



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 the next page](#) 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 the next page](#) 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:

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:

- Click **Refresh**.

## Quota Limits

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

Parameter	Description
<b>Hard Limit</b>	The <b>Hard Limit</b> is the absolute limit which file system space usage should not exceed. Any time the total allocated space is at or over the <b>Hard Limit</b> , all further allocations or write requests by the offending user or group will be denied.
<b>Soft Limit</b>	The <b>Soft Limit</b> is a lesser limit. When you exceed this limit (but not the <b>Hard Limit</b> ), allocations are still permitted, but a warning will be written to your console. When the <b>Soft Limit</b> has been overrun for longer than the <b>Total Grace Period</b> , the <b>Soft Limit</b> becomes a <b>Hard Limit</b> and any further allocations or write requests are denied. When the usage again falls below the <b>Soft Limit</b> , 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.

**i Note:** Limits are not enforced against super user accounts.

**i 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:

Type	Description
<b>User, or Group</b>	<b>User and Group Quotas</b> 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.
<b>Directory</b>	<b>Directory Quotas</b> are a little different. The system does not automatically keep track of the usage for each directory. The <b>snquota</b> 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.

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

**i 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.

---

## Rename a Standalone (unmanaged) StorNext File System

Use the following procedure to change the name of a StorNext file system:

**i 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.

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
```

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 5: Storage Manager Tasks

---

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

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

Task	Description
<b>Alternate Retrieval Location and Alternate Store Location</b>	<b>Alternate Retrieval Location</b> allows you to specify a remote retrieval location to use in situations where files stored on tape or a storage disk cannot be accessed. <b>Alternate Store Location</b> provides an automatic system for copying files from a main instance of StorNext to a remote instance of StorNext at the same time as copies are made to tertiary storage at the main site.
<b>Drive Replacement</b>	Allows you to update the drive serial number mappings.
<b>Client-side Encryption</b>	Reports and manages the master keys used for client side encryption.
<b>System Parameters</b>	Allows you to set and modify StorNext system parameters.

StorNext also features the **Active Vault Policy** feature. See [Active Vault Policy on page 204](#).

This chapter contains the following topics:

Storage Components .....	157
Drive Pools .....	159
Media Actions .....	161
Storage Exclusions .....	170
Truncation Exclusions .....	173
Tape Consolidation .....	175
Library Operator Interface .....	177
Software Requests .....	179
Scheduler .....	180
Alternate Store and Retrieval Location .....	183
Distributed Data Mover .....	195
Drive Replacement .....	203
Active Vault Policy .....	204
System Parameters .....	211

# Storage Components

The **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** page is divided into the following sections:

- Libraries
- Storage Disks
- Tape Drives
- Object Storage
- Object Storage Controllers
- Object Storage I/O Paths

## Access the Storage Components Page

On the **Tools** menu, click **Storage Manager**, and then click **Storage Components**.

## Parameters and Descriptions on the Storage Components Page

Section	Parameter	Description
<b>Libraries</b>	<b>State</b>	The current state of the library (for example, <b>Online</b> ).
	<b>Name</b>	The name of the library.
<b>Storage Disks</b>	<b>State</b>	The current state of the drive (for example, <b>Online</b> ).
	<b>Storage Disk</b>	The name of the storage disk.

Section	Parameter	Description
<b>Tape Drives</b>	<b>Serial Number</b>	The tape drive's serial number.
	<b>State</b>	The current state of the drive (for example, <b>Online</b> ).
	<b>Status</b>	The current status of the drive (for example, <b>Free</b> ).
	<b>User Alias</b>	The alias or label applied to a drive.
	<b>Mounted Media</b>	Indicates whether the drive currently has media mounted.
	<b>Dismount Delay</b>	The delay interval, if any, applied when the drive dismounts.
	<b>Compression</b>	Indicates whether compression is enabled ( <b>True</b> ) or disabled ( <b>False</b> ).
	<b>Clean</b>	Cleans the selected tape drive(s) as a job running in the background.
<b>Object Storage Appliances</b>	<b>State</b>	The current state of the Object Storage Appliance destination (for example, <b>Online</b> ).
	<b>Name</b>	The name of the Object Storage Appliance destination.
	<b>Provider</b>	The provider of the Object Storage Appliance destination.
<b>Object Storage Controllers</b>	<b>State</b>	The current state of the Object Storage Controller (for example, <b>Online</b> ).
	<b>Name</b>	The name of the Object Storage Controller.
	<b>Appliance</b>	The appliance of the Object Storage Controller.
<b>Object Storage I/O Paths</b>	<b>State</b>	The current state of the Object Storage I/O Path (for example, <b>Online</b> ).
	<b>Name</b>	The name of the Object Storage I/O Path.
	<b>IP Address</b>	The IP Address/Host name of the Object Storage I/O Path.
	<b>Controller</b>	The name of the parent Object Storage Controller.

## Set Devices Online and Offline

The process for setting devices online or offline is identical regardless of device type.

1. Select the library, storage disk or tape drive you want to place online or offline.
2. **(Optional)** You can also select multiple devices in each category.
3. After you are satisfied with your selections, click **Online** to place selected devices online, or click **Offline** to take selected devices offline.

## Additional Options for Tape Drives

There are four additional options available for tape drives:

Option	Description
<b>Dismount Delay</b>	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 <b>Dismount Delay</b> field, and then click <b>Dismount Delay</b> .
<b>Enable Compression</b>	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 <b>Enable Compression</b> .
<b>Disable Compression</b>	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 <b>Disable Compression</b> .
<b>Clean</b>	This option allows you to request that a drive be cleaned. Before choosing this option, make sure the tape drive is loaded with a cleaning cartridge. When you are ready to proceed, click <b>Clean</b> .

---

## 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.

### Information on the Drive Pools Page

Parameter	Description
<b>Drive Pool Name</b>	The name of the drive pools currently configured.



Parameter	Description
<b>Member Count</b>	The number of drives the pool contains.
<b>Member Drive ID List</b>	The internal StorNext ID of the drives the pool contains.

## View Drive Pool Information

1. On the **Tools** menu, click **Storage Manager**, and then click **Drive Pools**.
2. Select the drive pool whose information you want to see, and then click **View**. 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 (for example, LTO, Lattus, S3COMPAT, Q-Cloud, or a Storage Disk).
  - **Library**: The name of the library to which each drive belongs.
  - **Pool Name**: The name of the drive pool to which each drive belongs.
3. When you are finished viewing drive pool information, click **Done**.

## Add a Drive Pool

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

1. On the **Tools** menu, click **Storage Manager**, and then click **Drive Pools**.
2. Click **New** to add a new drive pool.
3. Enter the following fields:
  - **Drive Pool Name**: The name of the new drive pool you are creating.
  - **Available Drives**: Select one or more available drives for the new drive pool.
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.

## Edit a Drive Pool


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

1. On the **Tools** menu, click **Storage Manager**, and then click **Drive Pools**.
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. Perform one of the following:
  - Click **Add Drives** to add the selected drives from the drive pool.
  - Click **Remove Drives** to remove the selected drives from the drive pool.
5. Click **Done** to confirm your changes.
6. When the confirmation message appears, click **Yes** to proceed or **No** to abort.
7. After a message informs you that the drive pool was successfully modified, click **OK** to continue.

## Delete a Drive Pool

Before you begin, you must first remove all drives in the pool you want to delete. Follow the procedure in [Edit a Drive Pool on the previous page](#) to remove all drives from the pool.

---

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

1. On the **Tools** menu, click **Storage Manager**, and then click **Drive Pools**.
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** page, choose **Media Actions** from the **Tools > Storage Manager** menu.

## View 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 currently resides

- **Media Type:** The type of media
- **Media Format:** The format of the media
- **Formatted:** Displays whether the media is formatted (**true** if the media is formatted, or **false** if the media is not formatted).
- **Media Class:** The media class to which the media belongs
- **Policy Class:** The policy class to which the media belongs
- **Mark Status:** Displays whether the media is **marked** or **unmarked**.
- **Suspect:** Indicates whether the media is considered suspect (possibly unreliable or defective)
- **Write Protected:** Indicates whether the media is write protected
- **File Segment 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

## Filter Media

Most **Media Actions** screens contain a filtering feature that allows you to restrict the available media to those whose media ID match 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 match. Wildcards can be used in the string.
2. Click **Set Filter**.
3. Click **Refresh** to update the list of available media. Only media whose IDs match 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.

## Perform Media Actions

At the top of the page is a dropdown list of actions you can perform for selected media. First choose one of these options from the **Available Actions** list. After the desired action is selected, a new page appears displaying the necessary information to perform the action.

### Mount Media

Select this option to manually mount a piece of storage media into a drive.

1. Select from the **Library** dropdown list the library containing the media you want to mount.
2. Select the media to mount.

3. Select from the drive dropdown list the drive in which the media is to be mounted.
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 manually dismount a piece of previously mounted media from a drive.

1. Select the media you want to dismount.
2. 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. 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 [Scheduler on page 180](#) 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. 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 which do not contain data can be selected for removal. The media is removed from the system and is physically ejected from the library.

1. 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.
5. Use the Library Operator Interface feature to complete the actual physical removal of the media. (See [Scheduler on page 180](#) for more information.)

## 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.

---

 **Caution:** This media action removes all data and information about the data contained on the media prior to it being removed from the library. This information cannot be restored once it is removed.

1. 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.
5. Use the Library Operator Interface feature to complete the actual physical removal of the media. (See [Scheduler on page 180](#) for more information.)

## Reclassify Media

Select this option to change the media classification.

1. Select from the **Media Class** dropdown list the current media class designation you want to change. After you select the desired class, all media that currently have this classification appear in the Media section.

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. Select from the **Destination Media Class** dropdown list the new media class designation for the selected media. Select one of these options:
  - **DATA**: This media class means that media are candidates for read/write operations. Most media residing in the library have this classification unless they are full.
  - **ADDBLANK**: This is the default class with which media are associated when they are added to StorNext MSM. Running the Fsmedin command pulls media from this class and changes the classification to **DATA**.
  - **IMPORT**: Before running the fsmedin command on TSM-exported media, the classification should be changed to **IMPORT**.
  - **CHECKIN**: This classification is used for re-entering media which have been checked out. Media must be reclassified with **CHECKIN** prior to TSM performing fsmedin with the checkin option.
  - **MIGRATE**: TSM reclassifies media to this classification when the media becomes full according to the FS\_PERCENT\_FULL system parameter. Media with this classification can still be read.
  - **CLEAN**: Media in the class are cleaning media. If the barcode of a media ends with CLN, MSM imports the media into this class instead of **ADDBLANK**.
  - **REMOVE**: Media get reclassified to **REMOVE** when fsmedout is used.
  - **BACKUP**: Media with this classification were used for backups before backups were managed by StorNext storage polices. 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. This media action operates only on media that is currently blank, whether it is assigned to a policy class or is scratch. All media that contain data are filtered from this action.

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. Select from the **Destination Policy Class** dropdown list the policy class to which you want to assign the 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 to another media (such as scratch), 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.

The following table displays the support matrix for transcribe operations:

		Destination				
		Object Storage	Tape	Storage Disk	Q-Cloud Archive	Q-Cloud Vault
Source	Object Storage	Yes	No	No	Yes	Yes
	Tape	Yes	Yes	No	Yes	Yes
	Storage Disk	Yes	No	No	Yes	Yes
	Q-Cloud Archive	Yes	No	No	Yes	Yes
	Q-Cloud Vault	No	No	No	No	No

Based on the selected source media type(s), the StorNext GUI will display the appropriate transcribe options.

## Special Considerations for the Transcribe Media Action

The **Source Policy Class Filter** is grayed out when transcribing tapes as tapes can only contain a single policy class.

The **Destination Media Type** is specified by selecting one of the following:

- **Any Available Media** for the selected **Media Type**.
- **Any Available Blank Media** for the selected **Media Type**.
- **Specific Media ID** of the selected **Media Type**.
- Manually enter the desired **Media ID**.
  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. **(Optional)** In the **Transcribe Media Parameters** section, you can copy only data file(s) belonging to the specified class in the **Source Policy Class Filter** list. Alternatively, in the **Source Policy Class Filter** list, select **All Policy Classes** to copy all the files to the destination media. Additional parameters include the **Destination Media Type** list, the **Destination Media Format** list, the **Destination Media**

list, the **Encryption Type**, the **Encryption Master Key**, and the **Compression Type**.

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

## 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 click **All** to select all media.
3. Select from the **New Media State** dropdown the desired attribute to apply to the selected media.
4. Click **Apply**.
5. When the confirmation message appears, click **Yes** to move the selected media, or **No** to abort.
6. Repeat steps 3 - 5 to apply additional attributes to the selected media.

## Clean Media by Media ID

Select this option if you want to select media for cleaning based on media ID. When you run this function, the StorNext Storage Manager removes inactive files from the media that have not been accessed since the specified endtime. This process does not affect current file versions on the media. Periodic cleaning helps prevent inactive information from growing to an unmanageable size.

---

 **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 **End Time** field, enter the date and time you want the cleaning process to stop. The default is the



current date and time.

3. Click **Apply**.
4. 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. When you run this function, the StorNext Storage Manager removes from the file system inactive files that have not been accessed since the specified endtime. This process does not affect current file versions on the media. Periodic cleaning helps prevent inactive information from growing to an unmanageable size.

---

 **Caution:** Inactive file versions cleaned from the media cannot be recovered or used again.

1. Select from the **Managed and Mounted File Systems** dropdown list the file system to be cleaned.
2. At the **End Time** field, enter the date and time you want the cleaning process to stop. The default is the current date and 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 on the selected file system are cleaned. When you run this function, the StorNext Storage Manager removes inactive files associated with the selected policy class that have not been accessed since the specified endtime. This process does not affect current file versions on the media. Periodic cleaning helps prevent inactive information from growing to an unmanageable size.

---

 **Caution:** Inactive file versions cleaned from the media cannot be recovered or used again.

1. Select from the **Policy Classes** dropdown list the policy class whose media you want to clean.
2. At the **End Time** field, enter the date and time when you want the cleaning processing to stop. The default is the current date and 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. Depending on how your library works, you may need to manually load media into the library mailbox prior to running this action. Or, after running this action you may be prompted to load media into the mailbox. This action is library dependent.

1. In the **Library** list, select the library into which you want to add a media mailbox.
2. In the **Add Media Mailbox Parameters** section, select the **Port** from the list.
3. Click **Apply**.
4. When the confirmation message appears, click **Yes** to proceed, or **No** to abort.
5. 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. If an error message appears and indicates "Failed to import Media from Mailbox", the new media needs to be inserted into the mailbox prior to running this action. Click **OK**, add new media to the mailbox, and rerun this action.

If a message informs you that the operation was successful, click **OK**. After you see this message you are ready to load media into the library mailbox.

## Add Media Bulk Load

Select this option to add media to a library via bulk loading. Before running this action, you must first add the new media into the desired library through a manual process. The media must be inserted through internal slots, not into the mailbox. (That is, you must open the library door and put new media in the internal slots, close the library, allow it to do an audit, and then run this action.)

1. Select from the **Library** dropdown list 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.

## Set Object Storage Media Availability

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

1. In the **Available Media Actions** list, click **Set Object Storage Media Availability**. The GUI page displays a list of **namespaces** that exist under different Object Storage destinations.
2. In the **Media** table, click a **namespace**.
3. In the **New Object Storage Media State** list, click **Available** or **Unavailable**.
4. To accept your selection, click **Apply**.

### How to View Object Storage Media Availability

To view the availability of the Object Storage media, perform the following procedure:

1. In the **Reports** menu, click **Media** to display the details for all media (including cleaning media) in a selected library or all libraries.

2. To view a report for a particular piece of media, select the desired media from the list:
  - a. To select multiple media, hold down the CTRL key while you click additional media.
  - b. To select all media, click the checkbox to the left of the **Name** heading.
3. After you have selected media, click **View Media Information Report**. The **Media Information Report** page appears. This report allows you to see all files on the selected media. The availability attribute (**Available** or **Unavailable**) is displayed next to **Media Status**.
4. **(Optional)** Save the report output as a CSV file (Microsoft Excel format) by clicking **Download**.
5. When you are finished viewing the information report, click **Done**.

## Reassign Orphaned Media

Select this option to re-assign orphaned media (for example, media currently in transit, with no current or pending archive ID set) to the intended archive.

1. Select from the **Library** dropdown list the library containing the orphaned media you want to reassign.
2. Select one or more media, or check the box to the left of the **Media ID** heading to select all media.
3. At the **Reassign OrphanedMedia Parameters > Destination Library** field, select the destination library to which you want to reassign the selected media.
4. At the **Reassign OrphanedMedia Parameters > Current Media Location** field, select the location where the selected media is currently located.
5. Click **Apply**.
6. Complete the process by reassigning the media you specified to the destination library.

---

# Storage Exclusions

The **Storage Exclusions** option on the **Tools** menu 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.

---

**i 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.

## Access Storage Exclusions

On the **Tools** menu, click **Storage Manager**, and then click **Storage Exclusions**. The **Tools > Storage Manager > Storage Exclusions** page appears. Any previously saved exclusions are displayed.

## Add an Exclusion Pattern

1. On the **Tools** menu, click **Storage Manager**, and then click **Truncation Exclusions**. The **Tools > Storage Manager > Truncation Exclusions** page appears.
2. Click **Add New Exclusion**.
3. In the **Type** list, select one of the following types of exclusions:

Exclusion Type	Description
<b>Unknown</b>	This type appears until you select a different type, and remains if you select no type.
<b>Match</b>	<p>File paths which match the string in the <b>Pattern</b> 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.</p> <p>For example, enter <code>/sn/foodir/*</code> to exclude all files in the directory <code>/sn/foodir</code>.</p> <p><b>Note:</b> The difference between <b>Match</b> and <b>Match Path</b> is that slashes must be explicitly entered for the <b>Match Path</b> type. When using <b>Match Path</b> you cannot substitute a wildcard for a slash.</p>
<b>Match Path</b>	<p>File paths which match the string in the <b>Pattern</b> field are excluded. Wildcards can be used for this type.</p> <p><b>Note:</b> The difference between <b>Match</b> and <b>Match Path</b> is that slashes must be explicitly entered for the <b>Match Path</b> type. When using <b>Match Path</b> you cannot substitute a wildcard for a slash.</p>
<b>Begins With</b>	File paths beginning with the string in the <b>Pattern</b> field are excluded.
<b>Ends With</b>	File paths ending with the string in the <b>Pattern</b> field are excluded.
<b>Contains</b>	File paths containing the string in the <b>Pattern</b> field are excluded.
<b>Exact</b>	Only file paths that exactly match the string in the <b>Pattern</b> field are excluded.

4. At the **Pattern** field, enter the search string for the directory. Depending on the exclusion type you selected, this could be a whole or partial path name for a directory, or a string. If you selected the **Match** or **Match Path** type, you can use the following wildcards at the **Pattern** field:

Wildcard	Description
? (question mark)	Substitute any single character. For example, if you enter <b>t?p</b> , it will match “top” “tip” and “tap”.
* (asterisk)	Substitute one or more characters. For example, if you enter <b>f*I</b> , it will match “ful” “fail” “foil” “fall” and “frail”.
[ ] (brackets)	When you use this wildcard, a set of characters are specified. For example, if you enter <b>[abc]*</b> , any string beginning with the letters “a” “b” or “c” are matched. When using brackets you can also specify a range by using the - (dash) character. For example, if you enter <b>file[1-4]</b> , 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 <b>[!abc]*</b> , any string that does not begin with the letters “a”, “b” or “c” are matched.

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

## Delete an Exclusion

To delete an exclusion, click **Delete** to the right of the exclusion you want to delete.

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

## Examples of Exclusions

The following are some exclusion examples, including the **Type** selected and the string entered in the **Pattern** field to create the exclusion:

Exclusion	Type	Pattern
Exclude files in directories beginning with the letter <b>x</b> or <b>y</b>	<b>Match</b>	<b>*/[xy]*/*</b>
Exclude files in a directory named <b>/private/confidential</b>	<b>Match Path</b>	<b>/stornext/snfs1/lfs_ only/private/confidential/*</b>

**i Note:** For **Match Path**, the full directory path must be specified.

# Truncation Exclusions

The **Truncation Exclusions** option on the **Tools** menu 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.

**i 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 execute the command `fschfiat -t c` to remove the exclusion from the file.

## Access Truncation Exclusions

On the **Tools** menu, click **Storage Manager**, and then click **Truncation Exclusions**. The **Tools > Storage Manager > Truncation Exclusions** page appears. Any previously saved exclusions are displayed.

## Add an Exclusion Pattern

1. On the **Tools** menu, click **Storage Manager**, and then click **Truncation Exclusions**. The **Tools > Storage Manager > Truncation Exclusions** page appears.
2. Click **Add New Exclusion**.
3. In the **Type** list, select one of the following types of exclusions:

Exclusion Type	Description
<b>Unknown</b>	This type appears until you select a different type, and remains if you select no type.

Exclusion Type	Description
<b>Match</b>	<p>File paths which match the string in the <b>Pattern</b> 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.</p> <p>For example, enter <code>/sn/foodir/*</code> to exclude all files in the directory <code>/sn/foodir</code>.</p> <p><b>i Note:</b> The difference between <b>Match</b> and <b>Match Path</b> is that slashes must be explicitly entered for the <b>Match Path</b> type. When using <b>Match Path</b> you cannot substitute a wildcard for a slash.</p>
<b>Match Path</b>	<p>File paths which match the string in the <b>Pattern</b> field are excluded. Wildcards can be used for this type.</p> <p><b>i Note:</b> The difference between <b>Match</b> and <b>Match Path</b> is that slashes must be explicitly entered for the <b>Match Path</b> type. When using <b>Match Path</b> you cannot substitute a wildcard for a slash.</p>
<b>Begins With</b>	File paths beginning with the string in the <b>Pattern</b> field are excluded.
<b>Ends With</b>	File paths ending with the string in the <b>Pattern</b> field are excluded.
<b>Contains</b>	File paths containing the string in the <b>Pattern</b> field are excluded.
<b>Exact</b>	Only file paths that exactly match the string in the <b>Pattern</b> field are excluded.

- At the **Pattern** field, enter the search string for the directory. Depending on the exclusion type you selected, this could be a whole or partial path name for a directory, or a string. If you selected the **Match** or **Match Path** type, you can use the following wildcards at the **Pattern** field:

Wildcard	Description
? (question mark)	<p>Substitute any single character.</p> <p>For example, if you enter <code>t?p</code>, it will match “top” “tip” and “tap”.</p>
* (asterisk)	<p>Substitute one or more characters.</p> <p>For example, if you enter <code>f*I</code>, it will match “ful” “fail” “foil” “fall” and “frail”.</p>

Wildcard	Description
[ ] (brackets)	<p>When you use this wildcard, a set of characters are specified.</p> <p>For example, if you enter <b>[abc]*</b>, any string beginning with the letters “a” “b” or “c” are matched. When using brackets you can also specify a range by using the - (dash) character.</p> <p>For example, if you enter <b>file[1-4]</b>, the strings “file1” “file2” “file3” and “file4” are matched. You can also specify a complement of characters to not match.</p> <p>For example, if you enter <b>[!abc]*</b>, any string that does not begin with the letters “a”, “b” or “c” are matched.</p>

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

## Delete an Exclusion

To delete an exclusion, click **Delete** to the right of the exclusion you want to delete.

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

## Examples of Exclusions

The following are some exclusion examples, including the **Type** selected and the string entered in the **Pattern** field to create the exclusion:

Exclusion	Type	Pattern
Exclude files in directories beginning with the letter <b>x</b> or <b>y</b>	<b>Match</b>	<b>*/[xy]*/*</b>
Exclude files in a directory named <b>/private/confidential</b>	<b>Match Path</b>	<b>/stornext/snfs1/lfs_ only/private/confidential/*</b>

**i Note:** For **Match Path**, the full directory path must be specified.

# 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.




Previous releases of StorNext allowed you to consolidate tape space by manually executing the command `fsdefrag` (defragment file system), but you can schedule 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.

## Set Tape Consolidation Parameters

1. On the **Tools** menu, click **Storage Manager**, and then click **Tape Consolidation**. The **Tools > Storage Manager > Tape Consolidation** page appears.
2. Enter the following fields:

Parameter	Description
<b>Tape Full Threshold</b>	Specify the percentage at which tapes become candidates for consolidation. For example, enter <b>85</b> if you want tapes flagged for consolidation when they are 85% full.
<b>Fragmentation Threshold</b>	Specify the percentage of fragmentation at which tapes become candidates for consolidation. For example, enter <b>15</b> if you want tapes flagged for consolidation when 15% of the tape becomes fragmented.   <b>Note:</b> 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%.
<b>Clean Versions Expiration</b>	Specify the number of <b>Days, Weeks, Months</b> or <b>Years</b> after which you want to clean up versions for deleted or modified files.
<b>Max Active Tape Copy Operations</b>	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.
<b>Max Medcopy Attempts</b>	Specify the maximum number of attempts before StorNext stops trying to copy media experiencing copy failures.
<b>Max Media to Consolidate</b>	Specify the maximum number of tape media which can be consolidated during one <code>fsdefrag</code> process.

Parameter	Description
<b>Ignore Vaulted Media</b>	Enter <b>True</b> to ignore tape media in the vault, or <b>False</b> to include media in the vault.
	<b>i Note:</b> To be candidates for defragmentation, media must pass both the <b>Tape Full Threshold</b> and <b>Fragmentation Threshold</b> 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** page.
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.

## Schedule Tape Cleaning and Defragmentation

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

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

**i 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 180](#).

---

# Library Operator Interface

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

## Information on the Library Operator Interface Page

Parameter	Description
<b>Media ID</b>	The unique identifier for each piece of media.

Parameter	Description
<b>Current Library</b>	The name of the library where media currently reside. For <b>Move</b> or <b>Enter</b> actions, the source archive is displayed. For example, <b>Ejected From: vault123</b> .
<b>Pending Library</b>	The name of the destination library to which the media action will be carried out.
<b>Reason</b>	The reason motivating the action upon the media.
<b>Media Location</b>	The current physical location of the media.
<b>Action Required</b>	The action to be performed on selected media.
<b>Details</b>	Displays information or errors from the back-end system. If information is not available, the field is blank.

## Perform an Action on Media

1. On the **Tools** menu, click **Storage Manager**, and then click **Library Operator Interface**.
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. Due to the limitation of the port selection, the media chosen must all be related to the same action and library.
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.

**i 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**.

## Active Vault

Traditionally, vaulting enables a library to manage more media than it can physically hold, and is mainly used for data archival purposes.

With the recent introduction of Quantum Scalar i6000 Tape Library's Active Vault partition feature that can keep vaulted media inside the library, there is added value in using an Active Vault policy that works with the library's Active Vault partition feature.

With **Active Vault**, you can manage StorNext Management Suite (SNMS) licensed capacity more proactively. Active Vault offers a flexible, and fully configurable vault policy that helps you manage SNMS licensed capacity.

**Active Vault** is integrated with the SNMS scheduling automation feature, and i6000 Active Vault plug-in.

Each time a scheduled Active Vault policy runs, if SNMS used capacity exceeds a high water mark percentage of licensed capacity, qualified media will be vaulted (using the command `vsmove`) until used capacity is below a low water mark percentage of licensed capacity or qualified media is exhausted.

For additional details and configuration information, see [Active Vault Policy on page 204](#).

---

## Software Requests

The **Software Requests** functionality displays all outstanding MSM (Media Manager) subsystem requests, such as low level mounts, dismounts, media moves, imports, exports, checkins, checkouts, etc. For example, commands that will be sent to a tape library; not the higher level stores, retrieves, etc. You can cancel such requests using the Software Requests functionality.

For additional information, see the command `vscancelreq` in the *StorNext MAN Pages Reference Guide*.

---

**Note:** The Software Requests functionality is not for client downloads.

### View a Software Request

- On the **Tools** menu, click **Storage Manager**, and then click **Software Requests**. The **Software Requests** page appears and provides the following information.

Parameter	Description
<b>ID</b>	The software request's identifier.
<b>Type</b>	The type of software request currently running.
<b>State</b>	The current state of the request ( <b>Current</b> or <b>Pending</b> ).
<b>Time</b>	The time the software request was initiated.
<b>Priority</b>	The priority assigned to the software request. <b>Note:</b> <b>1</b> is the highest priority, and <b>32</b> is the lowest.

## Cancel a Software Request

1. On the **Tools** menu, click **Storage Manager**, and then click **Software Requests**. The **Software Requests** page appears.
2. If desired, click **Refresh** to update the displayed list of software requests.
3. Select the request you want to cancel.
4. Click **Cancel Request**.
5. When the confirmation message appears, click **Yes** to proceed or **No** to abort.
6. After a message informs you that the request was successfully canceled, click **OK**.

---

## Scheduler

StorNext events are tasks that are scheduled to run automatically based on a specified schedule.

**i 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.

## Events You Can Schedule

Parameter	Description
<b>Clean Versions</b>	This scheduled event cleans old inactive versions of files.
<b>Clean Info</b>	This scheduled background operation removes from StorNext all information about active and inactive versions of files on media that have had fsminfo processing run on them. Additionally, any files on the media will be removed from disk if they still exist on disk, have no disk data, and there are no other copies on other media.
<b>Rebuild Policy</b>	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.
<b>Partial Backup</b>	By default, a partial backup is run on all days of the week the full backup is not run. Partial backups include database journals, configuration files, and file system journal files.
<b>Full Backup</b>	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.
<b>Health Check</b>	By default, health checks are set up to run every day of the week, starting at 7:00 a.m.

Parameter	Description
<b>Tape Defragmentation</b>	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 page are included in the defragmentation process. For more information, see <a href="#">Tape Consolidation on page 175</a> .
<b>Active Vault</b>	Active Vault 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.
<b>Archive Compare</b>	The function of utility is to verify that Media Manager tape drive configuration matches the Tertiary Manager view and also verify those tape drives still exist within the actual archive. You do not select a specific library, as the Drives Validation Report and Drives Validation Update will be run on all configured libraries, excluding vaults.

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 [Add a Schedule below](#) or [Edit a Schedule on the next page](#).

## View a Schedule

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

1. On the **Tools** menu, click **Storage Manager**, and then click **Scheduler**.
2. Select the event you want to view, and then click **View**.
3. When you are finished viewing event details, click **Done**.

## Add a Schedule

1. On the **Tools** menu, click **Storage Manager**, and then click **Scheduler**.
2. Click **New**.
3. At the **Name** field, enter the name you want to assign to the new schedule.
4. Select one of the following event types from the **Feature** list:
  - **Clean Versions**
  - **Clean Info**
  - **Rebuild Policy**
  - **Partial Backup**
  - **Full Backup**
  - **Health Check**
  - **Tape Defragmentation**

- **Active Vault** (refer to section [Active Vault below](#) below for a brief description)
  - **Archive Compare**
5. At the **Period** field, select the execution interval for the new schedule: **Daily**, **Weekly** or **Monthly**. You can also select individual 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.

## Edit a Schedule

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

1. On the **Tools** menu, click **Storage Manager**, and then click **Scheduler**.
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.

## Delete a 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 can delete a default schedule, as long as you have created at least one other schedule of that type. For example, the schedule count has to be at least one for a given type.

1. On the **Tools** menu, click **Storage Manager**, and then click **Scheduler**.
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.

## Active Vault

Traditionally, vaulting enables a library to manage more media than it can physically hold, and is mainly used for data archival purposes.

With the recent introduction of Quantum Scalar i6000 Tape Library's Active Vault partition feature that can keep vaulted media inside the library, there is added value in using an Active Vault policy that works with the library's Active Vault partition feature.

With **Active Vault**, you can manage StorNext Management Suite (SNMS) licensed capacity more proactively. Active Vault offers a flexible, and fully configurable vault policy that helps you manage SNMS licensed capacity.

**Active Vault** is integrated with the SNMS scheduling automation feature, and i6000 Active Vault plug-in.

Each time a scheduled Active Vault policy runs, if SNMS used capacity exceeds a high water mark percentage of licensed capacity, qualified media will be vaulted (using the command `vsmove`) until used capacity is below a low water mark percentage of licensed capacity or qualified media is exhausted.

For additional details and configuration information, see [Active Vault Policy on page 204](#).

---

## Alternate Store and Retrieval Location

Task	Description
<b>Alternate Retrieval Location</b>	Allows you to specify a remote retrieval location to use in situations where files stored on tape or a storage disk cannot be accessed.
<b>Alternate Store Location</b>	Provides an automatic system for copying files from a main instance of StorNext to a remote instance of StorNext at the same time as copies are made to tertiary storage at the main site.

### Overview of 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, the Alternate Retrieval Location 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 permissions are changed based on the file on the alternate machine.

**Note:** 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.

## Overview of Alternate Store Location

The StorNext Alternate Store Location feature provides an automatic system for copying files from a main instance of StorNext to a remote instance of StorNext at the same time as copies are made to tertiary storage at the main site. The remote copies can serve as a copy-of-last resort with the Alternate Retrieval Location feature in StorNext. The feature also supports background copying of files that existed before the deployment of the feature. Background copy activity can be limited to avoid overwhelming the use of the StorNext system for new files.

---

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

The feature can be enabled per Storage Manager Class Policy per file system. After it is enabled, new files are automatically copied to the remote site. Pre-existing files can also be enabled for automatic copying by providing them as input to the `altstoreadd` command. On completion of the copy, the completion is recorded per file so that the completion status can be displayed with the `fsfileinfo` command. Prior to completion of the copy action, the `fsfileinfo` command reports if a file is enabled for copying, and reports that the copy action has not been performed.

After the copy action has been performed, the main-site instance of StorNext does not maintain a status of the remote-site copies. Ensuring that the remote-site copies are maintained as a read-only image of the main site is an administrative responsibility. Necessary actions to remedy any discrepancies that can accumulate at the remote site as a result of deletions or modifications caused by either user activities or equipment failure are also administrative responsibilities.

## Configure an Alternate Retrieval Location

1. On the **Tools** menu, click **Storage Manager**, and then click **Alternate Retrieval Location**.
2. At the field **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).
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**.

## Configure an Alternate Store Location

1. On the **Tools** menu, click **Storage Manager**, and then click **Alternate Store Location**.
2. At the field **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).
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**.

## Architecture

The StorNext Alternate Store Location feature runs an agent on both the main site and the remote site, which must be at the same StorNext release level. Each site can serve as an Alternate Store Location main site, remote site, or both at the same time. One StorNext instance can serve as the remote site for multiple main-site StorNext instances. Each main-site file system must have a distinct remote-site file system that cannot be shared with any other main site. A main site can only send copies to a single remote site.

The `fs_altstore` resident process of TSM provides both the main site and the remote site processing. It starts and stops with TSM, runs in an idle mode when it is not configured as a main site, and is always ready to serve as a remote site. Communication between the resident processes on the main and remote sites is with a TCP/IP socket connection to a designated port. The `altstore` service can be controlled with the `fsschedlock` utility.

Transferring of files between the main site and the remote site is done with a user-modifiable script. The default script provided with StorNext uses the Secure Shell `scp` command for maximum security. Other transfer schemes, such as, FTP or NFS can be substituted according to the security and performance requirements at customer sites. Files are transferred to a staging location at the mount point for each remote file system that stores remote copies. The name of the staging folder is `.AltStoreStagedir` under the mount point of the file system. Files are moved or copied from there to each file's target pathname. The per-file owner, group and `rxw` permissions are set to match their values at the main site. Files are read on the main site with root permission. They are sent to the staging directory with the `altstore` user's permissions. The remote agent does the final actions with root permission.

When a file is created under a policy with Alternate Store Location enabled, the file is added to a database table as a demand type request for remote copying. Files that were created before the feature was configured can be added to the table as a background type request for remote copying. All of the demand-type requests have priority and are processed before any background-type request is processed. There is a configurable limit on the number of simultaneous demand-copy transfer processes that can be started. There is a separate configurable limit on the subset of those transfer processes that can be used for background copies when they are not being used for demand copies. Experimenting with these limits can help to tune your system so that throughput is maximized while keeping system responsiveness at an appropriate level for other activities.

Truncated files that are enabled for copying and are in the database table of requests are automatically retrieved. The truncation status is checked at the point of calling the transfer script for the file. These files are changed to the background request type and placed on an internal queue for retrieval. This queue is allowed to grow to a maximum size for a limited amount of time to allow the retrieve operation to optimize the order of retrievals within a single request. The queue-size parameter is described in [ALTSTORE\\_RTRV\\_WAIT\\_COUNT on page 187](#), and the parameter for time is described in [ALTSTORE\\_RTRV\\_WAIT\\_TIME on page 187](#). When a large number of new demand-type requests get created, these can cause the retrieve-queue to be emptied to give priority to the demand requests. Eventually, the truncated background-type

requests will be reloaded and retrieved. When they have been retrieved, they are converted to demand-type requests for expedient copying to avoid being re-truncated before that can be completed.

When a file is enabled for copying, but has not been copied yet, manual requests for truncation with the **fsrmdiskcopy** command will state that “not all copies...are stored”, and will reject the truncation request. The retention policy for immediate file cleanup will not apply when the remote copy has not been made. However, truncation-policy processing that depends on file-system fill levels can truncate files that do not have a remote copy.

The intent of the Alternate Store Location feature is to create copies of file systems on a remote StorNext system with matching pathnames, user and group owners, and simple permissions (rwx). However, the feature is enabled in policy classes, which are applied to folders within file systems, and the mount points for the top-level directories of file systems can have different pathnames on the main and remote StorNext systems. Displayed file and directory owner and group names will be identical if the mappings of user and group names to UID and GID values are the same on both systems.

The indication in **fsfileinfo** output that a file copy exists is simply a record that a file copy completed successfully. Once a copy has been made, the Alternate Store Location feature does not perform any further tracking of the remote copy at the main site.

It is possible for a file to be marked for copying, but not to have a row in the database table to cause the copy to be performed. In that event, periodic rebuild processing will discover the discrepancy and will add a row to the table for the file to be copied.

## Components of the Alternate Store Location Feature

### fs\_altstore Resident Process

The **fs\_altstore** resident process is running whenever TSM is running. It is visible in process listings. There are no command-line parameters, and the process must not be run by hand. Its configurable options are controlled by the following editable settings as described in the `/usr/adic/TSM/config/fs_sysparm.README` file.

### Sysparm Parameters

#### ALTSTORE\_AGENT\_PORT\_NUMBER

Communication between the **altstore** processes running on the main and remote sites is through a TCP socket on the designated port number 12333. This port number must be allowed to be passed through any firewalls between the two sites. This value must be identical on both main and remote sites.

**Default: 12333**

#### ALTSTORE\_MAX\_CONNECTIONS

This specifies the maximum number of connections between a remote-site **Altstore** daemon and all of the main-site **Altstore** daemons it serves. For **Altstore** daemons acting as a remote-site, this value should be equal to or greater than the sum of the values (**ALTSTORE\_NUM\_TXFR\_STREAM + 1**) for each main-site. For **Altstore** daemons acting as a main-site only, the default value is recommended.

**Default:** (`ALTSTORE_NUM_TXFR_STREAM + 1`) or **31**, whichever is larger.

#### `ALTSTORE_NUM_BG_TXFR_STREAM`

This specifies the number of simultaneous transfer streams for background transfers in the Alternate Store Location feature. This is the maximum number of transfer streams that can be used for transferring background files when the streams are not being used for on-demand transfers. This limit ensures that on-demand transfers can be responsive when new transfers are needed. This number must be less than or equal to `ALTSTORE_NUM_TXFR_STREAM`.

**Default:** 2. The valid range is 1 through 30.

#### `ALTSTORE_NUM_SCRATCHPOOL`

This is the number of internal working-queue elements. It should be about 40 times the number of simultaneous transfer streams specified by `ALTSTORE_NUM_TXFR_STREAM`.

**Default:** 200. The valid range is 10 through 3000.

#### `ALTSTORE_NUM_TXFR_STREAM`

This specifies the number of simultaneous transfer streams for on-demand transfers in the Alternate Store Location feature. The optimal number depends on the number of CPU cores in the server computer, the capacity of the network hardware, the ability of the destination computer to handle traffic, and the contention with other services using the network. Experimentation can help to characterize system throughput and determine when this number reaches diminishing returns of total throughput.

**Default:** 5. The valid range is 1 through 30.

#### `ALTSTORE_POLL_TIME`

Time to wait (in seconds) before the `altstore` daemon wakes up to check for alternate store location candidates. The default value is 60 seconds. The valid range is 60 to 300 seconds. A longer value reduces the system resources consumed. A shorter value reduces the delay for the system to recognize the start or end of a lockout period as described in section: `fsschedlock`.

**Default:** 60

#### `ALTSTORE_RTRV_WAIT_COUNT`

When retrieving truncated files, the Alternate Store Location feature will wait for this number of retrievable files to accumulate if the `ALTSTORE_RTRV_WAIT_TIME` limit does not expire first.

**Default:** 10. The valid range is 1 through 30.

#### `ALTSTORE_RTRV_WAIT_TIME`

When retrieving truncated files, the Alternate Store Location feature will wait up to this many seconds for a batch of retrievable files to accumulate if the `ALTSTORE_RTRV_WAIT_COUNT` is not reached first.

**Default:** 30. The valid range is 1 through 900.

#### `ALTSTORE_TXFR_SCRIPT`

Pathname of the customizable script for transferring files with the Alternate Store Location feature.

**Default:** `/usr/adic/TSM/bin/fs_altstore_transfer`

### ALTSTORE\_TXFR\_USERNAME

The `altstore` user ID owns the staging directory on the remote site and is used as the target ID for file transfers. The value must be identical on both main and remote sites.


**Default:** `altstore`

### ALTSTORE\_VERSIONING

This value is configured on the source MDC and controls the behavior that occurs on the Alternate Store Location target MDC for existing files when they are being updated. The 'y' option causes the file to be overwritten, which creates another Storage Manager version of the file on the target MDC. This option incurs higher CPU overhead. The 'n' option causes the file to be re-created, which creates the first Storage Manager version of a new file on the target MDC. This option incurs lower CPU overhead.

**Default:** `y`

---

 **WARNING:** Use of the `fsversion` command to retrieve a different version of a file on either the main or remote site can result in a discrepancy between the sites. Use the `fsversion` feature with caution. When it is necessary to restore a main-site file to an older version, the main- and remote-site files can be kept in sync by rewriting a portion of the file on the main site, which will cause it to be stored as a new version on both the main and remote sites.

### ALTSTORE\_VERSIONING\_TIMEOUT\_FACTOR

This value is configured on the main-site MDC and controls the amount of time that is allowed for file overwrites to occur on the remote-site MDC before being considered an error. This value is only applicable when `ALTSTORE_VERSIONING` is set to 'y'. The versioning timeout factor is specified in seconds per gigabyte of data. This value should be increased if file overwrites are timing out.

**Default:** `60`. The valid range is 30 through 600.

## User Customizable Transfer Script

Quantum provides a default transfer script that uses SCP for the highest level of security. The choice of file-transfer technologies is dependent on local considerations for performance and security.

When the Alternate Store Location feature is configured, the same script is used for the Alternate Retrieval Location feature. Its two modes (store and retrieve) are distinguished by the value of the fourth parameter, which is set to `ALTSTORE` for the Alternate Store Location feature.

When the transfer script is in `ALTSTORE` mode, it uses the `root` user ID at the main site for sending files since it must be able to read all users' files, and uses the `altstore` user ID to write files into the remote site's staging directory. When the transfer script is also used for the Alternate Retrieval Location feature, it must use `root` privileges on the remote site as well to be able to read any file.

Instructions are provided inside the script for creating the `altstore` user and setting up Secure Shell (SSH) credentials for running the script. To test the credentials, while running as `root` on the main site, you should be able to use the `scp` command to copy a file to the `altstore` user's home directory on the remote site.

When both the Alternate Store Location feature and the Alternate Retrieval Location feature are configured, do not create the `altstore` user on the remote site. Instead, set up SSH credentials for the `root` user to access the remote site as `root`, and change the value of `ALTSTORE_TXFR_USERNAME` to `root` in the `/usr/adic/TSM/config/fs_sysparm_override` file.

If you prefer to use a different script, create a new executable file under a new name so your custom script is not overwritten when StorNext is upgraded. Set the `ALTSTORE_TXFR_SCRIPT` parameter in `/usr/adic/TSM/config/fs_sysparm_override` to the pathname of the executable. The Alternate Store Location parameters for the executable are as follows:

1. Main-site full pathname of the file to transfer.
2. Remote site IP address or Host Name (as configured with `fsaltnode`).
3. Remote-site full pathname of the file in the staging area.
4. The word "ALTSTORE", which indicates that the executable is being invoked for stores. When the parameter is missing, the executable is being invoked for retrieves.

## Commands

Following are the commands for operating and administering the Alternate Store Location feature. These are described in greater detail in their man pages:

### altstoreadd

Enables the Alternate Store Location remote-copy feature on files and adds them to the alternate store candidate list. It can also be used to verify that a remote copy has been made for a file or list of files.

### altstoremod

Display or manipulate the Alternate Store Location feature's alternate store candidate list.

### fsaltnode

Add, modify or delete Alternate Retrieval Location information in the Tertiary Manager database. The settings made with this command also apply to Alternate Store Location. The information from this command can also be set and displayed by the StorNext GUI as demonstrated in [Alternate Retrieval Location Configuration on page 193](#).

### fsfileinfo

Generate a report about files known to the Tertiary Manager. A field in the output describes if the Alternate Store Location feature is enabled and whether the copy action is pending or has been completed. Information from this command can also be displayed by the StorNext GUI as demonstrated in [Sample File Report on page 195](#).

### fschfiat

Modify the class attributes of a file. This can be used to disable or enable copying on a per-file basis.

## fsclassinfo

Report policy-class processing parameters, associated directory paths, and affinity lists. This command can display the alternate store location state per policy class. Information from this command can also be displayed by the StorNext GUI.

## fsschedlock

Command used for locking/unlocking some automated features. The **altstore** feature type controls the alternate store location feature. There may be a delay of a few minutes after the lockout schedule is reached before the **fs\_altstore** receives the command, completes in-process transfers, and stops processing new transfers.

## fsmodclass

Modify the processing parameters of a policy class. This command can change the alternate store location state per class policy. This information from this command can also be set and displayed by the StorNext GUI.

## fsrcopy

The **-r** option to **fsrcopy** invalidates the remote copy, if it exists, and a new copy is created.

## Log Files

Log, trace, and history information from the **fs\_altstore** resident process is recorded under the `/usr/adic/TSM/logs` directory in the `trace/trace_06`, `tac/tac_00` files, and in the `/usr/adic/TSM/history/hist_06` file.

## Installation, Configuration and Upgrade

### Installation

The Alternate Store Location feature is part of the StorNext system and does not require any extra steps to install. When it is not configured, it runs with negligible impact on system resources. After configuration, it is ready to provide both main- and remote-side functionality of the Alternate Store Location feature.

### Configuration

Following are the steps for configuring both sides of the Alternate Store Location feature:

1. Set up Secure Shell credentials for the transfer script or provide an alternative local transfer script as described in section [User Customizable Transfer Script on page 188](#).

---

**i Note:** It may be necessary to configure network firewalls as described in section [ALTSTORE\\_AGENT\\_PORT\\_NUMBER on page 186](#).

2. Configure the Alternate Retrieval Location parameters using the StorNext GUI, as demonstrated in section [Alternate Retrieval Location Configuration on page 193](#), or by using the `fssaltnode` command as described in section `fssaltnode`.
3. Enable the Alternate Store Location feature in Class Policies using the StorNext GUI, as demonstrated in section [Storage Manager Policy Configuration on page 193](#), or with the `fsmodclass` command as described in section `fsmodclass`.

## Upgrade

There are no special steps required for upgrading StorNext when the Alternate Store Location feature is configured.

## Operational Modes

Operation of the Alternate Store Location feature requires that the main site, the remote site, and the network between them are up and running. Whenever all of these are brought on-line, file transfers will resume automatically after a brief pause, typically 20 seconds.

Operation can be suspended during periods as needed for system performance or other reasons by use of the `fsschedlock` command as described in section `fsschedlock`. When the lockout period ends, there may be a few minutes delay before file transfers resume. The delay is affected by the setting described in `ALTSTORE_POLL_TIME`.

## Monitoring

The backlog of pending file copies can be viewed as described in section `altstoremod`.

The status of copy operations per file can be viewed as described in section `altstoreadd`.

## Copying Pre-Existing Files

Files that existed before their policies had the Alternate Store Location feature enabled are not copied without taking additional steps. It is recommended that the set of these files be added in subsets of the total list, rather than all at once. This allows for optimizing the rate of copy completions by organizing the retrieval of files for a minimum amount of tape handling. When the Alternate Store Location feature encounters a file that is truncated, it will retrieve the file automatically with a small amount of retrieval optimization, but it can end up bottlenecking file copies behind inefficient retrievals when there is a large set of truncated files in the list of files to be copied.



Best practices for the copying of pre-existing files is to:

1. Identify the set of all files you wish to copy that are not already enabled for Alternate Store Location copying.
2. Organize the full set into subsets that can be retrieved efficiently by subset.
3. Retrieve one complete subset of files to disk.
4. After the subset retrieval has completed, submit the subset to the input of the `altstoreadd` command as described in section `altstoreadd`, and start the retrieval of another subset of files.
5. Monitor the completion of background copies as described in section `altstoremod`. When a subset has been copied, truncate the files listed in the subset.

## Disaster Recovery

### Alternate Retrieval Location

The Alternate Retrieval Location feature works well with the Alternate Store Location feature. They share the configuration of the remote node name and remote paths per file system. When retrieving a truncated file that is missing all of its local copies, StorNext will automatically use the Alternate Retrieval Location feature to access the remote copy as a last resort. StorNext compares the expected and actual file sizes of the remote copy as a simple integrity check, and aborts the retrieval if they do not match. StorNext cannot guarantee that the remote copy has not been modified or is not current, but the remote copy should be current in most cases if the main and remote systems have been operating successfully up to the point of the disaster.

When both Alternate Retrieval Location and Alternate Store Location features are configured, a further integrity check is made for alternate retrieves. The retrieve will not proceed for a file when remote copying is enabled for the file and the remote copy has not been stored.

### System Replacement

In the event of a complete loss of the main site with its tertiary storage and backups, the remote site could be swapped into main-site status. With appropriate configuration and operation, the remote site can have a complete set of files with identical pathnames and rwx permissions. There may even be a complete set of per-file versions in tertiary storage if the remote files were overwritten as they changed as described in section `ALTSTORE_VERSIONING`.

## Troubleshooting

Log messages are written as described in section [Log Files on page 190](#). These may provide clues for the diagnosis of operational problems.

The most common problems preventing copies from being made are in the transfer operation. The user-customizable transfer script, as described in section [User Customizable Transfer Script on page 188](#),

depends on the proper setup of authentication credentials and the ability to make a connection through firewalls etc. Communication between the `fs_altstore` resident processes at the main and remote sites is also dependent on communication by TCP through firewalls, etc. as described in section `ALTSTORE_AGENT_PORT_NUMBER`.

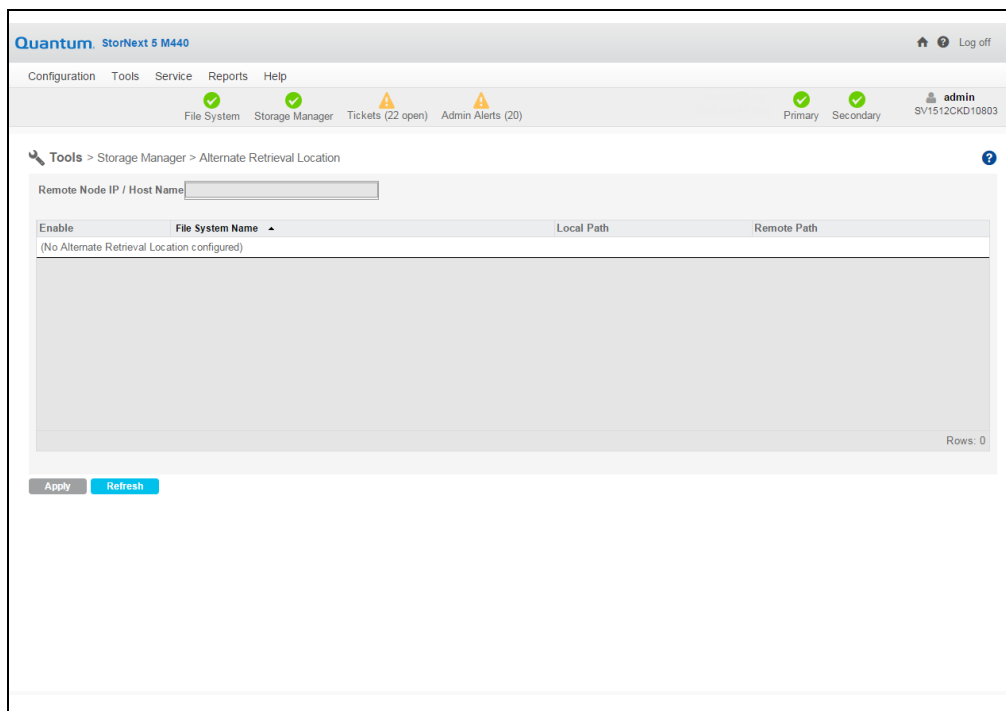
## Examples

### Alternate Retrieval Location Configuration

[Figure 25 below](#) displays the configuration that is shared between the Alternate Retrieval Location and the Alternate Store Location features. It is not possible to configure Alternate Store Location without Alternate Retrieval Location, but it is not necessary to enable the per-file-system options for retrievals.

Configuring the Remote Node, enabling local file systems, and specifying the remote mount-point path per file system are the first steps in configuring both features. The Remote Node - there can only be one - applies to all of the enabled file systems. This information can also be specified on the command line with the `fsaltnode` command.

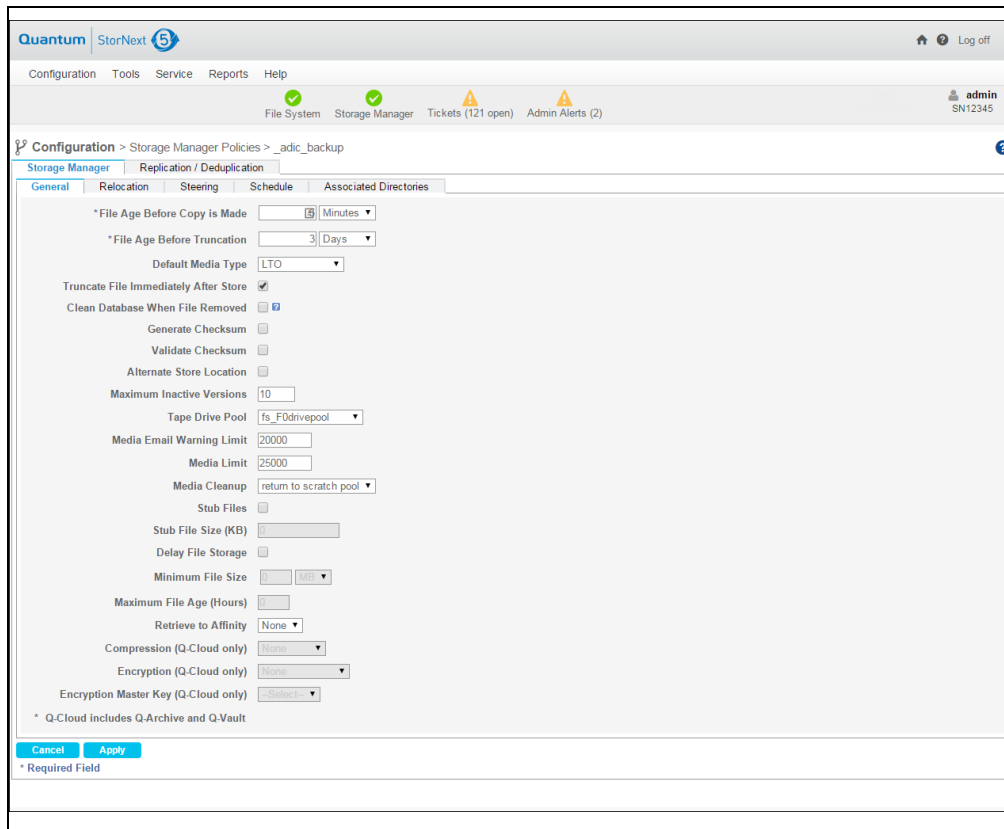
**Figure 25:** Tools > Storage Manager > Alternate Retrieval Location Page



### Storage Manager Policy Configuration

[Figure 26 on the next page](#) displays the Alternate Store Location configuration that is per policy class. Selecting the Alternate Store Location check box enables the feature for future files under the policy. Pre-existing files must be enabled and added to the candidate list individually as described in `altstoreadd`.

Figure 26: Configuration > Storage Manager Policies > General tab Page



## Sample altstoreadd Commands

The following examples show typical uses of the **altstoreadd** command. Additional information about the command is described in **altstoreadd**. The first script recursively finds all the files in a subdirectory hierarchy and submits them as candidates for background copying.

**i Note:** Full pathnames are required. Files that have already been copied will be noted with an error message.

```
cd /stornext/snfs1/policy1/subdir # example directory
find `pwd` -type f | altstoreadd
```

The next script finds all the files in the current directory (non recursive) and asks **altstoreadd** to provide a report on the status of remote copies for those files.

```
cd /stornext/snfs1/policy1/subdir
find `pwd` -maxdepth 1 -type f | altstoreadd -e
```

These two uses of the **altstoreadd** command can be used in a script that submits pre-existing files (those that existed before the Alternate Store Location feature was configured), in an optimal way for retrievals and for the management of primary-disk space.

## Sample fschfiat Command

In the event that remote-copy files have been modified or corrupted in some way, new copies can be made in two steps. The following command turns off the Alternate Store Location feature for all the files under a directory, which causes Stornext to discard its information about the remote copies.

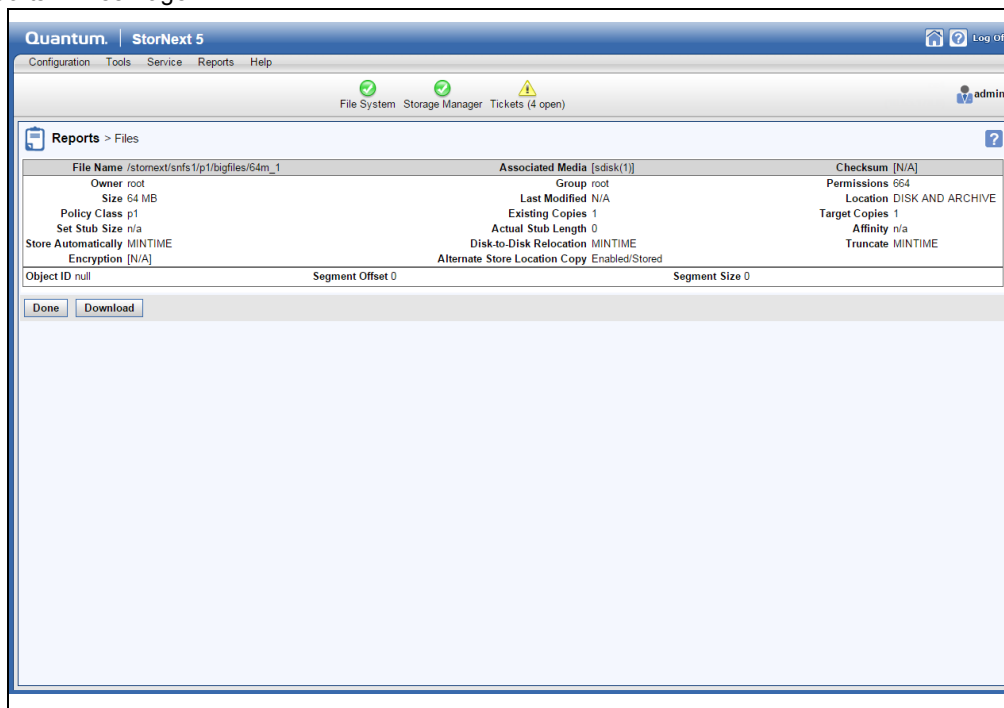
```
fschfiat -A n -R <directory>
```

Following that, use the command `fschfiat` to turn on the feature and cause new copies to be made for the files.

## Sample File Report

[Figure 27 below](#) displays the information that the GUI can report, per file with the Alternate Store Location information.

Figure 27: Reports > Files Page



## Distributed Data Mover

When the **Distributed Data Mover (DDM)** feature is enabled, 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.

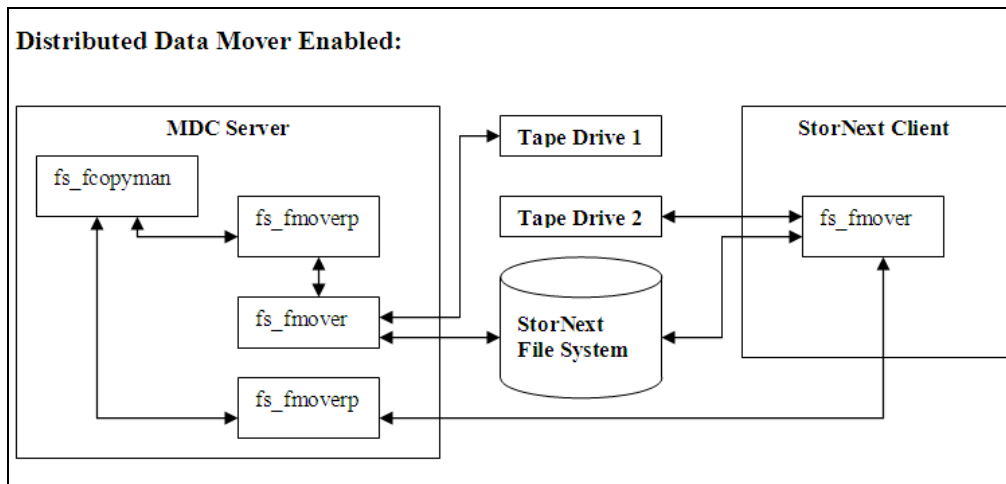
## Overview of DDM

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:



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

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

**⚠ Caution:** Using the StorNext Distributed Data Mover (DDM) feature can boost overall data movement performance to Storage Manager managed tiers by distributing data movement across multiple systems. By default, StorNext software requires the use of SCSI Persistent Reservations on StorNext metadata controllers and DDM clients. As SCSI persistent reservations control access to shared devices, such as tape, and ensure Storage Manager retains control of the tape device paths, even if a failover were to occur. For additional information, see [Using the StorNext Distributed Data Mover \(DDM\) feature can boost overall data movement performance to Storage Manager managed tiers by distributing data movement across multiple systems. By default, StorNext software requires the use of SCSI Persistent Reservations on StorNext metadata controllers and DDM clients. As SCSI persistent reservations control access to shared devices, such as tape, and ensure Storage Manager retains control of the tape device paths, even if a failover were to occur. For additional information, see Using the StorNext Distributed Data Mover \(DDM\) feature can boost overall data movement performance to Storage Manager managed tiers by distributing data movement across multiple systems. By default, StorNext software requires the use of SCSI Persistent Reservations on StorNext metadata controllers and DDM clients. As SCSI persistent reservations control access to shared devices, such as tape, and ensure Storage Manager retains control of the tape device paths, even if a failover were to occur. For additional information, see Using the StorNext Distributed Data Mover \(DDM\) feature can boost overall data movement performance to Storage Manager managed tiers by distributing data movement across multiple systems. By default, StorNext software requires the use of SCSI Persistent Reservations on StorNext metadata controllers and DDM clients. As SCSI persistent reservations control access to shared devices, such as tape, and ensure Storage Manager retains control of the tape device paths, even if a failover were to occur. For additional information, see Tape Devices and Persistent SCSI Reserve. on page 1. above. above.](#)

## 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:

Terminology	Description
<b>Mover</b>	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.
<b>Host</b>	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.
<b>Metadata Controller (MDC)</b>	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.
<b>Remote Client</b>	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_SCSI_RESERVE` in `/usr/adic/TSM/config/fs_sysparm.README` to direct the StorNext Manger to use SCSI-3 persistent reservations.

**⚠ Caution:** Using IBM Advanced Path Failover (APFO) requires SCSI Persistent Reservations are disabled, as device reservations are handled by IBM's software, not StorNext. If you configure a DDM and use the IBM APFO driver, set the "`FS_SCSI_RESERVE=multipath;`" configuration parameter. To configure the parameter, see [System Parameters on page 437](#).

## Verify 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.

1. Register the reservation keys by running the 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
```

2. List the reservation key by running the command:

```
[root@sfx13]# sg_persist -n -k /dev/sg81
```

3. Create reservation by running the command:

```
[root@sfx13]# sg_persist -n -d /dev/sg81 -o -R -T 3 -K 0x123456
```

4. Read reservation by running the command:

```
[root@sfx14]# sg_persist -n -d /dev/sg78 -r
```

5. Preempt reservation by running the command:

```
[root@sfx14]# sg_persist -n -d /dev/sg78 -o -P -T 3 -S 0x123456 -K 0xabcdef
```

6. Release reservation by running the command:

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

7. Delete key by running the 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.

## Install 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 require installing the following Quantum-supplied DDM packages:

- quantum\_curl
- quantum\_openssl
- snfs-mover
- quantum\_unixODBC
- snlfts



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

Follow the installation steps below for each client:

1. Log in as **root**.
2. Download the client with DDM package from the MDC.
3. Install the .rpm files 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:

```
/etc/init.d/cvfs fullstop  
rpm -Uhv *.rpm  
rpm -Uhv snfs*  
service cvfs start
```

## Access the Distributed Data Mover Page

On the **Tools** menu, click **Storage Manager**, and then click **Distributed Data Mover**. The **Tools > Storage Manager > Distributed Data Movers** page appears.

The DDM page displays any previously configured DDM-enabled hosts, managed file systems, tape drives, storage disks, and Object Storage, as well as the current status:

- Enabled
- Not Configured
- Not Enabled
- Internally Disabled

## Configured versus Enabled

Terminology	Description
<b>Configured</b>	Defines 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.
<b>Enabled</b>	Defines 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.

## Enable DDM

When DDM is enabled, data moving responsibilities are distributed among the hosts you specify as described in [Manage DDM Hosts on the next page](#).

1. Next to the **Distributed Data Mover** label, click **Enable** from the list.
2. **(Optional)** Click **Threshold**. Use the **Threshold** option only if you want most data moving operations to run locally on the MDC. When you select the **Threshold** option, the local host (in other words, 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.
3. Click **Apply**.

## Disable DDM

When DDM is disabled, data moving responsibilities rest solely on the primary server.

1. Next to the **Distributed Data Mover** label, click **Disable** from the list.
2. Click **Apply**.

## Manage DDM Hosts

The **Distributed Data Mover** page 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.

---

**i 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.

## Add a Host

1. On the **Distributed Data Mover** page, click **New**. Fields appear where you can enter host information
2. At the **Host** field, enter the host name you are adding.
3. Enter the remaining fields in the upper portion section of the page:

Parameter	Description
<b>Enabled</b>	This option determines whether DDM is currently enabled or disabled on the selected host. Check this box to enable DDM, or uncheck to disable.
<b>MDC Mover Threshold</b>	This option is used to set the threshold for when this host is the MDC. This option is valid only when this host is the MDC and the global settings is set to <b>Threshold</b> . This option defines how many mover process will run on the MDC before it starts remote ones. The default is zero. To change the default value, click the box and enter the desired number.
<b>Max Movers (active MDC)</b>	<p>This option defines the maximum number of simultaneous fs_fmover processes allowed on the host when it is running as the MDC server.</p> <p>There are two settings for this option:</p> <p><b>Unlimited</b> means there is no limit to the number of fs_fmover processes allowed on the host, and <b>Limited</b> sets a maximum number allowed at one time. The default value is <b>Unlimited</b>. To configure for unlimited processes, verify that this option is unchecked. The word <b>Unlimited</b> is displayed next to the check box. To set a limit, click checkbox. A second box will appear next to the checkbox. In this box enter the desired number of fs_fmover process that can be running at the same time on the host.</p>
<b>Max Movers (client or standby MDC)</b>	<p>This option defines the maximum number of simultaneous fs_fmover processes allowed on the host when it is running as a client or the standby MDC server.</p> <p>There are two setting for this option:</p> <p><b>Unlimited</b> means there is no limit to the number of fs_fmover processes allowed on the host, and <b>Limited</b> sets a maximum number allowed at one time. The default value is <b>Unlimited</b>. To configure for unlimited processes, verify that this option is unchecked. The word <b>Unlimited</b> is displayed next to the check box. To set a limit, click checkbox. A second box will appear next to the checkbox. In this box enter the desired number of fs_fmover process that can be running at the same time on the host when it is running as a client or standby MDC server.</p>

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

## Edit a Host or Devices

1. On the **Distributed Data Mover** page, in the **Hosts** table, click the host the host you want to edit.
2. Click **Edit**.
3. Modify the host configuration as desired.
4. If desired, select or remove managed file systems, tape drives and storage disks.

5. Click **Apply** to save your changes.
6. When the confirmation message appears, click **Yes** to proceed or **No** to abort.
7. After a message informs you that your changes were successfully saved, click **OK** to continue.

## Delete a Host

1. On the **Distributed Data Mover** page, in the **Hosts** table, click the host the host you want to delete.
2. 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.
3. When the confirmation message appears, click **Yes** to proceed or **No** to abort.
4. 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

**i 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 [Distributed Data Mover](#).

---

# Drive Replacement

The **Tools** menu's **Drive Replacement** option allows you to update the drive serial number mappings.

## Access the Drive Replacement Page

On the **Tools** menu, click **Storage Manager**, and then click **Drive Replacement**.

The **Tools > Storage Manager > Drive Replacement** page displays the following information:

Parameter	Description
<b>Candidate Drive</b>	The available candidate drive.
<b>Replacement Drive</b>	The available replacement drive.
<b>Cancel</b>	Click <b>Cancel</b> to abort the current operation.
<b>Apply</b>	Click <b>Apply</b> to submit your changes.

---

# Active Vault Policy

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 pro-actively.

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:

Parameter	Description
ACTIVEVAULT_ HIGH_USE	Specifies the percentage of used license capacity at which to begin the Active Vault Policy process (the <b>default</b> value is <b>95.0%</b> ).
ACTIVEVAULT_ LOW_USE	Specifies the percentage of used license capacity at which to stop the Active Vault Policy process (the <b>default</b> value is <b>90.0%</b> ).

**i Note:** 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:

Parameter	Description
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 <b>default</b> value is <b>90.0%</b> ).

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. See [Tools on page 210](#) for additional information.

## Configure Active Vault

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 [Configure External and Internal Vaults below](#) for details on how to configure external and internal vaults.
2. Create and schedule the vault policy; execute 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 [Command Syntax on page 209](#) for examples of some of the commands mentioned in this section.

3. Wait for a 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.
4. 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 [Library Operator Interface on page 177](#) for additional information about the Library Operator Interface.

## Configure External and Internal Vaults

### Configure an External Vault

1. On the **Configuration** menu, click **Storage Destinations**. The **Configuration > Storage Destinations > Libraries** page appears.

2. Click **New...** The **Configuration > Storage Destinations > Libraries > New** page appears.
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.

## 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** page appears.
2. Click **New...** The **Configuration > Storage Destinations > Libraries > New** page appears.
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.

---

**i 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.

---

**i 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:

---

**i 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 Library Management Console page. 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.

---

**i 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 Library Management Console - 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.

### List all scheduled Active Vault policies or a specific policy by name

```
fsschedule [-f activevault | -n name] [-1]
```

### Schedule an Active Vault policy

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

### Modify a scheduled Active Vault policy

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

## Active Vault Policy 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.

---

# System Parameters

The **Tools > Storage Manager > System Parameters** page allows you to set and modify a sub-set of StorNext system parameters for Object Storage.

## Information on the System Parameters Page

Parameter	Description
<b>Name</b>	Displays the name of the system parameter.  For additional information about system parameters, see the <code>fs_sysparm.README</code> file under the installation directory <code>/usr/adic/TSM/config/</code> .
<b>Value</b>	Displays the value of the system parameter. This also serves as a free-form text entry field so that you may update a system parameter value.   <b>Caution:</b> When editing values, ensure that the value you enter for a given system parameter is valid. An invalid value can interfere with proper functioning of Storage Manager. Updating a system parameter value restarts Storage Manager.
<b>Adjusted</b>	Displays <b>Yes</b> , or <b>No</b> , which signifies whether the default Storage Manager system parameter value has been overridden.
<b>Apply</b>	Click to apply the system parameter in the <b>Value</b> field for a given system parameter <b>Name</b> (row).
<b>Reset</b>	Click to clear all the values in the <b>Value</b> field for all system parameters in the table.
<b>Done</b>	Click to exit the <b>System Parameters</b> page.

## Edit or Update a System Parameter Value

1. On the **Tools** menu, click **Storage Manager**, and then click **System Parameters**. The **Tools > Storage Manager > System Parameters** page appears.
2. Select a system parameter under the **Name** column (click a row), and then input a valid system parameter into the text field under the **Value** column. If necessary, click **Reset** to clear all the values in the **Value** field for all system parameters in the table.
3. Click **Apply** to submit your system parameter changes. A confirmation dialog appears.
4. Click **Yes** to confirm your changes, or click **No** to cancel your changes and return to the **System Parameters** page.

## Exit the System Parameters Page

Click **Done** to exit the **System Parameters** page.



# 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 .....	214
Replication Terms and Concepts .....	217
Some Replication Scenarios .....	219
Configuring Replication .....	226
Running Replication Manually (Optional) .....	236
Replication Statuses and Reporting .....	237
Replication Target Relocating Procedures .....	239
Troubleshooting Replication .....	246
Data Deduplication Overview .....	247
Setting Up Deduplication .....	248
Data Deduplication Functions .....	250
Replication / Deduplication Removal Procedures .....	251

# 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*

---

**i 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.

---

**i Note:** If a replication policy is configured on the source side first, any replication request generated before the replication configuration on the target will fail with an error: **Disabled Target**.

If this occurs, even when the target is configured later, then subsequent replication requests will fail. To resolve this, execute the command `snpolicy -replicateforce=path` for each impacted source directory once. Once the replication succeeds, all subsequent replication requests will perform normally.

To avoid replication failure due to a disabled target, Quantum recommends you configure the target policy before you configure the source policy.

---

**i Note:** Replication does not support UID and GID remapping.

A file replicated to the target file system has the same UID and GID as the file in the source file system. If the target file system and source file system are located on different machines, in order to prevent unauthorized access to the replicated files, either:

- (a) The source and target machines must share account UIDs and GIDs for all replicated files and directories; or
- (b) The target machine should disallow regular users access to the replicated files.

Avoid granting users and groups access to files with disjoint UID or GID mappings. If users require access to replicated files, then the user accounts must be configured to share the same user and group account IDs between source and target machines.

Configuring replication is discussed in more detail in the section [Configuring Replication on page 226](#).

## 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.

---

**i 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 displays in simple terms how replicated data flows from the one replication source to one replication target.

**Figure 28:** 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 378](#).

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

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

# Replication Terms and Concepts

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

## Namespace Realization

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

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

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

## Blockpool

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

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

---

**i 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 displays 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.

---

**i 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.

---

**i 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 29:** 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 30:** 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`.

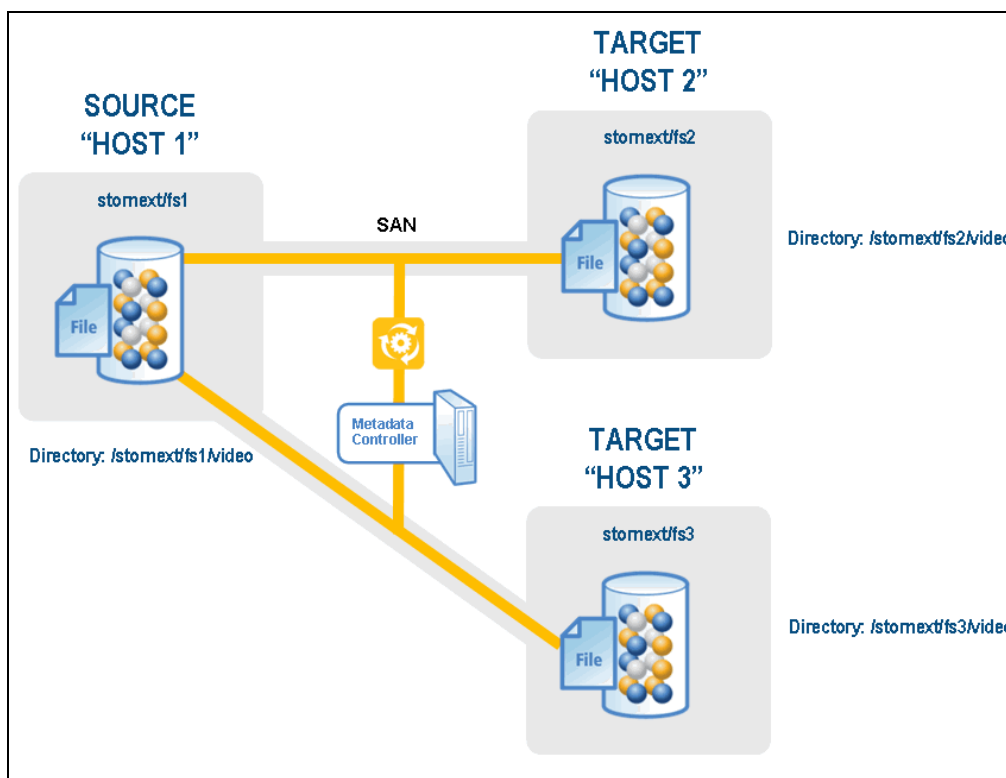
In this scenario we can also use the "Copies to Keep on Target" option. When "Copies to Keep on Target" is specified in a replication source policy, multiple copies are retained in each of the target file systems.

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

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

The following graphic illustrates replication scenario 3.

**Figure 31:** Replication Scenario 3

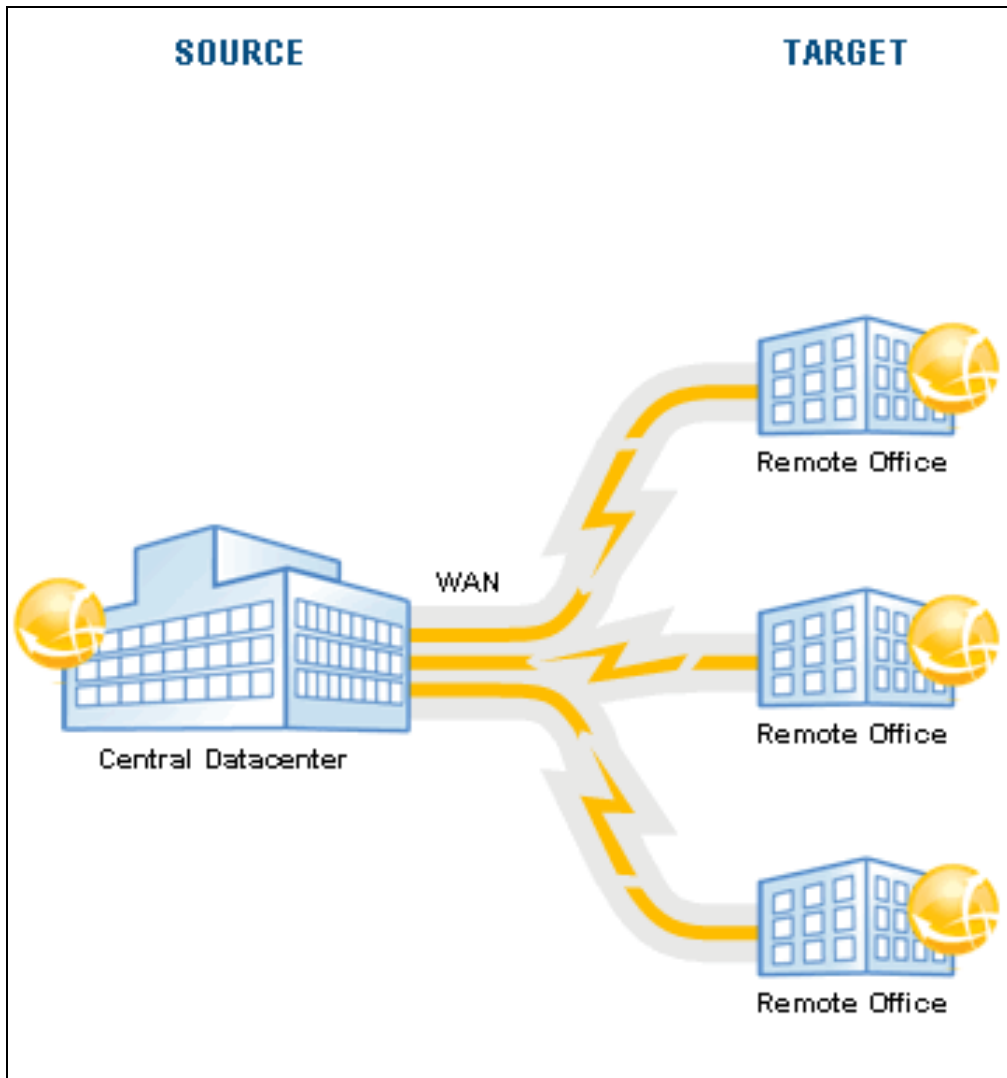


## Additional Replication Possibilities

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 32:** Replicating From One Source to Multiple Targets



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



**Figure 33:** 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.

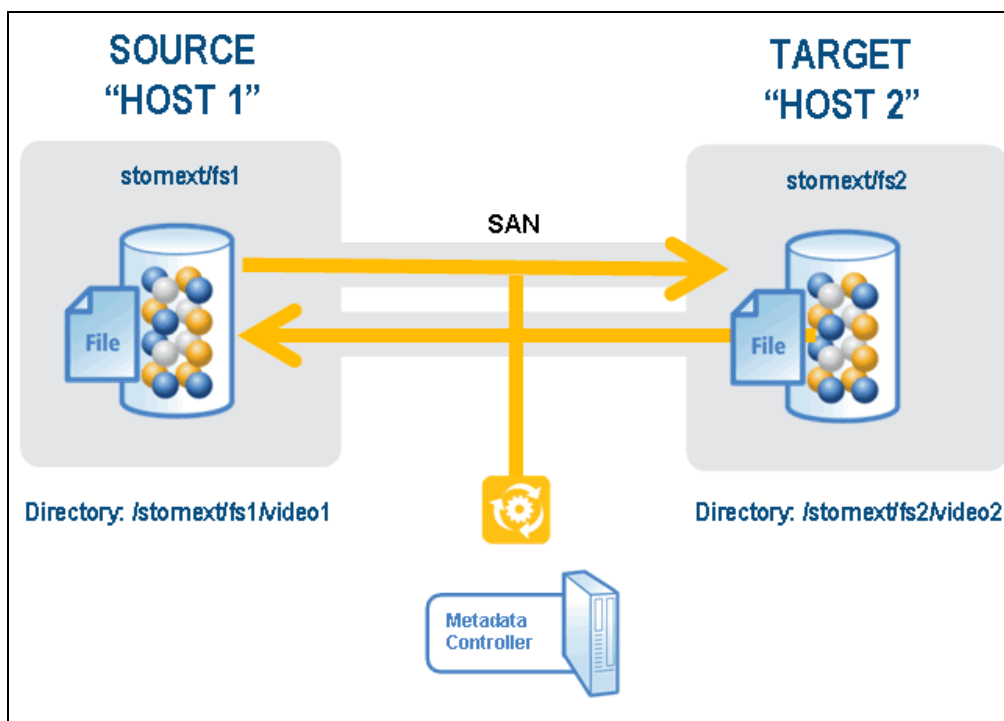
### Example

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 34:** Non-Supported Replication From Source to Target



## “Chained” Replication

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

Chained replication should not be confused with replicating from one source to multiple targets, which *is* supported.

For more information, see [Additional Replication Possibilities on page 222](#).

---

# 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 [The Configuration Wizard on page 20](#).

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

---

**i 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.

---

**i 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

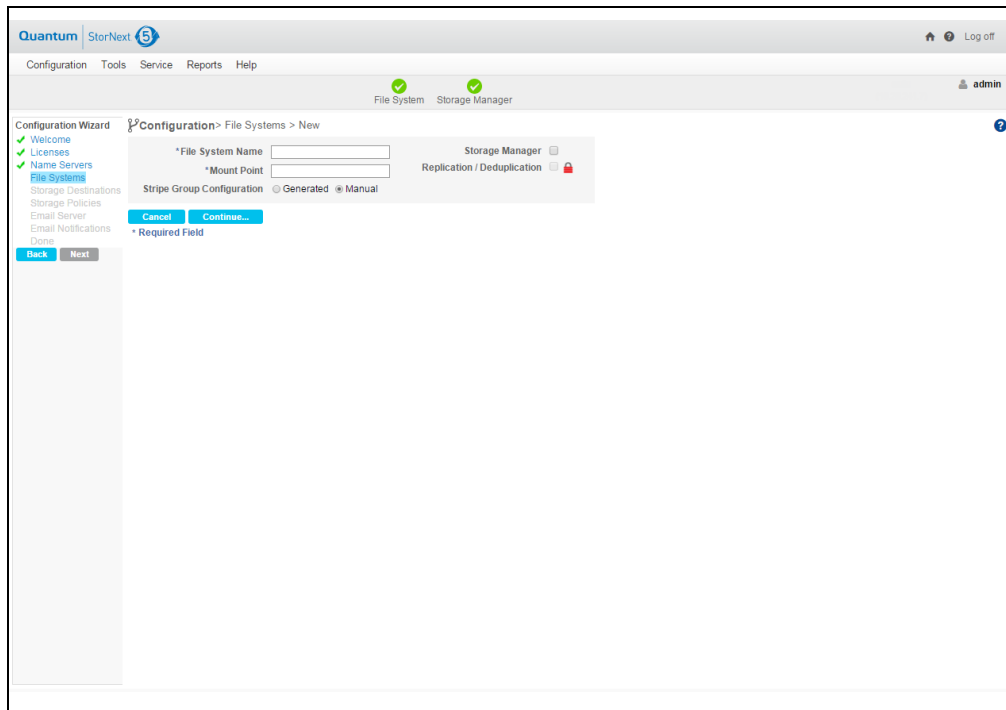
After you complete the first three Configuration Wizard steps, the first replication step is to create file systems: the blockpool file system(s), and the source and target file systems you plan to use.

---

**i 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.

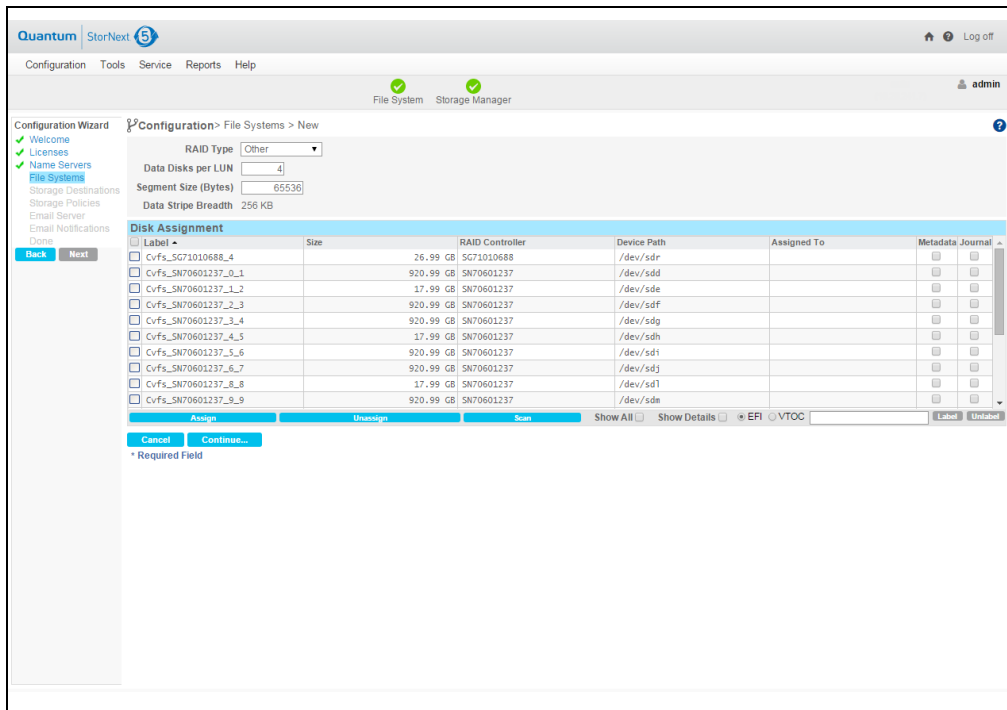
1. If you have not already done so, launch the StorNext Configuration Wizard and proceed through **Welcome**, **License** and **Name Servers** steps to the **File System** step.
2. The **Configuration > File System** page appears.
3. On the Configuration > File System screen, click **New**. The **Configuration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

**Figure 35:** Configuration > File System > New Page



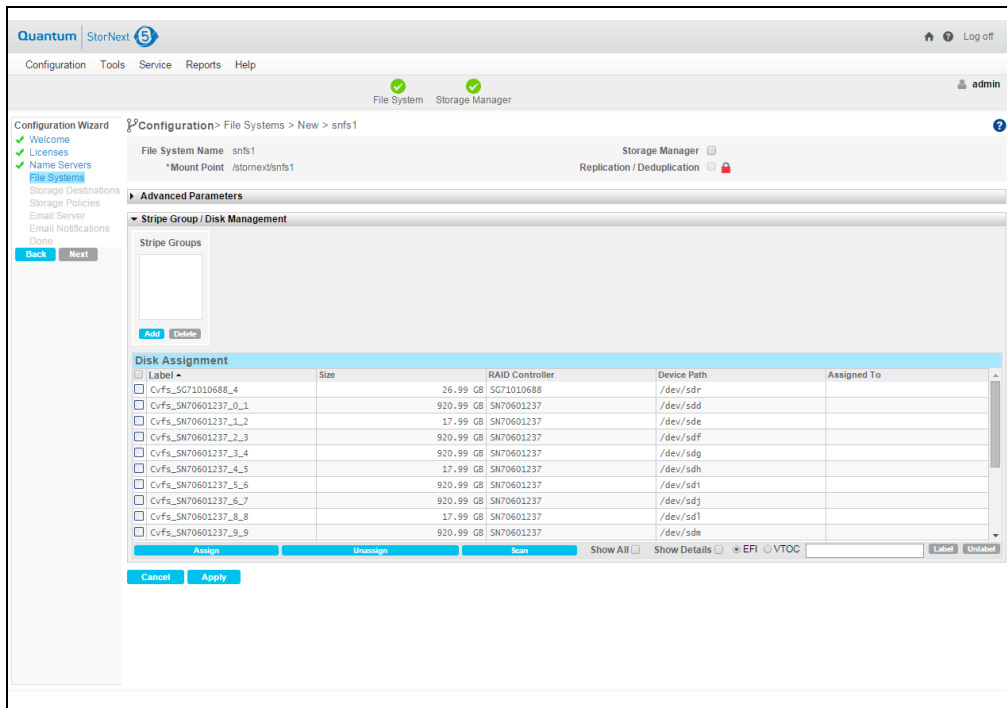
4. At the **File System Name** field enter the name of a file system to be used as a replication source. A default mount-point path automatically appears but you can change this mount point if you wish.
5. Choose the **Replication/Deduplication** option. A warning message alerts you that “A blockpool 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 on page 229](#).
6. Select the **Generate** option, and then click **Continue** to proceed.

Figure 36: Configuration > File System > New Page 2



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

Figure 37: Configuration > File System > New Page 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 [File Systems on page 34](#).

---

**i 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 recommends executing the command `snpolicy_gather -b > some_file` after upgrading to the latest version of StorNext. This saves a copy of your current configuration. The `-b` option creates 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.

---

**i 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.

---

**i 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

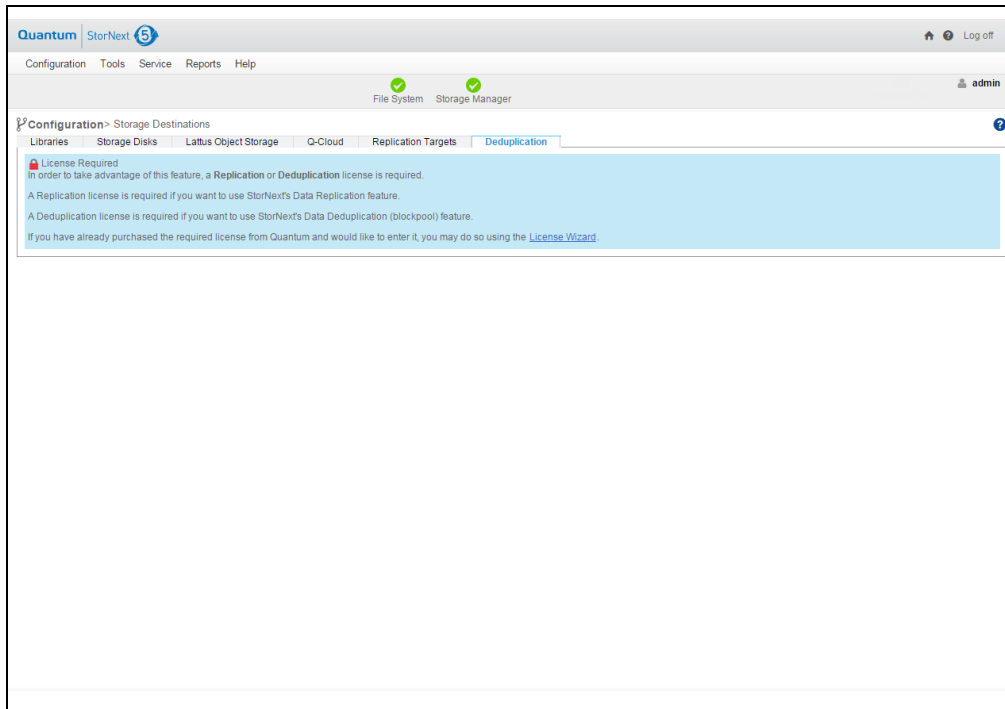
In this step, set up the blockpool on the blockpool file system you just created in the previous step. If the source file system and target file system reside on different machines, perform the following sub steps on both source machine and target machine.

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

**Figure 38:** Configuration > Storage Destinations > Deduplication Page (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.

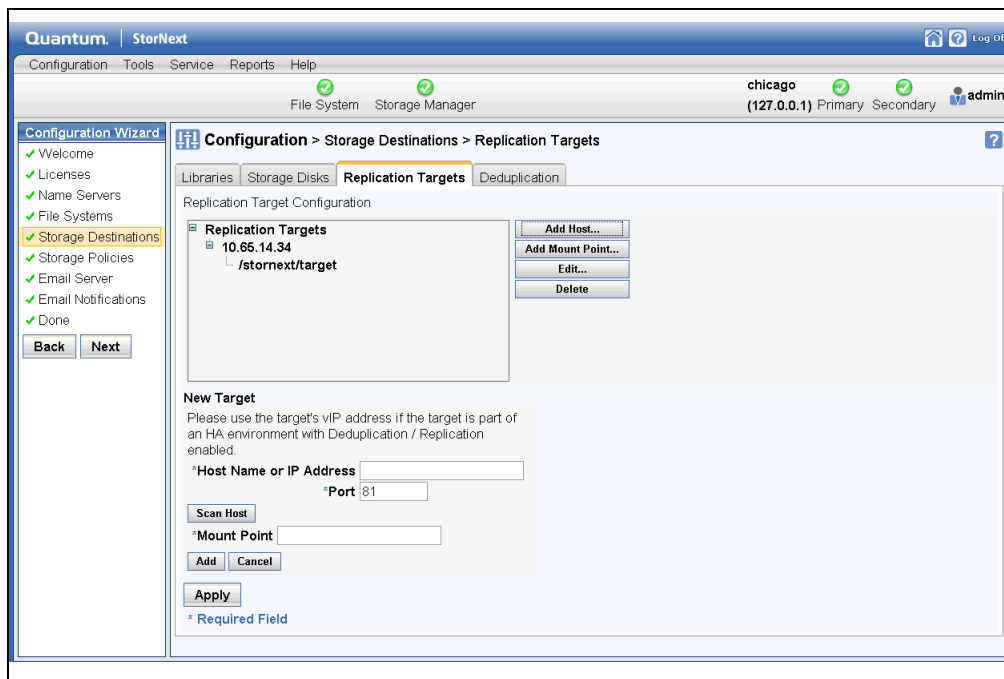
4. After you select the blockpool file system, click **Apply**. A background job is started to create the blockpool.

## 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** page appears.
2. Click **Add Host**.

Figure 39: Storage Destinations > Replication Targets Page



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 on page 226](#), and then click **Add**.
6. Click **Apply**. At this point you should see your file system listed as a replication target.

**i 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 233](#), 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.

See [Adding a Replication or Deduplication Policy on page 105](#).


## 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 40:** 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.

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** page appears.
3. Click the **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 219](#) for an expanded definition.

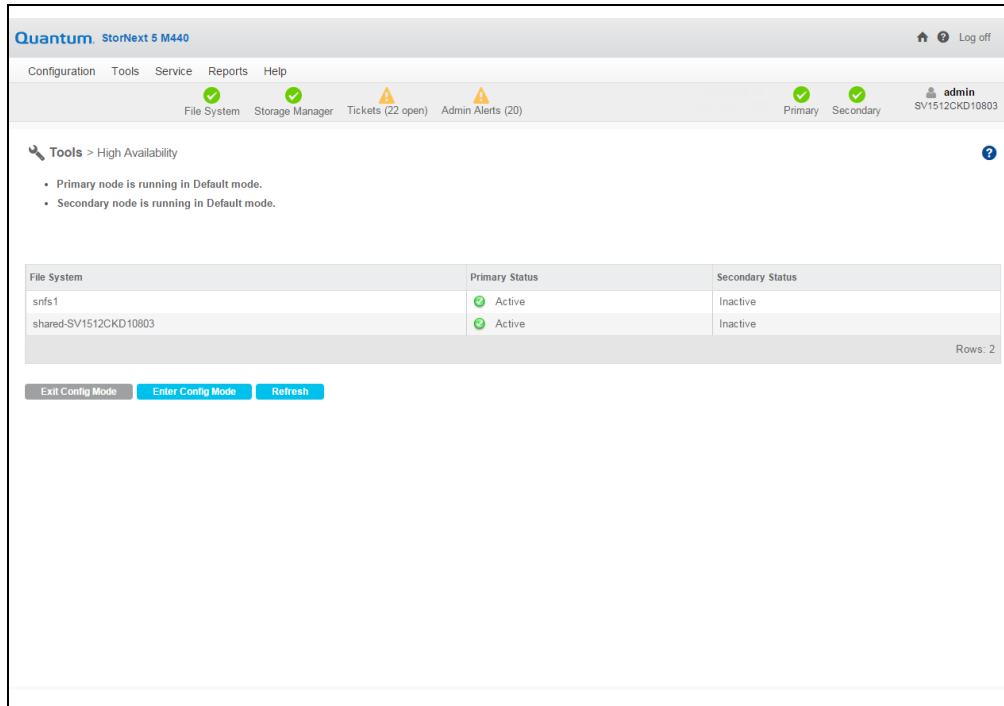
It is easiest to set up the vIP during the initial HA conversion. The vIP configuration items appear automatically at this time if a license exists for replication. It is not necessary to have a replication policy configured.

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 344](#).

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

Figure 41: Tools > HA Convert Page



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.

---

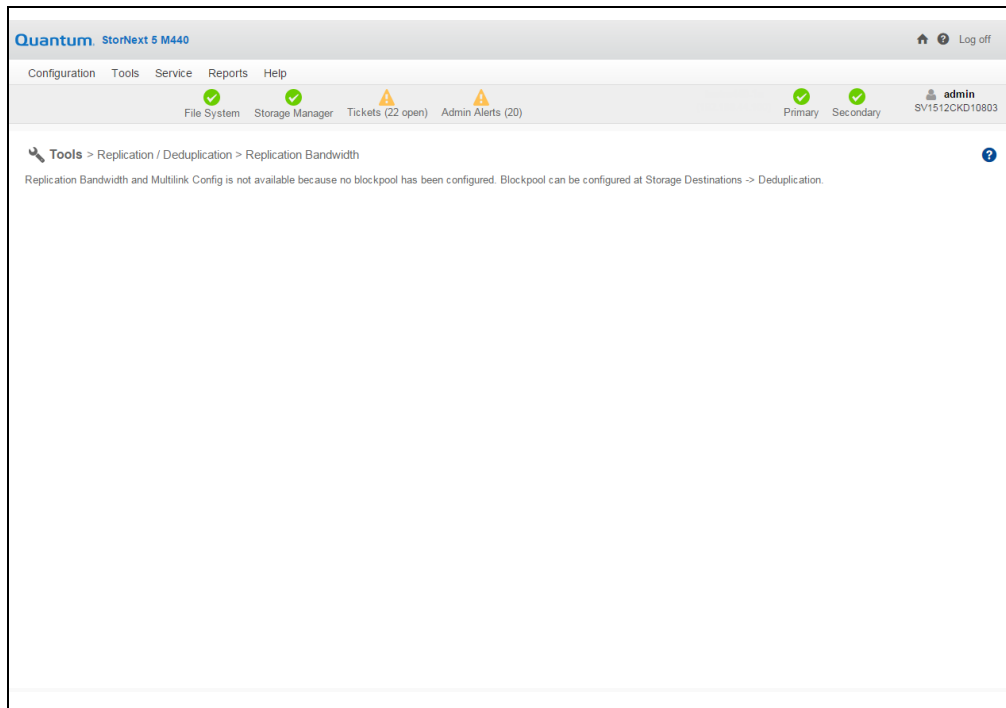
**i Note:** The IP address of the Blockpool Server in HA Clusters can also be configured from the tab **Outbound Replication** and **Inbound Replication** for the "global" policy. There is no difference in configuring from different tabs and the setting of the IP address for the Blockpool from one tab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

Figure 42: Tools > Replication > Bandwidth Page



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 on page 218](#).
  - **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 on page 218](#).
  - **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** page appears. See [Configuration Steps Summary on page 232](#).
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 below](#): View reports showing information pertaining to storage policies and replication targets.
- [Replication Administration below](#): View the current replication status.
- [StorNext Jobs on the next page](#): View currently running StorNext jobs, including replication.

## Replication Reports

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

- The **Policy Activity** report displays replication performance statistics.
- The **Policy Summary** report displays 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.

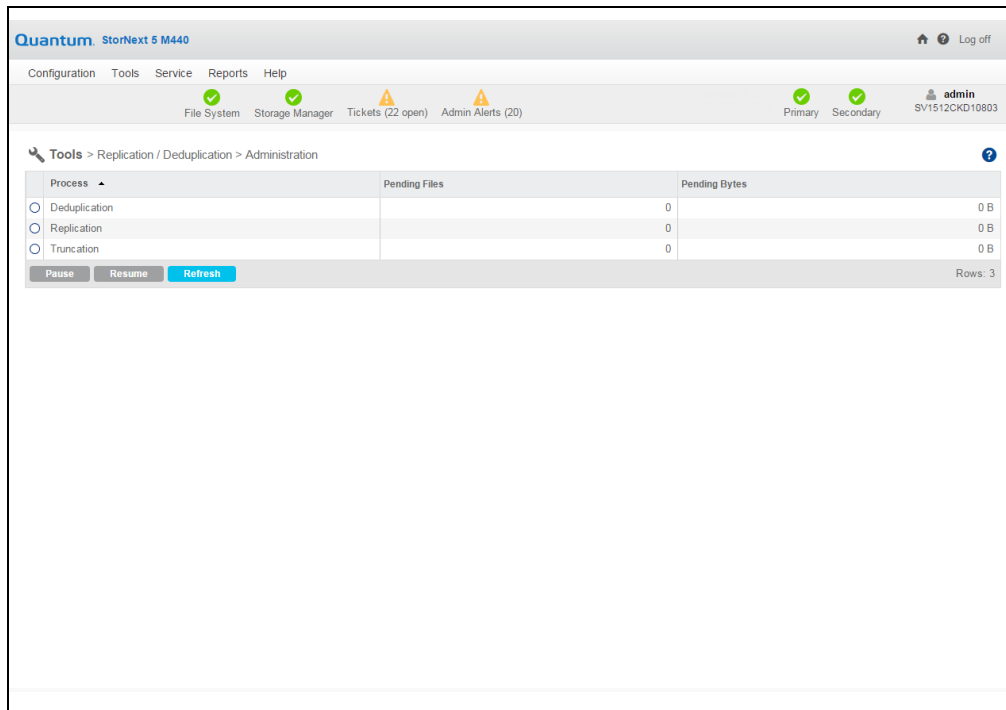
For more information about replication reports, see [Replication / Deduplication Reports on page 378](#).

## 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** page appears.

**Figure 43:** Tools > Replication/Deduplication > Administration Page



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

## Pause and Resume 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, click **Refresh** to update the status shown on the **Administration** page.

## StorNext Jobs

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

To view jobs, on the **Reports** menu, click **Jobs**. The **Reports > Jobs** report is displayed.

For more information about StorNext jobs, see [Jobs on page 365](#).

# 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.

## 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 or to Object Storage, 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 displays 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 on the next page](#)
2. [Relocating the Target File System to the Target Site on page 244](#)

**i 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 [Target Relocating Procedures on page 241](#) 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/fsm1ist` 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
  4. 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**.
  5. Add a new target to replication targets.
  6. 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.
  7. 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.

---

**i 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**.

---

**i 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**.

#### 8. 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
```

---

**i 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 [Target Relocating Procedures on page 241](#) 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 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 [Target Relocating Procedures on page 241](#).
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 [Target Relocating Procedures on page 241](#) 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 244](#) 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 628](#).

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.

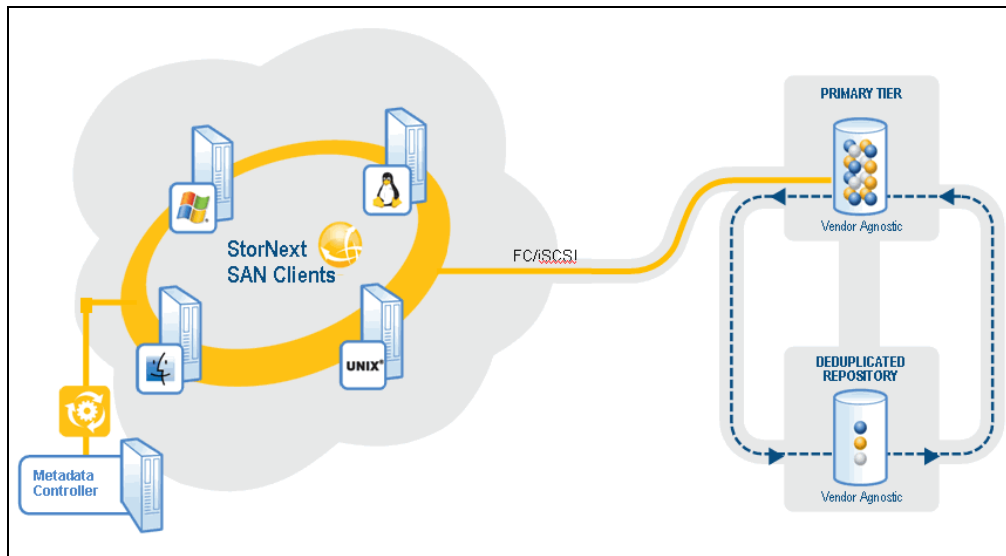
If several files contain the same blob, only one copy is stored in the blockpool.

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 44: 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 the tasks below to set up and enable deduplication:

## Create 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 [File Systems on page 34](#).

## Specify 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.

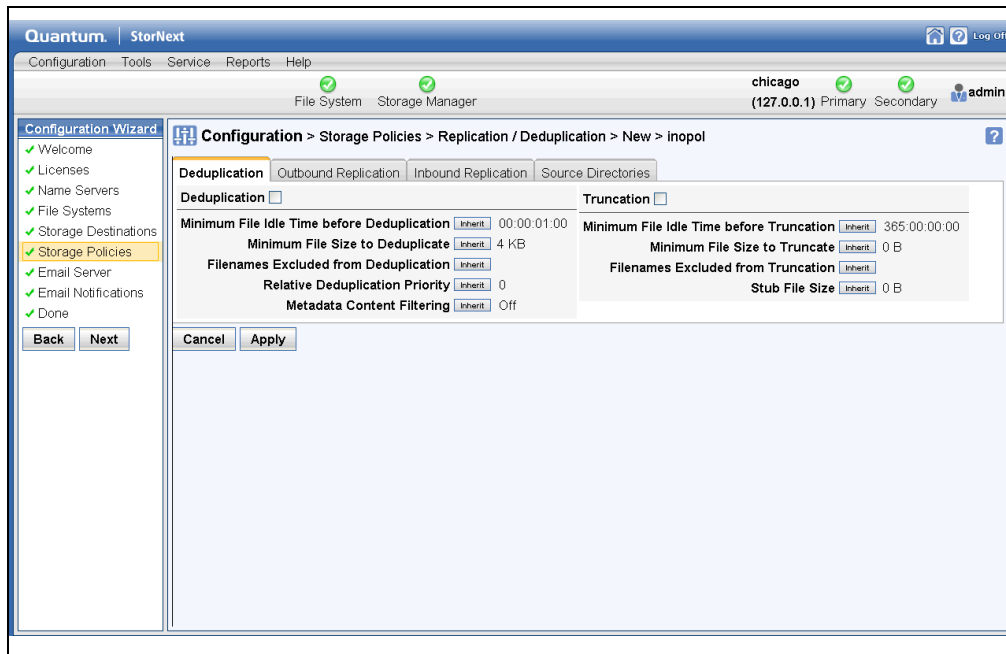
The process for specifying a blockpool for deduplication is identical to specifying a blockpool for replication. For more information, see [Step 2: Setting up the Blockpool on page 229](#) in the Configuring Replication section.

## Create 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** page appears.
2. If you are creating a new policy, click **New**. The **Storage Policies > New** page appears. See [Step 4: Create a Replication Storage Policy on page 232](#). 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:
  - a. **Policy Class**: The name of the new policy you are creating
  - b. **Policy Type**: choose **Replication /Deduplication** to create a deduplication storage policy.
4. Click **Configure**. The **Replication /Deduplication Policy** page appears.

Figure 45: Replication/Deduplication Policy Page



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 **Inherit** beside the desired field to enter your own values. For information about what to enter at each field, see the online help.

Below are some 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 [Replication Administration on page 237](#).

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 [Replication Administration on page 237](#).

## Deduplication Reports

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

- The **Policy Activity** report displays deduplication performance statistics.
- The **Policy Summary** report displays 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 378](#).

---

# 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 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.

## 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?

---

**i 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 the Command Line

### Obtain snpolicy-managed file systems

```
# /usr/cvfs/bin/snpolicy -listfilesystems=localhost  
fsname: snfs1 [replication dedup] up 110:49:07  
mount: /stornext/snfs1  
blockpool: Running up 110:49:07
```

### Obtain policy information

```
# /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.

## View the policy configuration

```
# /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.

## Check if a directory is associated with a TSM relation point

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

This will show the TSM policy class if the directory is associated with a TSM relation point. 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 the Command Line on page 253](#), 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).

**i 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 261](#).

Step	Action	Description
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: <pre># /usr/cvfs/bin/snpolicy_gather &amp;&gt;snpolicy_dump</pre> The configuration information is saved to file <code>snpolicy_dump</code> .



Step	Action	Description
2	Suspend Replication/Deduplication Activities	<p>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 <b>mnt_path</b> is the mount path of the snpolicy-managed file system.</p> <pre># /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 -rundelete=mnt_path -suspend # /usr/cvfs/bin/snpolicy -compact=mnt_path -suspend</pre> <p>In addition, disable <b>rep_input</b> by setting <b>rep_input</b> to “off” of policy “target”.</p> <pre># /usr/cvfs/bin/snpolicy -updatepolicy=mnt_path -name=target -policy='rep_input=false'</pre>

Step	Action	Description
3	Namespace and Private Directory Removal	<p>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.</p> <p>a. Obtain the realized namespace directories and target keys. Run:</p> <pre># /usr/cvfs/bin/snpolicy -listrepcopies=mnt_path # /usr/cvfs/bin/snpolicy - listrepcopies=/stornext/tgt1 source://stornext/snfs1@10.65.189.57:14500/test?key=2 231941 -&gt; target://stornext/tgt1@10.65.189.39:?key=350 0 -&gt; (Not Realized) source://stornext/snfs1@10.65.189.57:14500/test?key=9 465955003 -&gt; target://stornext/tgt1@10.65.189.39:?key=430 0 -&gt; /stornext/tgt1/repository/test</pre> <p>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.</p> <p>b. Determine whether the realized namespace is under a TSM relation point. Run:</p> <pre># /usr/adic/TSM/bin/fsdirclass /stornext/tgt1/repository/test FS0070 27 1002737568 fsdirclass completed: /stornext/tgt1/repository/test located in class smpol1</pre> <p>The output of the command indicates that the realized namespace was associated with a TSM relation point.</p> <p>c. Remove realized namespaces</p> <hr/> <p><b>i Note:</b> This is for StorNext 4.1 <b>only</b>. Skip this step if your StorNext version is 4.2.0 or greater.</p> <p>For each realized replication stream, run:</p>

Step	Action	Description
------	--------	-------------

```
# /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:

```
# /usr/cvfs/bin/snpolicy -repcleanup=/stornext/tgt1 -  
key=2162719  
I [0127 11:04:26.790013 25224] Removed 0 old files
```

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
```

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 displays 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.

Step	Action	Description
		<pre># /usr/cvfs/bin/snpolicy -listpolicies=/stornext/tgt1 NAME: default NAME: global inherits from: default NAME: target inherits from: global</pre> <p>The output above indicates that there is no policy or replication directory that has not been cleaned up.</p>
4	Stop Snpolicy Daemon and Blockpool Server	<p>In this step, the snpolicy daemon and blockpool server need to be stopped. Run the following commands:</p> <pre># /usr/cvfs/bin/cvadmin -e "stopd snpolicyd" # /usr/cvfs/bin/bp_stop</pre> <p>Check whether the process “snpolicyd” and “blockpool” have been stopped. You may run command “<b>ps -ef</b>” to check active processes. If the above commands still cannot stop them, try to kill them by running command “<b>kill -9 pid</b>”.</p>
5	Remove Snpolicy-managed Private Directories	<p>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:</p> <pre># /bin/rm -rf mnt_path/.rep_private</pre> <p>For each relation point that has realized namespace landed under it, run</p> <pre># /bin/rm -rf relation_point_path/.rep_private</pre>
6	Remove Replication History Logs	<p><b>i Note:</b> Skip this step if you want to retain the Replication History log files.</p> <p>Run the following command where <b>fsname</b> is the file system name of the snpolicy-managed file system.</p> <pre># /bin/rm -rf /usr/cvfs/data/fsname/rep_reports/* # /bin/rm -rf /usr/cvfs/data/fsname/policy_history</pre>

Step	Action	Description
7	Remove Related Event Files	<p>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.</p> <pre># /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</pre>
8	Remove Blockpool and its Configurations	<p>This step removes a blockpool repository and its configuration files. As mentioned in section <a href="#">Obtain Information from the Command Line on page 253</a>, the blockpool repository path can be found from file /usr/cvfs/config/blockpool_root. Run the following commands:</p> <pre># /bin/rm -rf blockpool_repository_path # /bin/rm -f /usr/cvfs/config/blockpool_root # /bin/rm -f /usr/cvfs/config/blockpool_config.txt</pre> <p>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.</p>


Step	Action	Description
9	Turn off the Snpolicy-managed Attribute in File System Configuration File	<p>This step turns off the “snpolicy-managed” attribute of the currently snpolicy-managed file systems. Perform the following for each snpolicy-managed file system:</p> <ol style="list-style-type: none"><li>Unmount the snpolicy-managed file system</li></ol> <pre># /bin/umount /stornext/tgt1</pre> <ol style="list-style-type: none"><li>Stop file system</li></ol> <pre># /usr/cvfs/bin/cvadmin -e "stop tgt1"</pre> <ol style="list-style-type: none"><li>Update the file system configuration file /usr/cvfs/config/tgt1.cfgx</li></ol> <pre># /bin/sed -e 's/&lt;snfs:snPolicy&gt;true/&lt;snfs:snPolicy&gt;false/g' /usr/cvfs/config/tgt1.cfgx &gt;tgt1.tmp # /bin/mv tgt1.tmp /usr/cvfs/config/tgt1.cfgx</pre> <ol style="list-style-type: none"><li>Start file system</li></ol> <pre># /usr/cvfs/bin/cvadmin -e "start tgt1"</pre> <ol style="list-style-type: none"><li>Mount file system</li></ol> <pre># /bin/mount -t cvfs tgt1 /stornext/tgt1</pre>
10	Restart the StorNext GUI	<p>The StorNext GUI service needs to be restarted in order to view the changed configurations. Run:</p> <pre># /sbin/service stornext_web restart</pre>

## 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 252](#).

Step	Action	Description
1	<b>Backup Replication/Deduplication Configurations</b>	This saves the replication/deduplication configurations on the source host. See <a href="#">Replication Removal on a Target Host on page 255</a> on how to back up.
2	<b>Suspend Replication/Deduplication Activities</b>	This is similar to <a href="#">Replication Removal on a Target Host on page 255</a> except that there is no need to change policy “target”.   <b>Note:</b> If you have multiple snpolicy-managed file systems, you must stop replication/deduplication activities for each file system.

Step	Action	Description
3	<b>Remove Policy from Snpolicy-managed Directories</b>	<p>For each snpolicy-managed directory obtained from section <a href="#">Obtain Information from the Command Line on page 253</a> (from the output of snpolicy command <b>listpolicies</b>), 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 <b>dir_path</b> is the path of the snpolicy-managed directory:</p> <pre># /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</pre> <p><b>Note:</b> If you have multiple snpolicy-managed file systems, run this command for every snpolicy-managed directory on each snpolicy-managed file system.</p> <p>After all snpolicy-managed directories have been removed the associated policies, run snpolicy command <b>listpolicies</b> to ensure no directory is left.</p> <pre>[root@ylu-rep-src1 ylu]\$ snpolicy - listpolicies=/stornext/snfs1 NAME: default NAME: global inherits from: default NAME: target inherits from: global NAME: rep_poll inherits from: global</pre>
4	<b>Remove Replication Targets from StorNext GUI</b>	In the StorNext GUI, click <b>Configuration &gt; Storage Destinations &gt; Replication Targets</b> , and delete all replication targets defined there.
5	<b>Stop Snpolicy Daemon and Blockpool Server</b>	Stop the snpolicy daemon and blockpool server on the source host. See <a href="#">Replication Removal on a Target Host on page 255</a> for details.
6	<b>Remove Snpolicy-managed Private Directories</b>	<p>This step removes the snpolicy-managed private directories. See <a href="#">Remove Snpolicy-managed Private Directories on page 259</a> for details.</p> <p><b>Note:</b> 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 <a href="#">Remove Snpolicy-managed Private Directories on page 259</a>. Also, the private directories for all snpolicy-managed file systems must be removed.</p>



Step	Action	Description
7	<b>Remove Replication History Logs</b>	See <a href="#">Remove Replication History Logs on page 259</a> for details.
8	<b>Remove Related Event Files</b>	See <a href="#">Remove Related Event Files on page 260</a> for details.
9	<b>Remove Blockpool and its Configurations</b>	See <a href="#">Remove Blockpool and its Configurations on page 260</a> for details.
10	<b>Turn off “Snpolicy-managed” Attribute in File System Configuration File</b>	See <a href="#">Turn off the Snpolicy-managed Attribute in File System Configuration File on page 261</a> for details.
11	<b>Restart StorNext GUI</b>	See <a href="#">Restart the StorNext GUI on page 261</a> for details.



# Chapter 7: Tools Menu Functions

---

This chapter contains the following topics:

Tools Menu Overview .....	266
User Accounts .....	268
Client Download .....	272
Firmware Upgrade .....	285
Appliance Release Notes .....	285
Dynamic Application Environment .....	287
System Control .....	290
Lattus Certificates .....	291
File and Directory Actions .....	303
S3 Buckets .....	312
StorNext Connect Ports .....	314
File Systems Overview .....	315
Storage Manager .....	324
Replication and Deduplication .....	324
High Availability (HA) .....	324
Upgrade Firmware .....	324

# Tools Menu Overview

The **Tools** menu contains options to control day-to-day operations of StorNext.

Menu Item	Description
<b>User Accounts</b>	Control user access to StorNext tasks.
<b>Client Download</b>	Download SNFS client software.
<b>Firmware Upgrade</b>	Perform a firmware upgrade for your HA system. <b>i Note:</b> This option is only visible on StorNext M660, M440, M330, Pro Foundation, Artico Metadata Appliances, and Xcellis Workflow Director systems.
<b>Dynamic Application Environment</b>	Displays the current status of your virtual machine environments. <b>i Note:</b> This option is only visible on StorNext M660, M440, Pro Foundation, Artico Metadata Appliances, and Xcellis Workflow Director systems.
<b>System Control</b>	Stop or start the file system or StorNext Storage Monitor, and specify whether to automatically start StorNext at system startup.
<b>Object Storage Certificates</b>	View, create, import, convert, download, and delete Object Storage certificates.
<b>File and Directory Actions</b>	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.
<b>S3 Buckets</b>	Enables you to scan for, add new, and delete existing S3 buckets.
<b>StorNext Connect Ports</b>	Enables you to scan for, add new, and delete existing S3 buckets. <b>i Note:</b> This option is only visible on StorNext M660, M440, Pro Foundation, Artico Metadata Appliances, and Xcellis Workflow Director systems.
<b>File Systems: Label Disks</b>	Label disk drives.
<b>File Systems: Check File System</b>	Run a check on your file system before expanding the file system or migrating a stripe group.
<b>File Systems: Affinities</b>	Configure affinities for your file system.
<b>File Systems: Migrate Data</b>	Migrate the file system's stripe group(s).

Menu Item	Description
<b>File Systems: Truncation Parameters</b>	Manage the file system's truncation parameters.
<b>File Systems: Manage Quotas</b>	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.
<b>Storage Manager: Storage Components</b>	View current status for libraries, storage disks, and tape drives; place one or more of these components online or offline.
<b>Storage Manager: Drive Pools</b>	Add, modify, or delete drive pools.
<b>Storage Manager: Media Actions</b>	Remove media from a library or move media from one library to another.
<b>Storage Manager: Storage Exclusions</b>	Specify types of file names to exclude from StorNext Storage Manager.
<b>Storage Manager: Truncation Exclusions</b>	Specify files or directories to exclude from the truncation process.
<b>Storage Manager: Tape Consolidation</b>	Enter parameters for automatically consolidating space on tape media.
<b>Storage Manager: Library Operator Interface</b>	Enter or eject media from the Library Operator Interface.
<b>Storage Manager: Software Requests</b>	View or cancel pending software requests.
<b>Storage Manager: Scheduler</b>	Schedule file system events including Clean Info, Clean Versions, Full Backup, Partial Backup, and Rebuild Policy.
<b>Storage Manager: Alternate Store and Retrieval Location</b>	<b>Alternate Retrieval Location</b> allows you to specify a remote retrieval location to use in situations where files stored on tape or a storage disk cannot be accessed. <b>Alternate Store Location</b> provides an automatic system for copying files from a main instance of StorNext to a remote instance of StorNext at the same time as copies are made to tertiary storage at the main site.
<b>Storage Manager: Distributed Data Mover</b>	Spread the distribution of data across several machines rather than the primary server.
<b>Storage Manager: Drive Replacement</b>	Allows you to update the drive serial number mappings.
<b>Storage Manager: Client-side Encryption</b>	Lists the master keys that can be used for client side encryption.

Menu Item	Description
<b>Storage Manager: System Parameters</b>	Allows you to set and modify StorNext system parameters.
<b>Replication/Deduplication: Administration</b>	View current replication process, or pause, resume, or stop replication.
<b>Replication/Deduplication: Replication Targets</b>	Add a host or directory for data replication, or edit existing replication targets.
<b>Replication/Deduplication: Replication Bandwidth</b>	Configure replication bandwidth limits and multi-link.
<b>High Availability: Convert</b>	Convert to a high availability configuration.
<b>High Availability: Manage</b>	Manage High Availability system parameters.

---

## User Accounts

The **User Accounts** option on the **Tools** menu allows you to add new StorNext users and modify permissions for existing users. **User Accounts** is also where you change the administrator's password. At any time you can update the information displayed by clicking **Refresh**.

### Add a New User

1. On the **Tools** menu, click **User Accounts**. The **Tools >User Accounts** page appears. All existing users and the admin are shown.
2. Click **New**. The **Tools >User Accounts > New** page appears.
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.
  1. In the **Confirm Password** field, re-type the **Password**.
5. In the **Session Timeout** field, type a number, and then select Minutes or Hours as unit of time measurement from the drop-down. The **Session Timeout** specifies the predetermined amount of time that should elapse before the user is logged out of the system. The default is 30 minutes, and the valid range is from 10 minutes to 12 hours.

---

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

6. **(Optional)** Select the **Change Password Policy** check box to edit the password policy. The default password policy must contain at least 1 character. By changing the default password policy, you can add more options such as **New Minimum Length**, **Uppercase Letter**, **Lowercase Letter**, and **Digit or Special Character** to the password policy. As you modify the password policy, the **New Password Policy** text is updated to reflect your changes.

**i Note:** Only accounts with the **Manage Password Policy** access privilege can edit the password policy. Any changes only affect future password updates.

7. 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**, **General User Defaults** or **Web Services**. Selecting one of these roles for the new user makes it easy for you to automatically add or remove functions.
8. Select all the different roles you want the new user to have:

Function	Privilege
Admin	Download Client Software
	Firmware Upgrade
	Manage Email Notifications
	Manage Email Server
	Manage File Systems
	Manage HA
	Manage Licenses
	Manage Logging
	Manage Network Settings
	Manage Object Storage Settings
	Manage Password Policy
	Manage Replication/Deduplication
	Manage Users
	Set Date and Time
	Stop/Start System Components

Function	Privilege
<b>Operator</b>	Cancel Software Requests
	Manage Admin Alerts
	Manage Backups
	Manage Drive Pools
	Manage Libraries
	Manage Storage Disks
	Manage Storage Manager
	Manage Tickets
	Perform File Actions
	Perform Library Operator Actions
	Perform Media Actions
	Perform Storage Component Actions
	Run Capture State
	Run Health Checks
<b>General User</b>	Run Reports
	Use Web Services

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**.

## View a User

1. From the **Tools > User Accounts** page, select the user whose information you want to view, and then click **View**. A page displays the parameters for the selected user.
2. When you are finished viewing user profile information, click **Back** to return to the **User Accounts** page.

## Edit a User

1. From the **Tools > User Accounts** page, select the user whose information you want to modify, and

then click **Edit**.

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

---

**i 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**.

## Delete a User

1. From the **Tools >User Accounts** page, select the user you want to delete, and then click **Delete**.
2. When the confirmation message appears, click **Yes** to proceed, or **No** to return to the **User Accounts > [admin name]** page without saving.

---

**i 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 a User Account

This feature allows you to enable or disable any GUI accounts on the system, including the default **admin** and **service** user accounts. The **Enable** and **Disable** features require 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.

---

**! Caution:** 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 enabled. To restore GUI access, contact Quantum Technical Support (see [Preface on page xxi](#)).

---

**i Note:** 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
```

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.



---

# Client Download

The StorNext client software lets you mount and work with StorNext file systems.

---

**i Note:** 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 491](#).

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 on page 195](#).

---

**i Note:** To ensure successful operation, before you install the client software verify that the client system meets all operating system and hardware requirements listed in the *StorNext 5 Compatibility Guide*.

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

After downloading the client software, install and configure it using the appropriate method for your operating system.

## Downloading Client Software

The client software can be downloaded from the StorNext installation DVD, the Web, or from a metadata controller (MDC) running the StorNext GUI. You can only download client software from an MDC if you have SNFS and SNSM installed and are running the StorNext GUI. You cannot download client software from a Red Hat 4 or Windows MDC, or from a file-system-only Red Hat 5 MDC.

### Download Client Software from the Installation DVD or Web

If your StorNext server is not running the StorNext GUI, you must download client software from the installation DVD or the Web.

If you have the installation DVD, retrieve the client software from the `fs_only` directory on the DVD. Likewise, if you are obtaining software from the FTP site, download the client software from the `fs_only` directory.

The names of the client install files begin with `snfs_client_<system>`, where `<system>` is the name of the operating system for which the install file is intended. Copy the appropriate install file to an accessible location on your computer and proceed to the installation instructions.

### Download Client Software from an MDC

To download the client software from an MDC, the client system must have network access to the MDC.

1. On the client system, point a web browser to the URL (host name and port number) of the MDC. For example, **http://servername:81**. Use one of the following web browsers to access the MDC (make


sure pop-up blockers are turned off):

- Internet Explorer 7.x or later 8.x
  - Mozilla Firefox 3.x
2. When prompted, type the username and password for the MDC, and then click **OK**. (The default username is **admin**, and the default password is **password**.) The StorNext home page appears.
  3. Choose **Client Download** from the **Tools** menu. The **Tools > Client Download** page appears.
  4. Select from the **Platform** list the desired operating system.
  5. Select from the **OS Version** list the desired operating system version corresponding to the platform you selected.
  6. When a window appears containing a link to the client software download location, click the link to begin downloading.
  7. Click **Download** to begin the process.
  8. When prompted, choose the **Save to Disk** option, and then click **OK**.
  9. Browse to the location where you want to save the file, and then click **Save**.
  10. After the client package has been saved, click **Done**.
  11. Continue with the installation procedure for your operating system as described in the following sections.
    - [Installing the Client on Linux Machines below](#)
    - [Installing the Client on Sun Solaris Machines on page 277](#)
    - [Installing the Client on HP-UX Machines on page 279](#)
    - [Installing the Client on IBM AIX Machines on page 281](#)
    - [Removing a Previous Version of StorNext](#)
    - [Running the Client Setup Wizard on page 283](#)
    - [Restoring a Previous Client Configuration on page 284](#)


## Installing the Client on Linux Machines

To run the StorNext client software on Red Hat Linux or SUSE Linux Enterprise, first install the client software package, and then configure the client.

---

 **Caution:** Before installing the StorNext client software, you must install the kernel source code. You can install the kernel source code by using the installation disks for your operating system.

---

 **Note:** Client software is extracted by default to the directory `/tmp/stornext`.

## Ensure You Have Enough Space to Extract the Installation Software

The Linux/Unix StorNext client installation software files are extracted into the `/tmp/stornext` directory by

default. The client software, when extracted, requires approximately 40 MB of space. Make sure there is enough free space in the default directory to extract the files. If there is not sufficient space, you may need to specify an alternative directory with the `-d` option.

To specify an alternate location before beginning the installation, execute the following command:

```
# ./<archive name> -d <dir>
```

where `<archive name>` is the name of the client software file you downloaded, and `<dir>` is the directory into which you want to extract the installation files.

## Other Requirements


The following requirements must be met before installing the StorNext client:

- The MDC does not have SELinux enabled.
- Quantum recommends that system clocks are synchronized using NTP for easier debugging, particularly in an HA environment.
- The following packages must be installed:
  - gcc
  - make
  - kernel-source (for systems running SUSE Linux)
  - kernel-devel (for systems running RedHat Linux)


The version of the kernel-source or kernel-devel package must correspond to the version of the booted kernel. In addition, the system must have basic utilities installed such as perl, bash, grep, the Bourne shell, etc. as well as basic libraries. In general, StorNext will not install on a stripped-down installation of Linux.

## Begin the Installation

---

 **Caution:** Before installing the StorNext client software, you must install the kernel source code. You can install the kernel source code by using the installation disks for your operating system.

---

 **Note:** Client software is extracted by default to the directory `/tmp/stornext`.

1. Log on to the client system as root.
2. Change to the directory where the client software archive file you downloaded is located.
3. Run the client software binary. At the command prompt, type:

```
./<archive name>
```

where `<archive name>` is the name of the client software archive file you downloaded. For example, the file for


**RH5 is snfs\_client\_RedHat50AS\_26x86\_64.bin.**

4. When you are presented with the StorNext End User License Agreement (EULA), press **Enter** to read the EULA. After reading the EULA enter `y` to accept the EULA, or enter any other key to decline. After you accept the EULA, the client software is extracted to `/tmp/stornext` (or to another directory if you specified an alternate location).

If the directory into which the files are being extracted does not have enough space, the installation will fail and you will receive an error message. You will need to identify an alternate location into which to extract the files by entering the following command:

```
# ./<archive name> -d <dir>
```

where `<archive name>` is the name of the client software archive file you downloaded, and `<dir>` is the directory into which you want to extract the installation files.

 **Caution:** If the directory into which the files are being extracted already exists, you will receive a warning message. If this occurs, either remove the directory and try the installation again, or change the installation directory by entering the following command:

```
# ./<archive name> -d <dir>
```

where `<archive name>` is the name of the client software archive file you downloaded, and `<dir>` is the directory into which you want to extract the installation files.

Quantum recommends you force using an existing directory (by using the `-f` option) because this could lead to additional problems installing the `.rpm` files.

5. Change directories to `/tmp/stornext` (or to the alternate directory if you specified one in the step above).
6. List the packages extracted from the software archive file. At the command prompt, execute the following command:

```
ls -l
```

7. Install the files whose names end in `.rpm`. At the command prompt, execute the following command:

```
yum install *.rpm
```

This command extracts the contents of both files into the directory `/tmp/stornext` (or an alternate

directory if you specified one in Step 5.)

8. Configure the boot order for system services.

---

**i Note:** Skip this step for Red Hat Enterprise Linux 7.x and SuSE Linux Enterprise Server 12.x clients. The boot order is determined by systemd.unit(5) unit file dependencies.

Perform one of the following:

- For **Red Hat Linux**, at the command prompt, type:

```
chkconfig --level 3456 cvfs on
```

- For **SUSE Linux Enterprise**, at the command prompt, type:

```
chkconfig -s raw 235  
chkconfig -s cvfs 345
```

9. Edit the `/usr/cvfs/config/fsnameservers` text file to contain the IP address of the MDC to which the client will connect.  
The `fsnameservers` file on the client must be exactly the same as on the MDC. If the `fsnameservers` file does not exist, use a text editor to create it.

---

**i Note:** The `fsnameservers` file needs to contain only the name of the MDC. For example, if the MDC is named `snsrver`, then the `fsnameservers` file should contain a single line with a single word: `snsrver`

10. Create a mount point for the file system. At the command prompt, type:

```
mkdir -p <mount point>  
chmod 777 <mount point>
```

where `<mount point>` is the directory path where you want the file system to be mounted. For example: `/stornext/snfs1`

11. Configure the file system to automatically mount after reboot. To do this, edit the `/etc/fstab` file so that it contains the following line

```
<file system> <mount point> cvfs verbose=yes 0 0
```

where `<file system>` is the name of the StorNext file system and `<mount point>` is the directory path created in Step 10.

## 12. Reboot the client system.

After reboot, the StorNext file system is mounted at the mount point you specified. To manually mount a file system, at the command prompt, type:

```
mount -t cvfs <file system> <mount point>
```

where **<file system>** is the name of the StorNext file system and **<mount point>** is the directory path where you want the file system to be mounted.

## Installing the Client on Sun Solaris Machines

To run the StorNext client software on Sun Solaris, first install the client software package, and then configure the client.

## Verify You Have Enough Space to Extract the Installation Software

The Linux/Unix StorNext client installation software files are extracted into the `/tmp/stornext` directory by default. The client software, when extracted, requires approximately 40 MB of space. Verify there is enough free space in the default directory to extract the files. If there is not sufficient space, you may need to specify an alternative directory with the `-d` option.

To specify an alternate location before beginning the installation, execute the following command:

```
# ./<archive name> -d <dir>
```

where **<archive name>** is the name of the client software file you downloaded, and **<dir>** is the directory into which you want to extract the installation files.

## Begin the Installation

1. Log on to the client system as **root**.
2. Change to the directory where the client software archive file you downloaded is located.
3. Run the client software binary. At the command prompt, type: `./<archive name>`

where **<archive name>** is the name of the software archive file you downloaded from the MDC. For example, the file for RH5 is `issn_dsm_linuxRedHat50AS_x86_64_client.bin`.

4. When you are presented with the StorNext End User License Agreement (EULA), press **ENTER** to read the EULA. After reading the EULA enter **y** (lower case only) to accept the EULA, or enter any other key to decline. After you accept the EULA, the client software is extracted to `/tmp/stornext` (or to another directory if you specified an alternate location).

---

**i Note:** If the directory into which the files are being extracted does not have enough space, the installation will fail and you will receive an error message. You will need to identify an alternate location into which to extract the files by executing the following command:

```
# ./<archive name> -d <dir>
```

where *<archive name>* is the name of the client software archive file you downloaded, and *<dir>* is the directory into which you want to extract the installation files.

---

**! Caution:** If the directory into which the files are being extracted already exists, you will receive a warning message. If this occurs, either remove the directory and try the installation again, or change the installation directory by executing the following command:

```
# ./<archive name> -d <dir>
```

where *<archive name>* is the name of the client software archive file you downloaded, and *<dir>* is the directory into which you want to extract the installation files.

It is NOT recommended that you force using an existing directory (by using the `-f` option) because this could lead to additional problems installing the `.rpm` files.

5. Change directories to `/tmp/stornext` .
6. Install the client software package. At the command prompt, type: `pkgadd -.`
7. Type **1** to select the **ADICsnfs** package.
8. Type **y** to confirm installation of the **ADICsnfs** package. When installation is complete, type **q** to quit the installation program.
9. Edit the `/usr/` text file to contain the IP address of the MDC to which the client will connect. The `fsnameservers` file on the client must be exactly the same as on the MDC. If the `fsnameservers` file does not exist, use a text editor to create it.

---

**i Note:** The `fsnameservers` file needs to contain only the IP address of the MDC. For example, if the IP address for the MDC is `192.168.143.1`, then the `fsnameservers` file should contain a single line with `192.168.143.1`

10. Create a mount point for the file system. At the command prompt, type:

```
mkdir -p <mount point>  
chmod 777 <mount point>
```

where *<mount point>* is the directory path where you want the file system to be mounted. For example: `/stornext/snfs1`

11. Configure the file system to automatically mount after reboot. To do this, edit the `/etc/vfstab` file so that it contains the following line: `<file system>-<mount point>cvfs 0 auto rw`

where `<file system>` is the name of the StorNext file system and `<mount point>` is the directory path created in step 10.

12. Reboot the client system. After reboot, the StorNext file system is mounted at the mount point you specified.

---

**Note:** To manually mount a file system, at the command prompt, type:

```
mount -F cvfs <file system> <mount point>
```

where `<file system>` is the name of the StorNext file system and `<mount point>` is the directory path where you want the file system to be mounted.

## Installing the Client on HP-UX Machines

To run the StorNext client software on HP-UX, first install the client software package, and then configure the client.

## Verify You Have Enough Space to Extract the Installation Software

The Linux/Unix StorNext client installation software files are extracted into the `/tmp/stornext` directory by default. The client software, when extracted, requires approximately 40 MB of space. Make sure there is enough free space in the default directory to extract the files. If there is not sufficient space, you may need to specify an alternative directory with the `-d` option.

To specify an alternate location before beginning the installation, execute the following command:

```
# ./<archive name> -d <dir>
```

where `<archive name>` is the name of the client software file you downloaded, and `<dir>` is the directory into which you want to extract the installation files.

## Begin the Installation

1. Log on to the client system as **root**.
2. Change to the directory where the client software archive file you downloaded is located.
3. Run the client software binary. At the command prompt, type: `./<archive name>`

where `<archive name>` is the name of the software archive file you downloaded from the MDC. For example, the file for RH5 is `sn_dsm_linuxRedHat50AS_x86_64_client.bin`.



4. When you are presented with the StorNext End User License Agreement (EULA), press **ENTER** to read the EULA. After reading the EULA enter **y** (lower case only) to accept the EULA, or enter any other key to decline. After you accept the EULA, the client software is extracted to `/tmp/stornext`.

---

**i Note:** If the directory into which the files are being extracted does not have enough space, the installation will fail and you will receive an error message. You will need to identify an alternate location into which to extract the files by executing the following command:

```
# ./<archive name> -d <dir>
```

where *<archive name>* is the name of the client software archive file you downloaded, and *<dir>* is the directory into which you want to extract the installation files.

---

**! Caution:** If the directory into which the files are being extracted already exists, you will receive a warning message. If this occurs, either remove the directory and try the installation again, or change the installation directory by entering the following command:

```
# ./<archive name> -d <dir>
```

where *<archive name>* is the name of the client software archive file you downloaded, and *<dir>* is the directory into which you want to extract the installation files.

It is NOT recommended that you force using an existing directory (by using the `-f` option) because this could lead to additional problems installing the `.rpm` files.

5. Change directories to `/tmp/stornext` (or to another directory if you specified an alternate location)
6. List the packages extracted from the software archive file. At the command prompt, type: `ls -`

Identify the correct package to install. The correct package begins with `snfs-client` and ends with the `.depot` file name extension.

7. Install the client software package. At the command prompt, type:  
`swinstall -s<package path and name>-x mount_all_filesystems=false \*`

where *<package path and name>* is the full path and name of the client software package you identified in step 6.

8. Edit the `/usr/cvfs/config/fsnameservers` text file to contain the IP address of the MDC to which the client will connect. The `fsnameservers` file on the client must be exactly the same as on the MDC. If the `fsnameservers` file does not exist, use a text editor to create it.

The `fsnameservers` file needs to contain only the name of the MDC. For example, if the MDC is named `snsnserver`, then the `fsnameservers` file should contain a single line with a single word: `snsnserver`

9. Create a mount point for the file system. At the command prompt, type:

```
mkdir -p <mount point>  
chmod 777 <mount point>
```

where *<mount point>* is the directory path where you want the file system to be mounted. For example: `/stornext/snfs1`

10. Configure the file system to automatically mount after reboot. To do this, edit the `/etc/fstab` file so that it contains the following line:

```
<mount point> <mount point>cvfs rw,<file system>0 0
```

where *<mount point>* is the directory path created in step 9, and *<file system>* is the name of the StorNext file system.

11. Reboot the client system. After reboot, the StorNext file system is mounted at the mount point you specified.

---

**i Note:** To manually mount a file system, at the command prompt, type:

```
mount -F cvfs <file system> <mount point>
```

where *<file system>* is the name of the StorNext file system and *<mount point>* is the directory path where you want the file system to be mounted.

## Installing the Client on IBM AIX Machines

To run the StorNext client software on IBM AIX, first install the client software package, and then configure the client.

## Verify You Have Enough Space to Extract the Installation Software

The Linux/Unix StorNext client installation software files are extracted into the `/tmp/stornext` directory by default. The client software, when extracted, requires approximately 40 MB of space. Make sure there is enough free space in the default directory to extract the files. If there is not sufficient space, you may need to specify an alternative directory with the `-d` option.

To specify an alternate location before beginning the installation, execute the following command

```
# ./<archive name> -d <dir>
```

where *<archive name>* is the name of the client software file you downloaded, and *<dir>* is the directory into which you want to extract the installation files.

## Begin the Installation

1. Log on to the client system as **root**.
2. Change to the directory where the client software archive file you downloaded is located.
3. Run the client software binary. At the command prompt, type: `./<archive name>`

where *<archive name>* is the name of the software archive file you downloaded from the MDC. For example, the file for RH5 is `sn_dsm_linuxRedHat50AS_x86_64_client.bin`.

4. When you are presented with the StorNext End User License Agreement (EULA), press **ENTER** to read the EULA. After reading the EULA enter **y** (lower case only) to accept the EULA, or enter any other key to decline. After you accept the EULA, the client software is extracted to `/tmp/stornext` (or to another directory if you specified an alternate location).

---

**i Note:** If the directory into which the files are being extracted does not have enough space, the installation will fail and you will receive an error message. You will need to identify an alternate location into which to extract the files by executing the following command:

```
# ./<archive name> -d <dir>
```

where *<archive name>* is the name of the client software archive file you downloaded, and *<dir>* is the directory into which you want to extract the installation files.

---

**! Caution:** If the directory into which the files are being extracted already exists, you will receive a warning message. If this occurs, either remove the directory and try the installation again, or change the installation directory by executing the following command:

```
# ./<archive name> -d <dir>
```

where *<archive name>* is the name of the client software archive file you downloaded, and *<dir>* is the directory into which you want to extract the installation files.

It is NOT recommended that you force using an existing directory (by using the `-f` option) because this could lead to additional problems installing the `.rpm` files.

5. Change directories to `/tmp/stornext`.
6. List the packages extracted from the software archive file. At the command prompt, type: `ls -l`  
  
Identify the correct package to install. The correct package begins with `snfs` and ends with the `.bff` file name extension.
7. Install the client software package. At the command prompt, type: `installp -ac -<package name>all`.  
  
where *<package name>* is the name of the client software package you identified in step 6.
8. Edit the `/usr/cvfs/config/fsnameservers` text file to contain the IP address of the MDC to which the client will connect. The `fsnameservers` file on the client must be exactly the same as on the MDC. If the `fsnameservers` file does not exist, use a text editor to create it.

---

**i Note:** The `fsnameservers` file needs to contain only the name of the MDC. For example, if the MDC is named `snsnserver`, then the `fsnameservers` file should contain a single line with a single word: `snsnserver`

9. Create a mount point for the file system. At the command prompt, type:

```
mkdir -p <mount point>  
chmod 777 <mount point>
```

where *<mount point>* is the directory path where you want the file system to be mounted. For example: /stornext/snfs1

10. Configure the file system to automatically mount. At the command prompt, type:

```
crfs -v cvfs -<file system>-a verbose=yes -a<mount point>
```

where *<file system>* is the name of the StorNext file system and *<mount point>* is the directory path created in step 9. The StorNext file system is mounted at the mount point you specified.

---

**i Note:** To manually mount a file system, at the command prompt, type:

```
mount <mount point>
```

where *<mount point>* is the directory path where you want the file system to be mounted

## Running the Client Setup Wizard

Before you begin the installation, you must first log on as an Administrator

---

**i Note:** If you are installing on Windows Vista, answer **Yes** to any messages asking if you want to run the installation process with administrative privileges. You must log on as an Administrator to install StorNext.

## Install the StorNext Client on Windows

1. Download the client software from the MDC as described in [Downloading Client Software on page 272](#).

**Exception:** If you are running SNFS-only on a Windows machine, you cannot download the client software from the MDC. Instead, you will need to obtain the client software from the StorNext installation DVD. Copy one of the following files from the DVD and place it in an accessible location on the client system, then proceed to the next step.

- **fs\_only/snfs\_client\_Windows\_i386.zip** (for 32-bit systems)
  - **fs\_only/snfs\_client\_Windows\_x86\_64.zip** (for 64-bit systems)
2. If necessary, unzip the client software file you downloaded.
  3. Open the unzipped folder and double-click the client software installer file. This file is named **SnfsSetup32.exe** (32-bit systems) or **SnfsSetup64.exe** (64-bit systems). The **StorNext**

**Installation** window appears.

4. Click **Install StorNext** to begin installation. The **StorNext File System License Agreement** window appears.
5. Read the license. If you agree to the terms, select **I accept the terms in the License Agreement** and click **Next** to continue the install. The StorNext setup wizard launches.
6. Click **Next** to continue. The **Customer Information** window appears.
7. Type your name and the name of your company or organization in the boxes, and then click **Next** to continue. The **Choose Setup Type** window appears.
8. Click one of the following installation options. Since this is a client-only install, you should select **Client Only** or **Custom**.
  - **Client Only**: Installs the StorNext client software and help files in the default location (**C:\Program Files\StorNext**).
  - **Custom**: Lets you choose which components to install and specify an installation location. When ready, click **Next**.
  - **Complete**: All installable components will be installed in the default location (**C:\Program Files\StorNext**). Since this is a client-only install, only the client software and help files will be installed (the same as if you selected the **Client Only** option). The **Ready to Install** window appears.
9. Click **Install** to continue. Wait while the setup wizard installs StorNext. When installation is complete, the **Finish** window appears.

---

**Note:** It may take several minutes for the installation to finish.

10. Click **Finish** to close the setup wizard. The **Installation Complete** dialog box opens.
11. Click **OK** in the **Installation Complete** dialog box. The **StorNext Installation** page displays again with choices to **Add/Remove Components**, **Remove StorNext**, or **Configure StorNext**.
12. The installation is complete. Click one of the following options:
  - **Add/Remove Components**: The components you selected during the install are already added, but if you changed your mind and want to select different components, you can do so now.
  - **Remove StorNext**: Removes StorNext file system.
  - **Configure StorNext**: Configures the StorNext client.
  - **Exit**: Exits without making any further changes at this time (you can always perform configuration later).

You can now configure StorNext File System as described in the *StorNext Installation Guide*.

## Restoring a Previous Client Configuration

When upgrading from a StorNext 3.0 or later release, client configurations are fully maintained and do not need to be imported or exported. Client configuration files are automatically saved when you remove StorNext, or can be manually saved for backup purposes. If you want to restore a saved a client

configuration file, you can import it using the StorNext Installation tool. This configures StorNext using the settings from a previous installation.

## Use the StorNext Installation Tool

1. Double-click the file **SnfsSetup32.exe** (32-bit systems) or **SnfsSetup64.exe** (64-bit systems). The **StorNext Installation** window appears.
2. Click **Configure StorNext**. The **StorNext Configuration** window appears.
3. Click **Import/Export Client Settings**. The **Import/Export Client Configuration** window appears.
4. Under **Import**, click **Browse**. Locate the client configuration (\*.reg) file to import, and then click **Open**. Client configuration files saved during removal of a previous version of StorNext are located in the following directory by default: **C:\Program Files\StorNext\config\**
  - Note:** Systems prior to StorNext 3.0 used the default directory **C:\SNFS\config\**.
5. Click an option for handling current configuration settings:
  - **Merge with current configuration:** The imported configuration is merged with the current configuration settings. Imported mount information is added to, or overwrites, existing information.
  - **Replace current configuration:** The current mount information is completely removed and replaced by the imported configuration.
6. Click **Import**, and then click **Yes** to confirm. A message appears informing you the configuration settings were successfully added to the registry.
7. Click **OK**, and then click **Quit** to close the **Import/Export Client Configuration** window.

---

## Firmware Upgrade

For instructions on upgrading your firmware, see [Upgrade the System \(Upgrade Firmware\)](#) on the *Appliance InfoHub Documentation Center* ([www.quantum.com/ApplianceInfoHub](http://www.quantum.com/ApplianceInfoHub)).

---

## Appliance Release Notes

Refer to the respective Release Notes document for important information you should know about your system.

- [Xcellis Foundation](#)
- [aiWARE for Xcellis](#)
- [Xcellis Workflow Extender](#)

- [Xcellis Workflow Director](#)
- [Artico](#)
- [Pro Foundation](#)
- [G300](#)
- [M660](#)
- [M440](#)

## Upgrade Firmware

1. Choose **Firmware Upgrade** from the **Tools** menu. The **Firmware Upgrade** page appears. Uploaded files from previous upgrades will be shown. A new system will not contain upgrade files.
2. Perform one of the following:
  - a. Select the **Auto Upload** option to upload the file immediately after you select it.
  - b. Do not select **Auto Upload**.
3. Click **Choose File**, and then navigate to the directory where the file resides. Firmware files are identifiable by a `.fw` extension. If you selected the **Auto Upload** option, the file is immediately uploaded. Proceed to **Step 5**.
4. 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.

---

**i Note:** Files are automatically validated after you click **Apply** (Step 5), but you won't receive a message telling you the file is valid.

5. 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.

---

**i Note:** After the upgrade completes, both nodes are rebooted. It could be a long time before you are able to log back in, so plan upgrade times accordingly.

## Delete Uploaded Files

Follow these steps to delete uploaded files you no longer need:

1. 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.
2. When a confirmation appears, click **Yes** to proceed or **No** to abort.
3. When a message informs you that the file was successfully deleted, click **OK**.

# Dynamic Application Environment

The **Dynamic Application Environment** page displays the current status of your virtual machine environments. To access the **Dynamic Application Environment** page, on the **Tools** menu, click **Dynamic Application Environment**. The **Tools > Dynamic Application Environment** page appears.

**i Note:** The **Dynamic Application Environment** option is **only** available on Xcellis systems, and is a licensed StorNext feature.

**i Note:** If your virtual machine environment is not displayed on the **Dynamic Application Environment** page, you may have to complete the configuration. See [Dynamic Application Environment Overview](#) in the Xcellis Workflow Director Documentation Center.

The table below provides the information displayed on the **Tools > Dynamic Application Environment** page:

Parameter	Description
<b>Virtual Machine Name</b>	Displays the name of the virtual machine.
<b>VNC Port</b>	Displays the Virtual Network Computing (VNC) port number.
<b>State</b>	Displays the current status of the virtual machine. The possible states are: <ul style="list-style-type: none"><li>• <b>Powered On</b></li><li>• <b>Powered Off</b></li><li>• <b>Suspended</b></li></ul>
<b>Change DAE Password...</b>	Click to change a Dynamic Application Environment password.
<b>Power On</b>	Click to power on a selected virtual machine.
<b>Shutdown</b>	Click to shut down a selected virtual machine.
<b>Power Off</b>	Click to power off a selected virtual machine.
<b>Suspend</b>	Click to suspend a selected virtual machine.
<b>Refresh</b>	Click to refresh the data on this page.

**i Note:** You may only have a set number of virtual machines (dictated by the system maximum number) powered on or suspended at a time. If you attempt to power on or suspend more than the system maximum number of virtual machines, then you will receive an error message prompting you to shut down another virtual machine that is powered on or suspended.



## Change a Dynamic Application Environment Password

**i Note:** This task allows you to update the default password used to [access the virtual machine console](#). The user account must have administrative privileges to update the password. The default username is **daadmin** and the default password is **password**.

1. In the **Virtual Machine Name** column, click the virtual machine.
2. Click **Change DAE Password....** The **Tools > Dynamic Application Environment > Change DAE Password** page appears.
  - a. In the **Current Password** field, enter the existing password for the virtual machine.
  - b. In the **New Password** field, enter a new password for the virtual machine.
  - c. In the **Confirm Password** field, enter the same password you entered in the **New Password** field.
3. Click **Apply** to save your changes and update the password, or click **Cancel** to return to the previous page.

## Power on a Virtual Machine

1. In the **Virtual Machine Name** column, click the virtual machine.
2. Click **Power On**. You are prompted to verify if you want to power on the selected virtual machine.
3. Click **Yes** to continue, or click **No** to cancel the operation and return to the **Dynamic Application Environment** page. A confirmation dialog appears.
4. **(Optional)** Click **Refresh** to update the **Dynamic Application Environment** page and display the updated virtual machine status in the **State** column.

## Shut down a Virtual Machine


**i Note:** This operation also gracefully shuts down the guest operating system.

1. In the **Virtual Machine Name** column, click the virtual machine.
2. Click **Shutdown**. You are prompted to verify if you want to gracefully shut down the selected virtual machine.
3. Click **Yes** to continue, or click **No** to cancel the operation and return to the **Dynamic Application Environment** page. A confirmation dialog appears.
4. **(Optional)** Click **Refresh** to update the **Dynamic Application Environment** page and display the updated virtual machine status in the **State** column.

## Power off a Virtual Machine

**! Caution:** This operation immediately powers off the virtual machine, and does not shut down the guest

---

 operating system. Quantum recommends you gracefully shut down your guest operating system and virtual machine before you power off the selected virtual machine (refer to the procedure on how to [Shut down a Virtual Machine on the previous page](#)). Failure to do so may result in data corruption.

1. In the **Virtual Machine Name** column, click the virtual machine.
2. Click **Power Off**. You are prompted to verify if you want to power off the selected virtual machine.
3. Click **Yes** to continue, or click **No** to cancel the operation and return to the **Dynamic Application Environment** page. A confirmation dialog appears.
4. **(Optional)** Click **Refresh** to update the **Dynamic Application Environment** page and display the updated virtual machine status in the **State** column.

## Suspend a Virtual Machine

1. In the **Virtual Machine Name** column, click the virtual machine.
2. Click **Suspend**. You are prompted to verify if you want to suspend the selected virtual machine.
3. Click **Yes** to continue, or click **No** to cancel the operation and return to the **Dynamic Application Environment** page. A confirmation dialog appears.
4. **(Optional)** Click **Refresh** to update the **Dynamic Application Environment** page and display the updated virtual machine status in the **State** column.

## Resume a Virtual Machine

1. In the **Virtual Machine Name** column, click the virtual machine.
2. Click **Resume**. You are prompted to verify if you want to resume the selected virtual machine.
3. Click **Yes** to continue, or click **No** to cancel the operation and return to the **Dynamic Application Environment** page. A confirmation dialog appears.
4. **(Optional)** Click **Refresh** to update the **Dynamic Application Environment** page and display the updated virtual machine status in the **State** column.

## Refresh the Virtual Machine State

Click **Refresh** to manually update the **Dynamic Application Environment** page and display the updated virtual machine status.

## Access the Virtual Machine Console

You can access a selected virtual machine via a Guacamole remote desktop console using your browser. For additional information about Guacamole, see <http://guacamole.incubator.apache.org/>.

1. In the **Virtual Machine Name** column, click the option button next to the virtual machine.
2. Below the table, in the text labeled "To access the virtual machine console **click here**", click the text "**click**

here". The Guacamole remote desktop console browser window appears and you are prompted for a username and password. The default username is **daadmin** and the default password is **password**.

---

**i** **Note:** To update the default password, see [Change a Dynamic Application Environment Password on page 288](#). The user account must have administrative privileges to update the password.

---

**i** **Note:** To display an on-screen keyboard for entering special key sequences such as **Ctrl+Alt+Delete**, press **Ctrl+Alt+Shift** to display the Guacamole options.

---

## System Control

Choose **System Control** from the **Tools** menu to access the **System Control** page. 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.

### Start or Stop the 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.

### Start or Stop the 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.

## Refresh the 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.

## Specify the 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**. For configuration details, see [HTTPS Configuration on page 300](#). If you are working on a Lattus system which does not have an existing SSL certificate, see [Work on a Lattus System and a StorNext Metadata Controller on page 300](#). The table below provides the information displayed for each certificate, on the **Tools > Lattus Certificates** page:

Heading	Description	Examples
<b>Public Certificate File Name</b>	<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"><li>• 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.</li><li>• There will always be a public file name, and a private if you use this feature to generate your certificates.</li></ul>	<p>accounts.mycompany.pem accounts.mycompany.der</p>
<b>Common Name</b>	<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>

Heading	Description	Examples
<b>Organization</b>	The legal name of your organization. This should not be abbreviated and should include suffixes such as Inc, Corp, or LLC.	Mycompany Corp
<b>Organizational Unit</b>	The division of your organization handling the certificate.	Information Technology IT Department
<b>Valid From</b>	The date the certificate is valid <b>from</b> , in the form of yyyy-mm-dd hh:mm:ss time zone.	1998-08-22 11:41:51 MST
<b>Valid To</b>	The date the certificate is valid <b>to</b> , in the form of yyyy-mm-dd hh:mm:ss time zone.	1998-08-22 11:41:51 MST

## 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:

Operating System	Default Repository Referenced by Storage Manager
<b>Debian</b>	<code>/etc/ssl/certs/ca-certificates.crt</code>
<b>Red Hat</b>	<code>/etc/pki/tls/certs/ca-bundle.crt</code> or <code>/etc/ssl/certs/ca-bundle.crt</code>
<b>SUSE</b>	<code>/etc/ssl/certs/</code>

## 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 Number	Description
1	<p>Use <code>/usr/cvfs/config/ssl</code> as your default certificate repository; set the <code>FS_OBJSTORAGE_CAPATH=/usr/cvfs/config/ssl</code> in the <code>/usr/adic/TSM/config/fs_sysparm_override</code> file, then copy all the root certificates from <code>/etc/ssl/certs</code> to <code>/usr/cvfs/config/ssl</code>.</p> <hr/> <p><b>Note:</b> Be sure to execute <code>c_rehash</code> on the directory that will be used as your default certificate repository afterward.</p>
2	<p>Do not use <code>/usr/cvfs/config/ssl</code>; instead use the default root certificate repository by copying your certificates from <code>/usr/cvfs/config/ssl</code> to <code>/etc/ssl/certs</code>.</p> <hr/> <p><b>Note:</b> Be sure to execute <code>c_rehash</code> on the directory that will be used as your default certificate repository afterward.</p>

## 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.

## Create a 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 a different filename extension or format, as this is the only format currently supported with the program that uses the certificate.

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 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
<b>File Name (.pem extension)</b>	<p>The Privacy Enhanced Mail (PEM) filename and its respective filename extension of the self-signed Lattus certificate.</p> <p><b>i Note:</b> Adding a certificate with the same name generates an error, instructing you to delete the certificate with that name first.</p>	<code>accounts.mycompany.pem</code>
<b>Password (at least 4 characters)</b>	<p>The <b>Password</b> input is an <b>optional</b> field. If a <b>Password</b> is entered, the input mimics the OpenSSL command password requirements as follows:</p> <ul style="list-style-type: none"> <li>• The <b>Password</b> input, and the <b>Confirm Password</b> input must match.</li> <li>• The <b>Password</b> input must be at least 4 characters, and can be all empty spaces or contain spaces.</li> </ul>	<code>mypassword1234</code>
<b>Confirm Password</b>	<p>See the requirements for the <b>Password</b> input.</p>	<code>mypassword1234</code>
<b>Expiration Date</b>	<p>The <b>Expiration Date</b> 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 <b>Years, Months, and Days</b>.</p> <p><b>i Note:</b> 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.</p>	<code>5 Years</code>
<b>Common Name</b>	<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>	<code>*.mycompany.com</code> <code>controller.mycompany.com</code>

Text Box	Description	Examples
<b>Organizational Unit</b>	The <b>division</b> of your organization handling the certificate.	Information Technology IT Department
<b>Organization</b>	The <b>legal name</b> of your organization. This should not be abbreviated and should include suffixes such as Inc, Corp, or LLC.	Mycompany Corp
<b>Location</b>	The <b>city</b> where your organization is located.	Englewood
<b>State</b>	The <b>state</b> where your organization is located. This should not be abbreviated.	Colorado
<b>Country</b>	The <b>two-letter ISO code</b> for the country where your organization is located.	US
<b>Subject Alternative Name</b>	The <b>Subject Alternative Name</b> is an <b>optional</b> 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.
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. The table below describes the various fields on the **Tools > Lattus Certificates > View** page:

Name	Description	Examples
<b>Public Certificate File Name</b>	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



Name	Description	Examples
<b>Private Key File Name</b>	<p>The filename of the private key in the Lattus certificate.</p> <hr/> <p><b>i Note: Note:</b> 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.</p>	<p>/usr/adic/gui/.ssl/myCert.pem</p>
<b>Issuer</b>	<p>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.</p>	<p>CN=mycert.mycompany.com, OU=StorNext Software, O=Mycompany Corp, L=Englewood, ST=CO, C=US</p>
<b>Common Name</b>	<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>mycert.mycompany.com</p>
<b>Organizational Unit</b>	<p>The <b>division</b> of your organization handling the certificate.</p>	<p>Information Technology IT Department</p>
<b>Organization</b>	<p>The <b>legal name</b> of your organization. This should not be abbreviated and should include suffixes such as Inc, Corp, or LLC.</p>	<p>Mycompany Corp</p>
<b>Serial Number</b>	<p>The serial number of the selected certificate.</p>	<p>F2:D8:5A:FA:C9:E6:11:CF</p>
<b>Valid From</b>	<p>The date the certificate is valid <b>from</b>, in the form of yyyy-mm-dd hh:mm:ss time zone.</p>	<p>2013-01-31 14:33:07 MST</p>
<b>Valid To</b>	<p>The date the certificate is valid <b>to</b>, in the form of yyyy-mm-dd hh:mm:ss time zone.</p>	<p>2018-01-30 14:33:07 MST</p>
<b>Location</b>	<p>The <b>city</b> where your organization is located.</p>	<p>Englewood</p>

Name	Description	Examples
<b>State</b>	The <b>state</b> where your organization is located. This should not be abbreviated.	Colorado
<b>Country</b>	The <b>two-letter ISO code</b> for the country where your organization is located.	US
<b>Signature Algorithm</b>	The algorithm used to create the signature of the certificate.	SHA1withRSA
<b>Signature Algorithm OID</b>	The object identifier (OID) identifies the type of signature algorithm used by the certificate.	1.2.840.113549.1.1.5
<b>Version</b>	The version number of the certificate.	V3
<b>Subject Alternative Name</b>	The <b>Subject Alternative Name</b> 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 the next page](#) 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.
- 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.
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, that is, anything that is not a .pem file format, then 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:
    - **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.
    - **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:
    - **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.
  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 a Certificate File

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 [Create a Self-signed Certificate on page 293](#) 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.
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 a Certificate File

### Notes and Considerations

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

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 yyyyMMdHHmmss.

1. On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears.
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 the File List

### 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.
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
objectstorage.openssl.binary	If the objectstorage.openssl.binary file is not installed, the <b>New...</b> button is disabled, and you will receive an error message.
objectstorage.c_rehash.binary	If the objectstorage.c_rehash.binary file is not installed, both the <b>New...</b> and <b>Import...</b> 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*.

## Work on a Lattus System and a StorNext Metadata Controller

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 [Basic Secure Sockets Layer \(SSL\) Guidelines on page 497](#).

See [Basic Secure Sockets Layer \(SSL\) Guidelines on page 497](#), as it outlines some standard information about using private and public certificates.

## Create a Private and Public SSL Certificate for use on a Lattus System and a StorNext MDC

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 [Create a Self-signed Certificate on page 293](#) section describes the various text boxes on the **Tools > Lattus Certificates > New** page:

---

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

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

## Update an Expired CA Root Certificate

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
```

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

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

---

## 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. To access the **Tools > File and Directory Actions** page, on the **Tools** menu, click **File and Directory Actions**. The following information is displayed for the available files:

Parameter	Description
<b>Name</b>	The name of the file.
<b>Owner</b>	The file owner.



Parameter	Description
<b>Group</b>	The group to which the file belongs.
<b>Size</b>	The size (in bytes) of the file.
<b>Last Modified</b>	The date and time when the file was last modified.

## Perform File and Directory Actions

At the top of the page is a drop-down list of **Available Actions** you can perform for files and directories. Select one of these options and follow the procedures below for the specific action.

For databases which contain more than one million files, Quantum recommends using the **Directory Filter** and **Filter** functionality for increased performance.

**i Note:** For any **File and Directory Action** function where you are prompted to input text into the **Directory Filter** field, **do not** include the **Traversed Path** (in other words, the mount point) within the text field. For example, if the entire directory path is `/stornext/snfs1/directory1`, do not input `/stornext/snfs1/` into the **Directory Filter** field.

**i Note:** For operations which include a **Filter** function, you may use wild-card syntax [an asterisk (\*) symbol, or the percent (%) symbol]. For example, `doc*`, `file000*`, etc.

## Store Files

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

1. To select the files to store, click the **Browse** button in the **Selected Files** field.
2. On the Browse page, do the following:
  - a. The managed relation points will be displayed. Click on a relation point. This displays any subdirectories and files under the relation point.
  - b. If files you want to store are visible, click **All** to select all visible files, or click the **Select** box next to each file you want to store.
  - c. If the files are in a subdirectory, click on the directory to view any subdirectories and files under the directory. Repeat step a.
  - d. Once all files you want store are selected, click **Continue** to return the selected files back to the **Tools > File and Directory Actions** page, or click **Cancel** to return to the **Tools > File and Directory Actions** page without saving selections.
3. To store the selected file according to policy, at the **Store Parameters** field, select **By Policy**.
4. Do one of the following:

- Click **Apply** to start a job to store the selected files. When the confirmation message appears, click **OK**.

---

**i Note:** The Job ID in the message and visit the Reports > Jobs page to see the result of the action.

- Click **Reset** to clear all selected files.
  - Click **Cancel** to exit the **Tools > File and Directory Actions** page without making changes.
5. To store the selected file according to custom parameters, at the **Store Parameters** field, select **Custom**.
  6. 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 **files**
    - **Minimum File Size:** Specify the minimum file size
    - **Media Type:** Specify the tape drive media type or sdisk
    - **Media Format:** Specify the files to store in the selected media format (**ANTF** or **LTFS**) for a **Custom** store selection. The list is only enabled for the LTO media type. For non-LTO media type, including LATTUS, S3COMPAT, Q-Cloud, and SDISK, the option is disabled and ignored.
  7. Do one of the following:
    - Click **Apply** to start a job to store the selected files. When the confirmation message appears, click **OK**.
    - Click **Reset** to clear all selected files.
    - Click **Cancel** to exit the **Tools > File and Directory Actions** page without making changes.

## Change File Version

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

1. To select the file to change, click the **Browse** button in the **Selected Files** field.
2. On the Browse page, do the following:
  - a. The managed relation points will be displayed. Click on a relation point. This displays any subdirectories and files under the relation point.
  - b. If file you want to change is visible, click the **Select** box next to that file.
  - c. If the file is in a subdirectory, click on the directory to view any subdirectories and files under the directory. Repeat step a.
  - d. Once the file you want change is selected, click **Continue** to return the selected file back to the

**Tools > File and Directory Actions** page, or click **Cancel** to return to the **Tools > File and Directory Actions** page without saving selections.

3. Do one of the following:
  - Click **Apply** to start a job to change the version of the selected file. When the confirmation message appears, click **OK**. If desired, make a note of the Job ID in the message and visit the **Reports > Jobs** page to see the result of the action.
  - Click **Reset** to clear all selected files.
  - Click **Cancel** to exit the **Tools > File and Directory Actions** page without making changes.
4. Repeat steps 1 - 5 to change versions for additional files.

## Recover Files

Select this option to recover previously deleted files.

1. To select the files to recover, click the **Browse** button.
2. On the Browse page, do the following:
  - a. At the **FileFilter** field, enter one or more characters. Click **Filter** to display all files whose names contain the character(s) you entered.
  - b. When search results are displayed, click the box next to the files to be recovered, or select all displayed files by clicking the box next to the **Name** heading.
  - c. The **Deleted Between** filter restricts the file list to only files which were deleted within a certain date range. To use this filter, position the cursor in the first field, and then click the calendar icon. Select the desired date and time, and then click the blue X icon to close the calendar. Repeat the process for the ending date. Click the **Filter** button to the right of the **Deleted Between** field to apply the filtering criteria you entered.
  - d. Click **Continue** to return the selected files back to the **Tools > File and Directory Actions** page, or click **Cancel** to return to the **Tools > File and Directory Actions** page without saving selections.
3. Do one of the following:
  - Click **Apply** to start a job to recover the selected files. When the confirmation message appears, click **OK**. If desired, make a note of the Job ID in the message and visit the **Reports > Jobs** page to see the result of the action.
  - Click **Cancel** to exit the **Tools > File and Directory Actions** page without making changes.

## Recover Directories

Select this option to recover previously deleted directories.

1. To select the directories to recover, click the **Browse** button.
2. On the Browse page, do the following:

- a. At the **Directory Filter** field, enter one or more characters. Click **Filter** to display all directories whose names include the character(s) you entered.
  - b. The **Deleted Between** filter restricts the directory list to only directories which were deleted within a certain date range. To use this filter, position the cursor in the first field, and then click the calendar icon. Select the desired date and time, and then click the blue X icon to close the calendar. Repeat the process for the ending date. Click the **Filter** button to the right of the **Deleted Between** field to apply the filtering criteria you entered.
  - c. When search results are displayed, click the box next to the directories to be recovered, or click the box next to the **Name** heading to select all displayed directories.
  - d. Click **Continue** to return the selected directories back to the **Tools > File and Directory Actions** page, or click **Cancel** to return to the **Tools > File and Directory Actions** page without saving selections.
3. 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 that 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.

4. Do one of the following:
- Click **Apply** to start a job to recover the selected directories. When the confirmation message appears, click **OK**. If desired, make a note of the Job ID in the message and visit the **Reports > Jobs** page to see the result of the action.
  - Click **Cancel** to exit the **Tools > File and Directory Actions** page without making changes.

## Retrieve Files

Select this option to retrieve files.

1. To select the files to retrieve, click the **Browse** button in the **Selected Files** field.
2. On the Browse page, do the following:
  - a. The managed relation points will be displayed. Click on a relation point. This displays any subdirectories and files under the relation point.
  - b. If files you want to retrieve are visible, click **All** to select all visible files, or click the **Select** box next to each file you want to retrieve.
  - c. If the files are in a subdirectory, click on the directory to view any subdirectories and files under the directory. Repeat step b.
  - d. Once all files you want retrieve are selected, click **Continue** to return the selected files back to the **Tools > File and Directory Actions** page, or click **Cancel** to return to the **Tools > File and Directory Actions** page without saving selections.
3. (Optional) If desired, select one of the files in the selected **Files** field. This enables the **Optional Retrieve Parameters** to be set. 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.

4. Do one of the following:
  - Click **Apply** to start a job to retrieve the selected files. When the confirmation message appears, click **OK**. If desired, make a note of the Job ID in the message and visit the **Reports > Jobs** page to see the result of the action.
  - Click **Reset** to clear all selected files.
  - Click **Cancel** to exit the **Tools > File and Directory Actions** page without making changes.

## Retrieve Directory

Select this option to retrieve directories.

1. To select the directory to use as the start for a recursive retrieve, click the **Browse** button in the **Selected Files** field.
2. On the Browse page, the managed relation points are displayed. Do one of the following:
  - If the relation point is to be used as the selected directory, Click the circle next to the relation point and proceed to step 3.
  - If the desired directory is a subdirectory of a relation point, click the relation point to display any subdirectories under the relation point.
  - If the directory you want to use is not visible, continue to click on the subdirectories until the desired one is visible. Once visible, click the circle next to the directory.

3. Once the desired directory is selected, do one of the following:
  - Click **Continue** to return the selected directory back to the **Tools > File and Directory Actions** page
  - Click **Cancel** to return nothing back to the **Tools > File and Directory Actions** page without saving selections
4. Do one of the following:
  - Click **Apply** to start a job to do a recursive retrieve. When the confirmation message appears, click **OK**. If desired, make a note of the Job ID in the message and visit the **Reports > Jobs** page to see the result of the action.
  - Click **Reset** to clear all selected directory.
  - Click **Cancel** to exit the **Tools > File and Directory Actions** page without making changes.
5. Repeat steps 1 - 3 to retrieve additional directories.

## Truncate Files

Select this option to truncate files.

1. To select the files to truncate, click the **Browse** button in the **Selected Files** field.
2. On the Browse page, do the following:
  - a. The managed relation points will be displayed. Click on a relation point. This displays any subdirectories and files under the relation point.
  - b. If files you want to truncate are visible, click **All** to select all visible files, or click the **Select** box next to each file you want to truncate.
  - c. If the files are in a subdirectory, click on the directory to view any subdirectories and files under the directory. Repeat step b.
  - d. Once all files you want truncate are selected, click **Continue** to return the selected files back to the **Tools > File and Directory Actions** page, or click **Cancel** to return to the **Tools > File and Directory Actions** page without saving selections.
3. Do one of the following:
  - Click **Apply** to start a job to truncate the selected files. When the confirmation message appears, click **OK**. If desired, make a note of the Job ID in the message and visit the **Reports > Jobs** page to see the result of the action.
  - Click **Reset** to clear all selected files.
  - Click **Cancel** to exit the **Tools > File and Directory Actions** page without making changes.

## Move Files

Select this option to move files.

1. To select the files to move, click the **Browse** button in the **Selected Files** field.
2. On the **Browse** page, do the following:
  - a. The managed relation points will be displayed. Click on a relation point. This displays any subdirectories and files under the relation point.
  - b. If files you want to move are visible, click **All** to select all visible files, or click the **Select** box next to each file you want to move.
  - c. If the files are in a subdirectory, click on the directory to view any subdirectories and files under the directory. Repeat step b.
  - d. Once all files you want move are selected, click **Continue** to return the selected files back to the **Tools > File and Directory Actions** page, or click **Cancel** to return to the **Tools > File and Directory Actions** page without saving selections.
3. 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. When moving files to a specific media type, check the **Blank Media** box to designate the file(s) to be moved to a blank media. The **Blank Media** checkbox is disabled if a specific **Media ID** is selected.
  - **Media Format**: Specify the files to store in the selected media format (**ANTF** or **LTFS**) for a **Custom** store selection. The list is only enabled for the LTO media type. For non-LTO media type, including LATTUS, S3COMPAT, Q-Cloud, and SDISK, the option is disabled and ignored.
4. Do one of the following:
  - a. Click **Apply** to start a job to move the selected files. When the confirmation message appears, click **OK**. If desired, make a note of the Job ID in the message and visit the **Reports > Jobs** page to see the result of the action.
  - b. Click **Reset** to clear all selected files.
  - c. Click **Cancel** to exit the **Tools > File and Directory Actions** page without making changes.

## Modify File Attributes

Select this option to modify attributes for the selected file.

1. To select the files to modify, click the **Browse** button in the **Selected Files** field.
2. On the **Browse** page, do the following:
  - a. The managed relation points will be displayed. Click on a relation point. This displays any subdirectories and files under the relation point.
  - b. If files you want to modify are visible, click **All** to select all visible files, or click the **Select** box next to each file you want to modify.
  - c. If the files are in a subdirectory, click on the directory to view any subdirectories and files under the

directory. Repeat step b.

- d. Once all files you want modify are selected, click **Continue** to return the selected files back to the **Tools > File and Directory Actions** page, or click **Cancel** to return to the **Tools > File and Directory Actions** page without saving selections.
3. At the **File Attributes** field, enter the following options. Required fields are marked with an asterisk.
    - **\*Number of Copies**: Specify to maintain 1, 2, 3, or 4 copies for the file
    - **Store**: Specify whether to store the file **By Policy**, or to **Neverstore** the file
    - **Relocate**: Specify whether to relocate the file **By Policy**, or to **Never** relocate the file
    - **Truncate**: Specify whether to truncate the file **By Policy**, **Immediately After Store**, or to **Never** truncate the file
    - **Stub File Size (KB)**: Specify **By Policy** to use the policy's stub size, or select **Custom** and enter the desired stub file size in the field to the right of **Custom**
  4. Do one of the following:
    - Click **Apply** to start a job to modify the selected files. When the confirmation message appears, click **OK**. If desired, make a not of the Job ID in the message and visit the **Reports > Jobs** page to see the result of the action.
    - Click **Reset** to clear all selected files.
    - Click **Cancel** to exit the **Tools > File and Directory Actions** page without making changes.

## View File Info

Select this option to view detailed information about the selected file.

1. To select the files to view, click the **Browse** button in the **Selected Files** field.
2. On the **Browse** page, do the following:
  - a. The managed relation points will be displayed. Click on a relation point. This displays any subdirectories and files under the relation point.
  - b. If files you want to view are visible, click **All** to select all visible files, or click the **Select** box next to each file you want to view.
  - c. If the files are in a subdirectory, click on the directory to view any subdirectories and files under the directory. Repeat step b.
  - d. Once all files you want view are selected, click **Continue** to return the selected files back to the **Tools > File and Directory Actions** page, or click **Cancel** to return to the **Tools > File and Directory Actions** page without saving selections.
3. Click **File Info** to view information about all selected files.
4. Click **Done** when you are finished viewing file information.



## Assign Affinities

Select this option to assign affinities to the selected directory.

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.
2. Choose **Assign Affinities** from the **Available Actions** drop-down list.
3. Select the directories you want to assign affinities to. If necessary, click **Browse** to navigate to the directory location, and then select the directory
4. Under the **Affinity Options** section, select the affinity you want to assign from the **Affinity to Assign** drop-down list.
5. Click **Apply**.
  - Click **Reset** to discard your changes, and reset the page.
  - Click **Cancel** to discard your changes, and return to the **Home** page.
6. When the confirmation message appears, click **Yes** to proceed, or **No** to abort.

---

## S3 Buckets

The **S3 Buckets** option enables you to **scan** for, **add**, and **delete** S3 buckets.

### Scan Available S3 Buckets

1. On the **Tools** menu, click **S3 Buckets**. The **Tools > S3 Buckets** page appears.
2. Enter the appropriate value into the following parameters:

**Note:** Parameters marked with an asterisk (\*) are required.

Parameter	Description
<b>Username*</b>	Enter the username of the S3 bucket.
<b>Password*</b>	Enter the password of the S3 bucket.
<b>Protocol*</b>	Select <b>http</b> or <b>https</b> as the protocol.
<b>Host*</b>	Enter the hostname of the S3 bucket.
<b>Port</b>	Enter a decimal integer to specify the port number of the S3 bucket.
<b>Bucket Name</b>	<b>(Optional)</b> Enter a name for the S3 bucket.

3. Click **Scan**. The S3 buckets table displays all available S3 buckets.

## Add a S3 Bucket

1. On the **Tools** menu, click **S3 Buckets**. The **Tools >S3 Buckets** page appears.
2. Enter the appropriate value into the following parameters:

**i Note:** Parameters marked with an asterisk (\*) are required.

Parameter	Description
<b>Username*</b>	Enter the username of the S3 bucket.
<b>Password*</b>	Enter the password of the S3 bucket.
<b>Protocol*</b>	Select <b>http</b> or <b>https</b> as the protocol.
<b>Host*</b>	Enter the hostname of the S3 bucket.
<b>Port</b>	Enter a decimal integer to specify the port number of the S3 bucket.
<b>Bucket Name</b>	<b>(Optional)</b> Enter a name for the S3 bucket.

3. Click **Add**. The S3 buckets table displays the newly added S3 bucket.

## Delete a S3 Bucket

1. On the **Tools** menu, click **S3 Buckets**. The **Tools >S3 Buckets** page appears.
2. Enter the appropriate value into the following parameters:

**i Note:** Parameters marked with an asterisk (\*) are required.

Parameter	Description
<b>Username*</b>	Enter the username of the S3 bucket.
<b>Password*</b>	Enter the password of the S3 bucket.
<b>Protocol*</b>	Select <b>http</b> or <b>https</b> as the protocol.
<b>Host*</b>	Enter the hostname of the S3 bucket.
<b>Port</b>	Enter a decimal integer to specify the port number of the S3 bucket.
<b>Bucket Name</b>	<b>(Optional)</b> Enter a name for the S3 bucket.

3. Click **Scan**. The S3 buckets table displays all available S3 buckets.
4. Select a S3 bucket from the S3 buckets table.

5. Click **Delete**. A confirmation dialog is displayed and prompts you to confirm if you want to delete the S3 bucket.
6. Click **Yes** to delete the S3 bucket, or click **No** to abort the operation. The S3 buckets table displays all available S3 buckets upon deletion.

---

## StorNext Connect Ports

The **StorNext Connect Ports** option allows you to enable HTTP ports for StorNext Connect, if they are currently disabled, and also query the current StorNext Connect ports status.

This feature simplifies the StorNext Connect installation so that administrators do not have to enable HTTP via the command line before downloading and installing the StorNext Connect bundle.

**Note:** The **StorNext Connect Ports** menu option is **only** available on Xcellis, Artico, Pro Foundation, and StorNext Metadata Appliances.

### Enable StorNext Connect Ports

**WARNING:** If you enable the HTTP ports for StorNext Connect, you cannot disable the ports from within the StorNext GUI.

1. On the **Tools** menu, click **StorNext Connect Ports**. The **Tools > StorNext Connect Ports** page appears.
2. Click **Enable**. A confirmation dialog box appears.
3. Click **Yes** to enable the HTTP ports, or click **No** to cancel the operation and return to the previous page.
4. **(Optional)** Click **Refresh** to update the current status of the HTTP ports.

### Query the Current StorNext Connect Ports Status

1. On the **Tools** menu, click **StorNext Connect Ports**. The **Tools > StorNext Connect Ports** page appears.
2. Click **Refresh**. The status indicator displays the following:

HTTP Port Status	Description
HTTP ports for StorNext Connect are UNKNOWN	StorNext is not able to retrieve the current HTTP port status, and it is currently not known.

HTTP Port Status	Description
HTTP ports for StorNext Connect are <b>DISABLED</b>	The status of the HTTP ports is currently disabled. If desired, you can enable the HTTP ports. See <b>Enable StorNext Connect Ports</b> .
HTTP ports for StorNext Connect are <b>ENABLED</b>	The status of the HTTP ports is currently enabled.

---

## File Systems Overview

The **Tools > File Systems** menu contains options that enable you to perform the following file system-related tasks:


Menu Option	Description
<a href="#">Label Disks below</a>	Apply EFI or VTOC label names for disk devices that will be used for StorNext File System in your StorNext libraries
<a href="#">Check File System on page 317</a>	Check for and repair StorNext File System metadata corruption due to a system crash, bad disk or other catastrophic failure.
<a href="#">Affinities on page 318</a>	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
<a href="#">Migrate Data on page 319</a>	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
<a href="#">Truncation Parameters on page 321</a>	Enter truncation parameters for your file systems managed by StorNext Storage Manager in order to free up file storage that isn't being actively used
<a href="#">Manage Quotas on page 322</a>	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.

## Label Disks

Each drive used by StorNext must be labeled.

**i Note:** 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. The two types of labels are **EFI** and **VTOC**.


Label Type	Description
<b>EFI</b>	EFI labels are required if you plan to create LUNs that are larger than 2TB.  <b>Note:</b> For Solaris, EFI labels are also required for LUNs with a raw capacity greater than 1TB. EFI labels will not work with the IRIX operating system. EFI labels are the default for StorNext.
<b>VTOC</b>	VTOC labels are used for all operating systems in previous StorNext and Xsan releases, and are still required for the SGI IRIX operating system, Solaris releases prior to Solaris 10 Update 2, and LUNs less than 1TB.

## Label 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 will result in a complete loss of data on that disk device.

1. Choose **Label Disks** from the **Tools > File Systems** menu.
2. **(Optional)** 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. Use the checkbox column to select disks, or hold down Shift while clicking the left mouse button to select multiple disks. If a disk device already has a label, continuing with this procedure overwrites the existing label.
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. If you have selected more than one disk to label, the label name is the base name used by the labeling process, and a unique number is appended to this base name for each disk that is labeled.
6. Click **Label**.
7. When the confirmation message appears, verify that the disk you are labeling is empty, and then click **OK** to proceed, or 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.

## Unlabel a Device

Follow this procedure to remove a label from a previously labeled device. The relabeling process is identical

to labeling initially as described in **Label a Device**.

**⚠ 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.

**i 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. Use the checkbox column to select 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.)

## Check File System

This operation checks for and repairs StorNext file system metadata corruption due to a system crash, bad disk or other catastrophic failure. If the file system is unmounted and stopped, a full check is done and repairs can be performed, if the file system is active, it may only be checked in a Read-only mode. In this mode, modifications are noted, but not committed.

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.

## Methods to Check File Systems

Method	Description
Checking while the file system is offline	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.
Checking while the file system is active	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.

**i 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.

## Perform a File System Check

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 [Manage on page 353](#).

1. Choose **Check File System** from the **Tools > File Systems** menu.
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. You may need to click **Refresh** periodically to see if progress has changed.
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.

## View a Check Report

After you have run at least one file system check, information about the process appears at the bottom of the page: 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 page and then click **Report**. When you are finished viewing the report, click **Done** to return to the previous page.

## Delete a Check 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**.

## Affinities

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 space from which space should be allocated for that file (or files created within the directory).

## Associate an Affinity with a File System's Stripe Group

1. On the **Tools** menu, click **File Systems**.
2. Select the file system.
3. Click **Edit**.
4. Click the **Stripe Group** tab. On this page you can associated an affinity with a stripe group.

For more information about affinities, see the *StorNext 5 User's Guide*.

## Add an Affinity

1. On the **Tools** menu, click **File Systems**, and then click **Affinities**. The **Tools > File Systems > Affinities** page appears.
2. Click **New**. The **New Affinity** page appears.
3. At the **Affinity** field, enter the name of the new affinity.
4. At the **File System** 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.

## Delete an Affinity

1. Select the affinity you want to delete.
2. Click **Delete**.
3. When asked to confirm the deletion, click **Yes** to proceed or **No** to abort.
4. When a message notifies you that the affinity was successfully deleted, click **OK** to continue.

## Migrate Data

This page can be used to either move data files from a source stripe group to other stripe groups, thus freeing the source stripe group so it can be removed from an existing StorNext file system, or to migrate metadata from one disk to another.

During data stripe group migration you indicate a source stripe group from which to move data.

The time it takes to complete the migration process depends on the amount of data being moved. 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 stripe group. All operations (including metadata operations) are handled normally, but no new writes are allowed to the source stripe group, which will be marked read-only.



## Migrate Metadata and Journal Data

1. On the **Tools** menu, click **File Systems**, and then click **Migrate Data**. The **Tools > File Systems > Migrate Data** page appears.
2. Click a **File System**.
3. Click a **Metadata/Journal Stripe Group** in the **File System to migrate**.
4. Click **Migrate**. A new page appears.
5. Click a disk from the **Choose Disk to Migrate From** table.
6. Click a disk from the **Choose Disk to Migrate To** table.
7. Click **Migrate**. The following occurs:
  - The file system stops and is **unmounted**.
  - The GUI runs the command **sndiskmove**.
  - The source disk/LUN is relabeled to **\$LABEL.old**.
  - The destination disk/LUN is relabeled to **\$LABEL**.
8. Click **Refresh** to manually update the status.
9. When completed, start and mount the file system using the GUI.

## Migrate User Data

**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.
2. Click a **File System**.
3. Click a **Stripe Group** in the **File System to migrate**.
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.
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, start and 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**.

## Truncation Parameters

This page enables you to view or change the following information pertinent to the truncation feature as it pertains to StorNext Storage Manager:

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

Parameter	Description
<b>Run</b>	Indicates the current status of the truncation feature ( <b>Enabled</b> or <b>Disabled</b> ).
<b>Mount</b>	Indicates whether the file system is currently mounted.
<b>File System</b>	Displays the name of the truncation-enabled file system.
<b>Mount Point</b>	Shows the mount point for the truncation-enabled file system
<b>Truncation Parameters</b>	Shows the current truncation setting, such as Time-based 75%.

## Update Truncation Parameters

1. On the **Tools** menu, click **File Systems**, and then click **Truncation Parameters**.
2. Click the line containing the file system whose truncation parameters you want to change. Parameters appear at the bottom of the page. As desired, update any of the following fields:

Parameter	Description
<b>Enable Truncation</b>	Choose this field to enable or disable truncation for the selected file system.
<b>Truncation Mode</b>	Select either <b>Time-based</b> or <b>Space-based</b> truncation mode.
<b>Minimum Usage (%)</b>	Enter the usage percentage at which the file system must be before truncation starts taking place on the file system. This parameter is used so that files are not truncated if the file system is not at least this full.
<b>Low Water (%)</b>	Enter the percentage of occupied disk space a file system must reach before StorNext stops applying truncation.
<b>High Water (%)</b>	Enter the percentage of occupied disk space a file system must reach before StorNext starts applying truncation.

3. Click **Apply** to save your changes.
4. When a confirmation message appears, click **Yes** to continue or **No** to abort without saving.

**i 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.

**⚡ WARNING:** If you enable quotas on an existing file system, the file system becomes unavailable while the changes are applied.

For more information about quotas, see the StorNext User's Guide.

This page offers the following functionality:

- You can **view**, **edit**, **delete** or configure **new** quota values as it pertains to the StorNext file system.

The table on this page allows you to **view** the following quota values:

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

Quota Property	Description
<b>Grace Period Remaining</b>	Specifies the grace period remaining on the system. The <b>Grace Period Remaining</b> is the amount of time remaining to free up space to return your usage under the <b>Soft Limit</b> quota value.
<b>Number of Files</b>	Specifies the number of files on the system.

## Edit Current Quota Values for a Specified File System

1. Select the **Directory Namespace** whose quota values you want to edit.
2. Click **Edit...**
  - Configure the following quota value properties:
  - In the **Hard Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down.
  - In the **Soft Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down.
  - In the **Time Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down.
3. Click **Update** to confirm and save your selection, or click **Cancel** to abort and return to the **Manage Quotas** page.

## Configure New Quota Values for a Specified File System

1. Click **New...**
2. Configure the following quota value properties:
  - In the **Hard Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down.
  - In the **Soft Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down.
  - In the **Time Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down.
  - 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 **Manage Quotas** page.

## Delete a Configured Quota Value

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.

## Refresh the Configured Quota Values

Click **Refresh**.

---

# Storage Manager

Storage Manager tasks are described in [Storage Manager Tasks on page 155](#).

---

# Replication and Deduplication

Replication and deduplication tasks are described in [Replication and Deduplication on page 213](#).

---

# High Availability (HA)

HA tasks are described in [Converting to HA on page 344](#).

---

# Upgrade Firmware

For instructions on upgrading your firmware, see [Upgrade the System \(Upgrade Firmware\)](#) on the *Appliance InfoHub Documentation Center* ([www.quantum.com/ApplianceInfoHub](http://www.quantum.com/ApplianceInfoHub)).

# Appliance Release Notes

Refer to the respective Release Notes document for important information you should know about your system.

- [Xcellis Foundation](#)
- [aiWARE for Xcellis](#)
- [Xcellis Workflow Extender](#)
- [Xcellis Workflow Director](#)
- [Artico](#)
- [Pro Foundation](#)
- [G300](#)
- [M660](#)
- [M440](#)



# Chapter 8: Service Menu Functions

---

This chapter contains the following topics:

Service Menu Overview .....	326
Health Check .....	327
Capture State and Capture DSET .....	329
Capture State .....	332
System Backup .....	337
Admin Alerts .....	337
Tickets .....	339
Logging .....	342
Web Services (V2) .....	342

---

## Service Menu Overview

The **Service** menu contains options to monitor and capture system status information.

Menu Option	Description
<b>Health Check</b>	Perform one or more health checks on StorNext and view recent health check results.
<b>Capture State</b>	Obtain and preserve detailed information about the current StorNext system state.
<b>Capture DSET</b>	<p>Obtain and preserve detailed information about the current StorNext system state.</p> <p><b>i Note:</b> The information in this section is applicable for releases up to StorNext 5 release 5.3.x (in other words, prior to StorNext 5 release 5.4.0). Refer to <a href="#">Capture State on page 332</a> for the new information, beginning with StorNext 5 release 5.4.0.</p> <p>Obtain and preserve detailed information about the current state of your StorNext hardware.</p> <p><b>i Note:</b> This option is only visible on StorNext M660, M440, M330, Pro Foundation, Artico Metadata Appliances, and Xcellis Workflow Director systems.</p>
<b>System Backup</b>	Run a backup of StorNext software.
<b>Admin Alerts</b>	View informational messages about system activities.
<b>Tickets</b>	View, edit, or close service tickets generated for the system.
<b>Logging</b>	Enables robust debugging mode for advanced tracing.
<b>Web Services (V2)</b>	Specify your encryption and authentication options.

---

## Health Check

The **Health Check** feature enables you to run various diagnostic checks on your StorNext system. This page displays the following information:

### Parameters on the Health Check Page

Parameter	Description
<b>Test</b>	The name of the test.



Parameter	Description
<b>Start</b>	The time the test began running.
<b>Finish</b>	The time the test completed.
<b>Status</b>	The current or final status for each test.

## Diagnostic Tests Available

Parameter	Description
<b>Archive</b>	Verify that all configured archives are online.
<b>Config</b>	Verify that affinities are configured correctly in SNSM for managed file systems, and that SNSM-managed file systems are identified and configured correctly.
<b>Disk Space</b>	Verify that enough disk space exists for the SNSM database tables, logging, and other functions.
<b>Drive</b>	Verify that all configured drives are online.
<b>Media</b>	Verify that there are enough media available for all policies to store all file copies, and that SNSM media are configured correctly.
<b>Network</b>	Verify SNAPI connectivity.
<b>Policies</b>	Verify that SNSM is keeping up with file system events and store candidate processing.

## Run a Health Check

1. On the **Service** menu, click **Health Check**. The **Service > Health Check** page appears.
2. Select one or more tests to run by clicking the desired check. To select multiple checks, press and hold **Control** 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**.
4. If desired, check job status by clicking **Jobs** on the **Reports** menu. For additional information, see [Jobs on page 365](#).

## View the Health Check Results

After a test has been run successfully (as indicated by **Success** in the **Status** column), you can view test results.

- To view results for one or more tests, select the desired tests and then click **View Selected**.
- To view results for all successfully completed tests, click **View All**.

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.

## View the 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 page appears.
3. When you are finished viewing history information, click **Done** to return to the **Service > Health Check** page.

---

# Capture State and Capture DSET

The **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.

**i Note:** The information in this section is applicable for releases up to StorNext 5 release 5.3.x (in other words, prior to StorNext 5 release 5.4.0). Refer to [Capture State on page 332](#) for the new information, beginning with StorNext 5 release 5.4.0.

Running **Capture State** creates a log file in this format:

```
snapshot-machinehostname-YYYYMMDDHHMMSS.tar.gz
```

This file contains a summary report that is produced by executing the `pse_snapshot` command on all component config/filelist files.

If desired, you can download or delete a previously captured file.

**i Note:** If you run the snapshot directly on each node from the **Service** menu, the process creates a log file in this format:

```
snapshot-YYYYMMDDHHMMSS.tar.gz
```

Similarly, the Capture Dell System E-Support Tool (**DSET**) feature enables you to create a DSET report which captures current status for your StorNext hardware.

The DSET report contains an array of status information about the StorNext hardware. Capture DSET is similar to the Capture State tool, except the Capture DSET log assists Quantum support personnel in identifying and diagnosing problems specific to the StorNext hardware.

Running Capture DSET creates a log file named using the format:

```
"dset-machinehostname-YYYYMMDDHHMMSS.zip
```

If desired, you can download or delete either capture state files or captured DSET files.

- i Note:** The **Capture DSET** menu option is **only** available on Xcellis, Artico, Pro Foundation, and StorNext Metadata Appliances.

## Create a Capture State Log

1. On the **Service** menu, click **Capture State**. The **Service > Capture State** page appears. Any previously captured snapshots are shown.

**i Note:** Captures are stored in the directory `/usr/adic/HAM/shared/capture_state/`. On a non-HA system, captures are stored in the directory `/usr/adic/gui/logs/capture_state` on the StorNext server.

2. Click **Capture**. The **Capture State Status** page appears. The capture file appears after the process completes.
3. If desired, check job status by clicking the **Reports** menu, then clicking **Jobs**. See [Jobs on page 365](#) for more information.
4. At any time you can click **Refresh** to update the capture state status shown on the page.
5. When the capture completes, click **Download** to save the generated file. If the download does not start automatically, click the supplied link.
6. To view the file, click **Open With** and then click **Browse** to navigate to an application such as WinZip capable of reading tar.gz files.
7. To save the file, select **Save to Disk** and then navigate to the location where you want to save the file.

## Delete a Previous System State Capture

1. On the **Service** menu, click **Capture State**. The **Service > Capture State** page appears. All previously captured snapshots are shown.
2. Select the file you want to delete, and then click **Delete**.
3. When a confirmation page prompts you to confirm that you want to delete the file, click **Yes** to continue or **No** to abort.
4. After the status page informs you that the file was successfully deleted, click **OK**.

## Create a Capture State for a HA Secondary Node

- Click **Capture Secondary**.

- i Note:** When you use the Capture State feature on an HA system, system information for the primary node is captured by default.

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.

## Create a Capture DSET File

1. On the **Service** menu, click **Capture DSET**. The **Service > Capture DSET** page appears. Any previously captured DSET snapshots are shown.
2. Click **Capture**. The **Capture DSET Status** page appears. The capture DSET file name appears in the list after the process completes.

## 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**.
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.

## Delete a Previous Capture DSET File

1. On the **Service > Capture DSET** page, select the DSET file you want to delete.
2. Click **Delete**.
3. When a confirmation page prompts you to confirm that you want to delete the file, click **Yes** to continue or **No** to abort.
4. After the status page informs you that the file was successfully deleted, click **OK**.

# Capture State

The **Capture State** feature enables you to create a log file that captures the current state of the system. The log file assists Quantum Support personnel analyze and debug problems with your system.

Assuming your system is connected to the internet and you have an open support case, you can automatically, or manually upload your log files to Quantum Support via Secure File Transfer Protocol (SFTP).

Additionally, you can capture, download and upload just file system logs and platform logs. All logs are included in a **Full Snapshot**, but in some cases you can create a log that consists of a smaller sub-component.

## General Considerations

- **Upload Progress:** The GUI reminds you that you can check the progress of the upload by reviewing the job on the **Reports > Jobs** page.
- **Dark Sites and Non-Internet Accessed Servers:** If either of these scenarios apply to your environment, you must continue to download captured log files to a local client with internet access using the **Download** button in the GUI.

The table below provides the information displayed on the **Service > Capture State** page:

Parameter	Description
<b>Capture Type</b>	<p>Specifies the type of capture file to create. The available options are:</p> <ul style="list-style-type: none"><li>• <b>Full Snapshot:</b> Provides a capture file containing all logs available within the individual capture types (<b>File System Only</b>, <b>Platform Only</b>, and <b>DSET</b>).</li><li>• <b>File System Only:</b> Provides a capture file containing the system's StorNext file system debug logs, configuration, version information and disk devices. The file contains client debug information on client machines and server information on server machines, as well as port-mapper information from all machines. System log files as well as SNFS log files are included.</li><li>• <b>Platform Only:</b> Provides a capture file containing log files and configuration files for the StorNext server operating system and Quantum application logs. The StorNext software specific logs files are not contained in this capture.</li></ul> <p><b>i Note:</b> This menu option is only available on Xcellis, Artico, Pro Foundation, and StorNext Metadata Appliances.</p> <ul style="list-style-type: none"><li>• <b>Dell System E-Support Tool (DSET), :</b> Provides a capture file containing an array of status information about the StorNext hardware and assists Quantum Support personnel in identifying and diagnosing problems specific to the StorNext hardware.</li></ul> <p><b>i Note:</b> This menu option is only available on Xcellis, Artico, Pro Foundation, and StorNext Metadata Appliances.</p>

Parameter	Description
<b>File System</b> (only displayed if the <b>Capture Type</b> is <b>File System Only</b> )	Displays the available file systems on your server ( <b>managed, shared, unmanaged</b> ). Select a file system, or <b>Ctrl-click</b> to select multiple file systems for which to create a capture file.
<b>Target Node</b>	If your system is configured for High Availability (HA), this option specifies the server ( <b>Both Nodes, Primary</b> only, or <b>Secondary</b> only) for which the capture file is being generated for.
<b>Serial Number</b>	Specifies the serial number of your system.
<b>Enable Quantum Upload</b>	<p>If you have a live internet connection and an open support case, check this box to <b>automatically</b> upload captured log files to Quantum Support using SFTP after the capture file is generated.</p> <hr/> <p><b>i Note:</b> You must have or obtain a <b>Service Request/Case ID</b> number to upload your file to the Quantum Support server. If you do not have a <b>Service Request/Case ID</b> number available, call Quantum Support as an open <b>Service Request/Case ID</b> number is required. This feature also requires a live internet connection on your server.</p>
<b>Server</b> (only displayed if <b>Upload to Quantum</b> is checked)	Specifies the Quantum Support server where capture files are uploaded to.
<b>Service Request / Case ID</b> (only displayed if <b>Enable Quantum Upload</b> is checked)	<p>This field allows you to input a <b>Service Request</b> number provided to you by Quantum Support.</p> <hr/> <p><b>i Note:</b> You must have or obtain a <b>Service Request/Case ID</b> number to upload your file to the Quantum Support server. If you do not have a <b>Service Request/Case ID</b> number available, call Quantum Support as an open <b>Service Request/Case ID</b> number is required. This feature also requires a live internet connection on your server.</p>

Parameter	Description
<b>File</b>	<p>Specifies the capture file name. Creating a capture generates a log file in the following format:</p> <p><i>&lt;Serial Number&gt;_&lt;Machine Hostname&gt;_&lt;Capture Type&gt;_&lt;Time Stamp&gt;.&lt;Filename Extension&gt;</i>, where the <i>Time Stamp</i> is in the format <i>YYYYMMDDHHMMSS</i>.</p> <p>Example:</p> <pre>SV123ABCF12345_xps-sles11_snapshot_20160831072620.tar.gz</pre>
<b>Type</b>	<p>Specifies the type of capture file:</p> <ul style="list-style-type: none"><li>• <b>Full Snapshot</b></li><li>• <b>File System Only</b></li><li>• <b>Platform Only</b></li><li>• <b>Dell System E-Support Tool (DSET)</b></li></ul>
<b>Date</b>	<p>Specifies the date and time on which the capture file was created.</p>
<b>Size</b>	<p>Specifies the file size of the capture file.</p>
<b>Capture</b>	<p>Click to manually generate a capture file.</p>
<b>Capture &amp; Upload</b>	<p>Click to automatically upload a selected capture file via SFTP to a Quantum Support server.</p> <p><b>i</b> <b>Note:</b> You must have or obtain a <b>Service Request/Case ID</b> number to upload your file to the Quantum Support server. If you do not have a <b>Service Request/Case ID</b> number available, call Quantum Support as an open <b>Service Request/Case ID</b> number is required. This feature also requires a live internet connection on your server.</p>
<b>Download</b>	<p>Click to download a selected capture file to a local destination.</p>
<b>Upload to Quantum</b>	<p>Click to manually upload a selected capture file via secure FTP to a Quantum Support server.</p> <p><b>i</b> <b>Note:</b> You must have or obtain a <b>Service Request/Case ID</b> number to upload your file to the Quantum Support server. If you do not have a <b>Service Request/Case ID</b> number available, call Quantum Support as an open <b>Service Request/Case ID</b> number is required. This feature also requires a live internet connection on your server.</p>
<b>Delete</b>	<p>Click to delete the selected capture file(s).</p>
<b>Refresh</b>	<p>Click to update the <b>Capture State</b> page.</p>

## Create a Capture State File

1. On the **Service** menu, click **Capture State**. The **Service > Capture State** page appears. Any

previously generated capture state files are displayed.

2. In the **Capture State Options** section:
  - a. In the **Capture Type** list, select the type of capture state file to generate.

---

**i Note:** If you select **File System Only**, the File System option appears and displays the available file systems on your server (**managed, shared, unmanaged**). Select a file system, or **Ctrl-click** to select multiple file systems for which to create a capture file.

- b. In the **Target Node** list, select the server (**Both Nodes, Primary** only, or **Secondary** only) for which the capture file is being generated for.
3. Click **Capture**. A confirmation box appears.
4. Click **OK**. The capture state file name appears in the list after the process completes.
5. **(Optional)** On the **Reports** menu, click **Jobs** to monitor the job progress.

## Create a Capture State File and Automatically Upload to the Quantum Support Server

---

**i Note:** You must have or obtain a **Service Request/Case ID** number to upload your file to the Quantum Support server. If you do not have a **Service Request/Case ID** number available, call Quantum Support as an open **Service Request/Case ID** number is required. This feature also requires a live internet connection on your server.

1. On the **Service** menu, click **Capture State**. The **Service > Capture State** page appears. Any previously generated capture state files are displayed.
2. In the **Capture State Options** section:
  - a. In the **Capture Type** list, select the type of capture state file to generate.

---

**i Note:** If you select **File System Only**, the File System option appears and displays the available file systems on your server (**managed, shared, unmanaged**). Select a file system, or **Ctrl-click** to select multiple file systems for which to create a capture file.

- b. In the **Target Node** list, select the server (**Both Nodes, Primary** only, or **Secondary** only) for which the capture file is being generated for.
  - c. Click **Enable Quantum Upload**. The **Server** and **Service Request / Case ID** items appear.
  - d. In the **Service Request / Case ID** field, input the number provided to you by Quantum Support.
3. Click **Capture & Upload**. A confirmation box appears.
4. Click **OK**. The capture state file name appears in the list after the process completes.
5. **(Optional)** On the **Reports** menu, click **Jobs** to monitor the job progress.



## Upload a Capture State File Manually to the Quantum Support Server

- i Note:** You must have or obtain a **Service Request/Case ID** number to upload your file to the Quantum Support server. If you do not have a **Service Request/Case ID** number available, call Quantum Support as an open **Service Request/Case ID** number is required. This feature also requires a live internet connection on your server.
1. On the **Service > Capture State** page, select the file you want to upload.
  2. In the **Capture State Options** section, check **Enable Quantum Upload**. The **Server** and **Service Request / Case ID** items appear.
  3. In the **Service Request / Case ID** field, input the number provided to you by Quantum Support.
  4. Click **Upload to Quantum**. You are prompted to verify your upload.
  5. Click **Yes** to begin the upload, or click **No** to cancel the operation and return to the previous page. If you click **Yes**, a notification appears informing you that the upload has initiated.
    - Click **OK**.
  6. **(Optional)** On the **Reports** menu, click **Jobs** to monitor the job progress.

## Download a Capture State File

1. On the **Service > Capture State** page, select the file you want to download.
2. Click **Download**.
3. When prompted, click the file name to begin the download process.
4. Navigate to the location where you want to save the file.
5. After the file is saved, click **Done**.
6. **(Optional)** If you were instructed by Quantum Support personnel to send the capture file, email the file to the email address provided to you by the Quantum Support representative.

## Delete a Capture State File

1. On the **Service > Capture State** page, select the file you want to delete. Alternatively, you can select multiple files.
2. Click **Delete**.
3. When a confirmation page prompts you to confirm that you want to delete the file(s), click **Yes** to continue or **No** to abort.
4. After the status page informs you that the file(s) was/were successfully deleted, click **OK**.

## Refresh the Capture State Page

Click **Refresh** to update the **Capture State** page.

---

# System Backup

The **System Backup** option allows you to perform a full or partial 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. Also by default, a partial backup is run on all other days of the week that the full backup is not run. This backup includes database journals, configuration files, and file system journal files.

---

**i Note:** StorNext must be configured with a storage media such as tape, sdisk, or Lattus in order to perform a system backup, or an error will occur.

---

**i Note:** Quantum recommends making two or more backup copies to minimize vulnerability to data loss in the event of hardware failure.

## Parameters on the System Backup Page

Parameter	Description
<b>Copy</b>	The backup copy number
<b>Date</b>	The date the backup was initiated
<b>Type</b>	The type of backup ( <b>Full</b> or <b>Partial</b> ).
<b>Status</b>	The current or final status of the backup operation (for example, PASS or FAIL)
<b>Media</b>	The media ID associated with the backup

## Perform a System Backup

1. On the **Service** menu, click **System Backup**. The **Service > System Backup** page appears.
2. Click **Full Backup** to perform a full backup, or click **Partial Backup** to perform a partial backup.
3. If desired, check job status by choosing **Jobs** from the **Reports** menu. See [Jobs on page 365](#) for more information.
4. At any time you can click **Refresh** to update the backup status shown on the page.
5. 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.

## Parameters on the Service > Admin Alerts Page

Parameter	Description
<b>Alert ID</b>	The unique identifier for the alert.
<b>Date</b>	The date the alert was generated.
<b>Request ID</b>	The unique StorNext identifier which generated the alert.
<b>Identifier</b>	The specific StorNext software component which generated the alert.
<b>Detail</b>	A description of the alert.

## Possible Conditions that Generate an Admin Alert

- Health Checks disk space warning
- Intrusive Health Checks when drives are mounted
- Media console errors
- Drive dismount request when drive is already dismounted
- Media audit failures

## View an Admin Alert

On the **Service** menu, click **Admin Alerts**. The **Service > Admin Alerts** page appears. View a specific alert by scrolling to the right of the page (if the alert is longer than can be displayed on one page)

## Refresh the Admin Alert List

Click **Refresh** to update the list.

## Delete an Individual Admin Alert

1. Click the radio-button on the left to select an individual alert.
2. Click **Delete**.
3. Click **Yes** to confirm and delete the selected alert, or click **No** to cancel and return to the **Service > Admin Alerts** page.

## Delete Multiple Admin Alerts

1. Click **Delete All**.
2. Click **Yes** to confirm and delete the selected alerts, or click **No** to cancel and return to the **Service > Admin Alerts** page.

---

## 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 page 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.

## Change the Table View

There are three ways to display ticket information:

Menu Option	Description
<b>Page</b>	View tickets page by page, using the navigation controls at the lower right side of the table.
<b>Scroll</b>	Use the scroll bar at the right side of the table to change the current view.
<b>All</b>	Similar to <b>Scroll</b> mode, except the table expands to show all tickets. In <b>Scroll</b> mode the table remains the same size.

To change the current view, select **Page**, **Scroll** or **All** from the **Table View** drop-down field above the table on the right side of the screen.

## How to Use 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.

---

**i Note:** When everything is selected, **Filter Options** are **Not Active**.

1. On the **Service > Tickets** page, click the arrow to the right of the **Filter Options** heading to display the options.
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.

**i Note:** The **Filter Options** status changes to **Active** after you make changes.

## View a Ticket

1. On the **Service** menu, click **Tickets**. The **Service > Tickets** page appears and provides the following information:

Parameter	Description
<b>Ticket</b>	The RAS ticket number, displayed in the order in which it was created.
<b>Status</b>	The ticket's current status (Open or Closed)
<b>Priority</b>	The ticket's priority based on system impact (High, Medium, or Low).
<b>Last Update</b>	The date of the last system status update.
<b>Summary</b>	A short summary of the fault that triggered creating the RAS ticket.
<b>Event</b>	A short summary of the fault event that triggered creating the RAS ticket, such as <b>Data Corruption</b> .

2. If desired, change the display by choosing **Show All Tickets**, **Show Closed Tickets**, or **Show Open Tickets** in the list at the bottom of the page.
3. Highlight the ticket you wish to view, and then click **View**. The **Service > Tickets > View Ticket > [number]** page appears. This page provides the following information:

Parameter	Description
<b>Ticket Number</b>	The number of the ticket in the displayed ticket list
<b>Date Opened</b>	The date and time the ticket was created
<b>Status</b>	The current status of the ticket ( <b>Open</b> or <b>Closed</b> ).
<b>Priority</b>	The ticket's priority based on system impact ( <b>High</b> , <b>Medium</b> , or <b>Low</b> ).
<b>Summary</b>	A brief description of the ticket.

Parameter	Description
<b>Event Details</b>	Detailed information about event that triggered the ticket, including a link that allows you to view recommended actions which helps correct the fault or condition.
<b>Analysis</b>	Any user-entered comments pertaining to the fault or condition, such as a recommended action

4. To see recommended actions for the ticket, click **View Recommended Actions**. The **Recommended Actions** page provides information and steps to correct the condition or fault that generated the RAS ticket. Follow the instructions on the page 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** page.

## Edit a Ticket

1. Select the desired ticket and then click **Edit**. The **Service > Tickets > Edit Ticket > [number]** page appears.
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 page, click **Close**, or click **Cancel** to return to the previous page without saving your changes.

## Close a Ticket

- To close a specific ticket, select the desired ticket and then click **Close**.

## Close Multiple Tickets

- To delete all tickets, click **Close All**.

## Delete a Ticket

- Select a ticket from the tickets table, and then click **Delete**. This procedure allows you to delete one ticket at a time.

## Delete Multiple Tickets

- 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 **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 additional information about logs, see [Logs on page 364](#).

**i 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.

## Enable Logging

1. On the **Service** menu, click **Logging**. The **Service > Logging** page appears.
2. As necessary, enable logging for any of the following system components:

System Component	Description
<b>Application Debug Logging</b>	This option enables debugging for StorNext.
<b>Replication/Deduplication Debug Logging</b>	This option enables debugging for StorNext policies such as SNPPolicy.
<b>Web Services Debug Header Logging</b>	This option enables debugging for Web-related components specific to Web headers.
<b>Web Services Debug Content Logging</b>	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.

---

# Web Services (V2)

The **Service** menu's **Web Services (V2)** option allows you specify your encryption and authentication options.

## Display the Web Services (V2) Page

- On the **Service** menu, click **Web Services (V2)**.

## Specify Your Encryption and Authentication Options

1. In the **State** list, select the value you wish to apply. The table below provides a description of the values you can apply.
2. In the **Protocol** list, select the value you wish to apply. The table below provides a description of the values you can apply.
3. In the **Authentication Type** list, select the value you wish to apply. The table below provides a description of the values you can apply.
4. Click **Apply** to confirm your changes, or click **Cancel** to discard your changes.

Parameter	Description
<b>State</b>	Displays the state of web service. The valid values are <b>On</b> and <b>Off</b> . The default value is <b>Off</b> .
<b>Protocol</b>	Displays the request protocol that will be used to access the web service. The valid values are <b>http</b> , <b>https</b> , and <b>http or https</b> . The default value is <b>http</b> .
<b>Authentication Type</b>	Displays the authentication type that will be used to access the web service. The valid values are <b>None</b> , and <b>User</b> . The default value is <b>None</b> . If the <b>State</b> is set to <b>Off</b> , the <b>Protocol</b> is automatically updated to <b>http</b> and the <b>Authentication Type</b> is updated to <b>None</b> .





# 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 [High Availability Systems on page 548](#).

This chapter contains the following topics:

HA Overview .....	344
HA Terms and Concepts .....	345
Preparing for HA Conversion .....	346
Converting to HA .....	348
HA Statuses and Reporting .....	354
Troubleshooting HA .....	355

---

## 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 below](#) and [Secondary Node below](#). For additional information, see [FSM Failover In HA Environments on page 573](#).

### 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:

- The vIP should be static (currently StorNext supports only static IP for HA).
- The NIC should have a *physical* IP address assigned.
- The vIP should be a real and unique IP address.
- 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 563](#).

## 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 5 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. Using the StorNext GUI, enter the StorNext license information on both server nodes.

---

**i Note:** The StorNext license information on both server nodes must be identical.

5. Using a web browser, launch the StorNext GUI on one server; use the server to perform the HA conversion process.
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 [File Systems on page 34](#).
  - 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 563](#).

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 below](#).

## 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

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.

- i 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 30](#).
- 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.
- StorNext operating files will be moved to the HaShared file system, and this move cannot easily be reversed.
- The Reduplication/Deduplication bullet item applies to all MDCs used in an HA configuration, including StorNext Metadata Appliances. 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

## Configure HA

1. Choose **Tools > High Availability > Convert**.

---

**i 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.

2. 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.
3. **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.)
4. Click **Convert** to convert the primary MDC node to HA and click **Yes** to confirm the conversion.
5. Once the primary MDC node has been converted, **Status** will change to **Converted**. Click **OK** to continue.
6. 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).
7. 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.)

---

**i 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.

8. Click **OK** to convert the secondary node.
9. 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 node provides a message stating it is not the primary node and a link to launch the primary node.

---

**! Caution:** Do not login to the GUI of the secondary 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 **Start** to restart the Storage Manager components.
- Wait until the system icons for both nodes of the system as well as File System and Storage Manager are green, which indicates normal operation.
- 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.

## Initiate the Failover of a System that Has Just Been Converted to HA

1. Open an SSH connection to the MDC node operating as the **primary**.
2. Login to the command line of the **primary** MDC node.
3. Confirm that the MDC node is operating as the **primary** by entering the following at the command prompt:

```
snhamgr -m status
```

4. Verify the output is (bold used for clarification):

```
:default:primary:default:running:
```

5. Initiate an HA failover on the **primary** MDC node to the **secondary** MDC node.
  - a. Issue the following command to stop StorNext services on the **primary**MDC node:

```
adic_control stop
```

- b. Then, issue the following command to initiate the failover:

```
snhamgr force smith
```

6. Wait 3 minutes until the MDC node previously operating as the **secondary** becomes the **primary**.
7. Open an SSH connection to the MDC node now operating as the **primary**.
8. Login to the command line of the **primary** MDC node.
9. Confirm that the MDC node is operating as the **primary** by entering the following at the command prompt:

```
snhamgr -m status
```

10. Verify the output is:

```
:default:primary:default:running:
```

11. Repeat if desired to fail over to the original system operating as the **primary**.

**i Note:** When the **force smith** command is used on an MDC node, the system will reboot, and it may take a significant amount of time for the MDC node to come back online, so plan for this additional delay. Wait until the MDC node reboot has completed before initiating another fail over to the MDC node originally operating as the **primary**.

12. Once the failover has completed, restart SNFS services on any clients that were stopped earlier.

13. Mount the SNFS file systems on each client machine, if needed.

14. Verify that all clients have full access.

15. Verify access to all file systems and move files to/from disk and tape.

## Initiate the Graceful Failover of an HA Pair

1. Open an SSH connection to the node operating as the **primary**.
2. Login to the command line of the **primary** node as root.
3. Confirm that the node is operating as the **primary** by entering the following at the command line:

```
snhamgr -m status
```

4. Verify the output is (bold used for clarification):

```
:default:primary:default:running:
```

5. On the node operating as the **primary**, initiate an HA failover to the node operating as the **secondary**.

```
service cvfs stop
```

6. Wait until the **secondary** node becomes the **primary**, and leave your SSH connection to this node open. Time may vary.



7. Open an SSH connection to the node now operating as the **primary**.
8. Login to the command line of the **primary** node as root and enter the following:
9. Confirm that the node is operating as the **primary** by entering the following at the command line:

```
snhamgr -m status
```

10. Verify the output is:

```
:default:primary:default:stopped:
```

11. From the SSH connection to the node now operating as the **secondary**, enter the following:

```
service cvfs start
```

12. Confirm that the node is operating as the **secondary** by entering the following at the command line:

```
snhamgr -m status
```

13. Verify the output is:

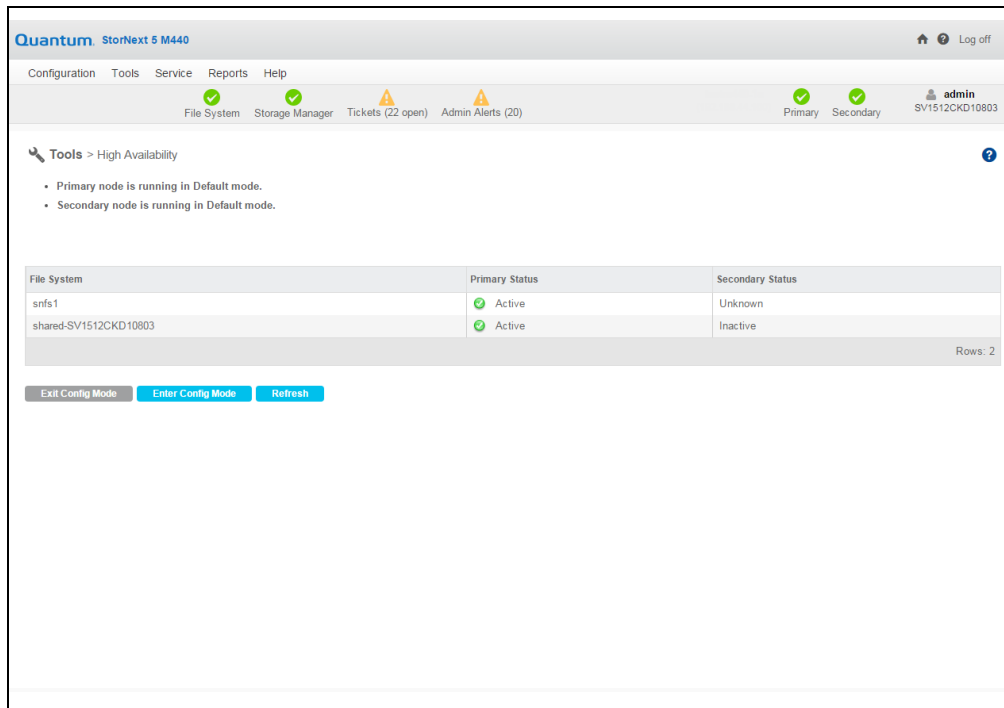
```
:default:running:default:primary:
```

14. Repeat if desired to fail over to the original system operating as the **primary**.
15. Verify that all clients have full access.
16. Test access to all file systems.

## 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** page appears.

Figure 46: Manage HA Page



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.

## Manage

The **Manage** option enables you to view the current status of the file systems on your HA systems. Specifically, you can view whether the file systems on your primary and secondary nodes have a status of **Active**, **Inactive**, **Standby**, or **Unknown**.

The page 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.

This page also enables you to lock the HA cluster for administration purposes, placing the cluster into **Config** (configuration) mode so your system administrator can make configuration changes and other modifications. This mode allows the HA Manager to prevent HA resets while making configuration changes or stopping the primary server.

---

**⚠ 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.

---

**⚠ Caution:** Setting `haFsType` to **HaUnmonitored** disables the HA monitor timers used to guarantee against split brain. When two MDCs are configured to run as an HA pair but full HA protection is disabled in this way, it is possible in rare situations for file system metadata to become corrupt if there are lengthy delays or excessive loads in the LAN and SAN networks that prevent an active FSM from maintaining its branding of the ARB in a timely manner.

## Lock the HA Cluster and Enter Config Mode (and Subsequently Exit Config Mode)

**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.

---

**i 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. On the **Tools** menu, click **High Availability**, and then click **Manage**.
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** to start both nodes of the HA cluster in **default** mode.
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

This section contains troubleshooting suggestions for issues which pertain to StorNext HA (high availability) systems. For an in-depth look at HA systems and operation, see [High Availability Systems on page 548](#).

### 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 570](#) in the HA topic.

## 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 'fsmList'` command, which displays the local FSMPM's internal FSM and process table. A remote FSMPM's table can also be viewed with `'fsmList on <MDC name or address>'`.

Finally, the message: 'Admin Tap Connection to FSM failed', is an intermittent response that occurred because the timing of the `cvadmin select` command was during the period after the FSM failed the fifth usurpation attempt and before the FSM was restarted (a ten-second delay). At other times, the response will show an activating FSM. Note that the `cvadmin`-displayed asterisk simply indicates that the FSM has been told to activate, not that it has been successful at usurpation and activation.

**Question:** Using the same configuration above (Shasta and Tahoe), an HA Reset occurs if I pull the fibre connection from Shasta when it is the Primary MDC, but it takes 30-40 seconds. Why does it take so long?

**Answer:** When the fibre connection is lost, Shasta's FSMs cannot maintain their brands on their ARB blocks, so the HA timers do not get restarted in the *read, write, restart-timer ARB branding* loop. After five seconds the timers would expire and reset the MDC. However, there is a second method for resetting the timers that uses the LAN.

Every two seconds, the 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 fsm`list` 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
```



# Chapter 10: StorNext Reports

---

This chapter contains the following topics:

Reports Menu Overview .....	362
Logs .....	364
Jobs .....	365
Files .....	369
Drives .....	369
Media .....	370
Q-Cloud Object Storage Media Usage .....	372
Relations .....	373
File Systems .....	374
SAN Devices .....	374
Tape Consolidation .....	375
SAN and LAN Clients .....	376
LAN Client Performance .....	377
Replication / Deduplication Reports .....	378
Distributed Data Mover .....	382
Hardware Status .....	383

Gateway Metrics .....	385
StorNext Metrics .....	389

---

# Reports Menu Overview

The **Reports** menu contains options to view StorNext reports.


Menu Option	Description
<b>Logs</b>	Access logs of StorNext operations.
<b>Jobs</b>	View a list of pending and completed jobs on the system.
<b>Files</b>	View information about specific files, such as the owner, group, policy class, permissions, and copy information.
<b>Drives</b>	View information about the drives in your libraries, including the serial number and current state and status.
<b>Media</b>	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.
<b>Q-Cloud Object Storage Media Usage</b>	View the usage report for Q-Cloud object store media.
<b>Relations</b>	View the names of the policy classes which correspond to the managed directories in your system.
<b>File Systems</b>	View file system statistics including active clients, space, size, disks, and stripe groups.
<b>SAN Devices</b>	View information about devices currently attached to your SAN, including disks/LUNs, libraries, and tape drives.
<b>Tape Consolidation</b>	View statistics on the tape consolidation (defragmenting) process.
<b>SAN and LAN Clients</b>	View statistics for StorNext clients, including the number of connected clients and LAN Clients, and client performance.
<b>LAN Client Performance</b>	View information about LAN Clients and servers, including read and write speed.
<b>Replication/Deduplication: Policy Activity</b>	View replication and deduplication performance statistics.


Menu Option	Description
<b>Replication/Deduplication: Policy Summary</b>	View replication and deduplication information for each policy.
<b>Hardware Status</b>	View up-to-date information about the system board and network ports for both nodes in your HA system, plus storage arrays. <b>Note:</b> This option is only visible on StorNext M660, M440, M330, Pro Foundation, Artico Metadata Appliances, and Xcellis Workflow Director systems.
<b>Distributed Data Mover</b>	View activity related to the Distributed Data Mover feature.
<b>Gateway Metrics</b>	View information and activity related to your gateways, clients, and file systems.
<b>StorNext Metrics</b>	View information and activity related to StorNext Metrics. The StorNext Metrics reports provide performance data logging and visual reporting and graphing features for StorNext systems. The StorNext Metrics reports are a visual reporting tool that combines comprehensive performance data logging with powerful visual reporting and analysis tools to help you identify potential problems and optimize system operations.


## Report Navigation Controls


If a log or report spans more than one page, navigation controls at the bottom of the page 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. If there are more than 10 pages, clicking this button moves the display 10 pages backwards from the currently selected page.

Click  to go to previous page

Click  to go to the next page

Click  to skip ahead ten pages. If there are more than 10 pages, clicking this button advances the display 10 pages forward from the currently selected page.

Click  to go to the last page

Click a specific page number  to go to that page

---

# Logs

The **Reports** menu option enables you to access and view any of the logs described in the following section.

## Types of Report Logs

Type of Logs	Description
<b>StorNext Logs</b>	Logs about each configured file system
<b>File Manager Logs</b>	Logs that track storage errors, etc. of the Storage Manager
<b>Library Manager Logs</b>	Logs that track library events and status
<b>Server System Logs</b>	Logs that record system messages
<b>Web Server Logs</b>	Various logs related to the web server
<b>Database Logs</b>	Logs that track changes to the internal database
<b>Replication/Deduplication Logs</b>	Logs that track the progress and status of data replication or data deduplication operations.
<b>HA Logs</b>	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.
<b>Shared Firmware Upgrade Logs</b>	Logs that track firmware upgrade status of both nodes.
<b>Local Firmware Upgrade Logs</b>	Logs that track firmware upgrade status of the current primary node.
<b>Quota Logs</b>	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.
- **fsaddrrelation** failure due to an inconsistency with policy attributes between replication/deduplication and **fspolicy**.

## Access the StorNext Log Files

The process is the same regardless of the type of log you are viewing.

1. On the **Reports** menu, click **Logs**. The **Reports > Logs** page appears.
2. On the left side of the page, 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 page, use the navigation controls at the bottom of the page as described in [Reports Menu Overview on page 362](#).

---

## Jobs

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 page if there are multiple screens of jobs.

## Parameters on the Jobs Page

Parameter	Description
<b>ID</b>	The job ID number.
<b>Job</b>	The job name assigned by StorNext for the type of action performed (for example, "FileSystem Make").
<b>Attributes</b>	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.
<b>User</b>	The logged in user who initiated the job.
<b>Start and End Time</b>	The times the job was started and ended.
<b>Status</b>	The job's final or current status, such as Success or Failure.

## View 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 page. When you are finished viewing that job's information, click **Done**.

## Filter Options

The **Status Filter** at the top of the page 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 default **Status Filter** is **Working**.

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	Close RAS Ticket	Library Operator	Rename File System
Add DDM Host	Configure Alternate Retrieval Location	Make File System	Replicate
Add Drive Pool	Convert to HA	Manage Logging	Reset Schedule
Add Drives to Drive Pool	Create Blockpool	Manual Move Media	Restart System Service
Add Email Contact	Delete DDM Host	Modify Email Contact	Resume Deduplication

Add Email Notification	Delete Drive Pool	Modify Email Notification	Resume Replication
Add Email Server	Delete Email Contact	Modify Email Server	Resume Truncation
Add File System	Delete Email Notification	Modify File Attributes	Retrieve Directory
Add Library	Delete Email Server	Modify File System	Retrieve Files
Add Media	Delete File System	Modify Media	Run Store Policy
Add Object Storage	Delete Job	Modify Object Storage	Scan File System For New Storage
Add Policy	Delete Library	Modify Policy	Set Directory Affinity
Add Schedule	Delete Media	Modify Q-Cloud Configuration	Set Object Storage Media Availability
Add Storage Disk	Delete Object Storage	Modify Schedule	Start File System
Add Tape Drive	Delete Policy	Modify Storage Disk	Start System Service
Add User	Delete Schedule	Modify Tape Drive	Stop Blockpool
Assign Media to Policy Class	Delete Storage Disk	Modify User	Stop Deduplication
Audit Library Discrepancy Report	Delete Tape Drive	Mount File System	Stop File System
Audit Library Physical Inventory	Delete User	Mount Media	Stop Replication
Cancel Request	Dismount Media	Move Files	Stop System Service
Capture DSET	Drive Device Path Audit	Move Media	Stop Truncation
Capture State	Drives Validation Report	Move StripeGroups	Store Files
Capture State Remote	Drives Validation Update	Pause Deduplication	System Backup
Change File Version	Eject Media	Pause Replication	Tape Consolidation Report
Change Media State	Enter Media	Pause Truncation	Test Replication
Check File System	Expand File System	Purge Media	Transcribe Media



---

Clean Drive	Force Replication	Reassign Orphaned Media	Truncate Files
Clean Media by File System	HA Manager Request	Reclassify Media	Unknown
Clean Media by Media ID	Health Check	Recover Directories	Unmount File System
Clean Media by Policy Class	Import Media Bulk Load	Recover Files	Update DDM Host
Clear Q-Cloud Configuration	Import Media From Mailbox	Remove Drives from Drive Pool	Update RAS Ticket Analysis
Close All RAS Tickets	Library Modify	Remove Media	Validate Library Slot Count

---


## Delete Jobs

Click **Delete** to delete the selected job and all log files. This will free up disk space as a large number of jobs may cause problems on HA systems. Deleting old jobs may also improve the GUI performance on this page as it does not have to process as many database items.

---

 **Caution:** This cannot be undone.

---

 **Note:** Jobs that are currently in a **Working** state cannot be deleted.

## Delete Filtered Jobs

Click **Delete Filtered** to delete all currently shown jobs that are in a non-working state and all log files.

---

 **Caution:** This cannot be undone.

## Update Displayed Job Information

Click **Refresh** to manually update the displayed job information.

## Exit the Jobs page

Click **Done** to exit the **Jobs** page.

# Files

The **Files Report** provides general information about selected files, as well as specific details if you require more granular information.

## Run the Files Report

1. On the **Reports** menu, click **Files**. The **Reports > Files** report appears. The **Report > Files** page displays the following information about your drives:
  - **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.
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. Press and hold **Shift** and click to select contiguous files, or press and hold **Control** and click to select multiple non-contiguous files.
4. Click **Continue** to proceed and return to the **Reports > Files** page.
5. Click **File Info** to view detailed information for the files you selected. The **File Info** page appears.
6. To download the report, click **Download**.
7. When you finished viewing report information, click **Done**.

---

# Drives

The **Drives Report** provides a list of drives in your system and enables you to view details about selected drives.

## Run the Drives Report

1. On the **Reports** menu, click **Drives**. The **Reports > Drives** report appears. The **Report > Drives** page displays the following information about your drives:

Parameter	Description
<b>Serial Number</b>	The drive's serial number.
<b>State</b>	The current state of the drive, such as <b>Online</b> or <b>Offline</b> .
<b>Device Path</b>	Displays the path of the device.
<b>Status</b>	The drive's current status, such as <b>Free</b> or <b>Mounted</b> .
<b>User Alias</b>	The user identifier for the drive.
<b>Mounted Media</b>	The mounted media number.
<b>Dismount Delay</b>	The interval of time before the drive is dismounted.
<b>Compression</b>	Specifies whether compression is enabled ( <b>True</b> ) or disabled ( <b>False</b> ).

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.
3. To download the report, click **Download**.
4. When you finished viewing report information, click **Done**.

---

## Media

The **Media Report** displays a list of details for all media (including cleaning media) in a selected library or all libraries.

### Run the Media Report

1. On the **Reports** menu, click **Media**. The **Media** report page appears.
2. Select from the **Destinations** list one of these options:

Parameter	Description
<b>Show All Destinations</b>	Select this to view information for all media in all destinations.
<b>Show All Libraries</b>	Select this to view information for all media in all libraries.
<b>Show All Storage Disks</b>	Select this to view information for all media in all storage disks.
<b>Show All Object Storage</b>	Select this to view information for all media in all Object Storage destinations.

3. Select from the **Library** list one of these options:

Parameter	Description
<b>All Libraries</b>	Select this to view information for all media in all libraries.
<b>A selected library</b>	Select a specific library whose media information you want to view.

4. Select from the **Media Class** list one of these options:

Parameter	Description
<b>Show All Media Classes</b>	Display information for all media in all classes
<b>Show Data Media Class</b>	Display only information for media available to be used for data migration that are not considered full.
<b>Show Migrate Media Class</b>	Display only information for media used for data migration which are considered full.
<b>Show Blank Media Class</b>	Display only information for blank media.
<b>Show Clean Media Class</b>	Display only information for cleaning media.
<b>Show Backup Media Class</b>	Display only information for media used for backup purposes.

5. 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.

- The **N/A (Spans Multiple)** filter only applies to Storage Disks and Lattus Media. If a storage disk or Lattus 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).
- 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:

Parameter	Description
<b>Media ID</b>	The unique identifier for the media.
<b>Library</b>	The name of the library in which the media currently resides.
<b>Media Type</b>	The type of media.
<b>Media Format</b>	The format of the media.
<b>Formatted</b>	Displays whether the media is formatted (true if the media is formatted, or false if the media is not formatted).
<b>Status</b>	Displays whether the media is Available or Unavailable.
<b>Media Class</b>	The media class to which the media belongs.
<b>Policy Class</b>	The policy class to which the media belongs.
<b>Mark Status</b>	Displays whether the media is marked or unmarked.
<b>Suspect</b>	Indicates whether the media is considered suspect (possibly unreliable or defective).
<b>Write Protected</b>	Indicates whether the media is write protected.
<b>File Segment Count</b>	The number of files saved on the media.
<b>% Used</b>	Indicates the percentage of the media which is currently used.
<b>Copy</b>	Indicates the policy class copy number on the media.
<b>Mounted in Drive</b>	Indicates whether the media is currently mounted in a drive.
<b>Last Accessed</b>	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, press and hold **Control** 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.
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**.

---

## Q-Cloud Object Storage Media Usage

The **Q-Cloud Object Storage Media Usage** report displays the usage report for object store media.

## Filter Criteria

The data in the table can be filtered based on **Policy** or **Media**. When an option is selected, the corresponding list of available policies or media is populated in the **Select Filter** Items box. You can select one or many from the list. Click **Apply Filter** to apply the filter criteria, or click **Reset** to discard your changes.

## Understand the Q-Cloud Object Storage Media Usage Table Data

The following table provides a description for each of the parameters in the table.

Parameter	Description
<b>Media ID</b>	Displays the Media identifier of the medium.
<b>Class ID</b>	Displays the policy class associated with this media.
<b>Total Object Count</b>	Displays the total number of files written to the medium.
<b>Encrypted Object Count</b>	Displays the total number of encrypted files written to the medium.
<b>PreCompress Size</b>	Displays the total number of bytes before compression written to the medium.
<b>PostCompress Size</b>	Displays the total number of bytes after compression written to the medium.

## Show Total Sum Across All Q-Cloud Object Store Media And/or Policy Classes

When selected, the total sum across all object store media and/or policy classes is shown in the table.

---

# Relations

The **Relations Report** displays the pathname of the managed file system's directory and the corresponding policy class name.

## Run the Relations Report

1. On the **Reports** menu, click **Relations**. The **Reports > Relations** report appears.
2. When you are finished reviewing the report output, click **Done**.

---

# File Systems

The **File Systems Report** provides information about your file systems.

## Run the File Systems Report

1. On the **Reports** menu, click **File Systems**. The **Reports > File Systems** page appears. The **File Systems Report** provides the following information:

Parameter	Description
<b>File System Name</b>	The name of the file system.
<b>Mount Point ("Mounted on")</b>	The mount point location for the file system.
<b>Status</b>	The file system's current status, indicated a green check mark icon ( <b>Active</b> ), a yellow exclamation mark icon ( <b>Warning</b> ), or a red X icon ( <b>Stopped</b> ).
<b>Stripe Group</b>	Name of stripe group.
<b>Content</b>	Content of stripe group, <b>M</b> (Metadata), <b>J</b> (Journal), <b>U</b> (User).
<b>Read</b>	Read access. Either <b>Enabled</b> or <b>Disabled</b> .
<b>Write</b>	Write access. Either <b>Enabled</b> or <b>Disabled</b> .

2. Click **Refresh** to manually update (refresh) the report data.
3. Click **Done** when you are finished viewing the report.

---

# SAN Devices

The **SAN Devices Report** displays a list of details for all currently configured devices attached to your SAN.

## Run the SAN Devices Report

1. On the **Reports** menu, click **SAN Devices**. The **SAN Devices** page appears. The **SAN Devices Report** provides the following information:

Disks and LUNs	Description
<b>Serial Number</b>	The disk's or LUN's serial number or path name.
<b>Type</b>	The device type.
<b>Label</b>	The label, if any, assigned to the device.
<b>Size</b>	The total capacity for the device.
<b>Status</b>	The device's current status. Statuses include
<b>Used</b>	Indicates whether the device is currently in use (true or false).
<b>File System</b>	The name of the file system with which the device is associated.

Libraries and Tapes Drives	Description
<b>Serial Number</b>	The serial number of the library or tape drive.
<b>Product ID</b>	The model number or product name of the library or tape drive.
<b>Device Type</b>	The type of device, Tape Library or Tape Drive.
<b>Device Path</b>	The path name for the device.

2. **(Optional)** Click **Refresh** to manually update (refresh) the report data.

---

## Tape Consolidation

The **Tape Consolidation Report** displays information about the tape consolidation process, also known as defragmentation.

### Run the Tape Consolidation Report

1. On the **Reports** menu, click **Tape Consolidation**. The **Reports > Tape Consolidation** report appears. Enter the following fields which determine report parameters:



Parameter	Description
<b>Max Media to Report</b>	The maximum number of media included in the consolidation process.
<b>No Maximum for Reported Media</b>	Indicate that there is no limit for the number of media included in the report.
<b>Include Unavailable Media</b>	Specify whether to include currently unavailable media in the report.
<b>Fragmentation Summary Only</b>	Specify whether to provide only a high-level summary of consolidation results.

2. Click **Apply** to save and apply the report parameters you just entered.
3. Click **Done** to generate the report as a job. To view the report output, on the **Reports** menu, click **Jobs**, and then select a **Tape Consolidation Report** job.
4. When you are finished review the report output, click **Done**.

---

## SAN and LAN Clients

The SAN and LAN Clients Report provides statistics for StorNext clients, including the number of StorNext SAN clients and LAN Clients, and client performance.

### Run the Clients Report

1. On the **Reports** menu, click **SAN and LAN Clients**. The **Reports > SAN and LAN Clients** page appears.  
The **SAN and LAN Clients Report** provides the following information:

Parameter	Description
<b>File System</b>	The name of the file system supporting the clients.
<b>Mounted on</b>	The name of the file system mount point.
<b>Status</b>	The file system's current status (Normal, Error, or Warning)
<b>SAN Clients</b>	The total number of physically connected StorNext SAN clients, and the IP address of the current client.

Parameter	Description
<b>LAN Clients</b>	The total number of StorNext LAN Clients.
<b>Gateway Server</b>	The total number of Gateway Servers for the file system.
<b>Server</b>	The names of the Gateway Servers.
<b>Listening Interface (IP: Port)</b>	The IP address and port number through which the Gateway Server communicates with StorNext.
<b>TCP Window Size</b>	The TCP window size (in KB) used by the Gateway Server. The default is 64.
<b>Transfer Buffer Size</b>	The transfer buffer size (in KB) used by the Gateway Server. A larger buffer may increase performance for larger files. The default is 256.
<b>Transfer Buffer Count</b>	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. The default is 16.
<b>Server Buffer Count</b>	The number of server buffers used by the Gateway server.
<b>Daemon Threads</b>	The maximum number of daemon threads used by the Gateway Server. The default is 8.

2. Click **Refresh** to manually update (refresh) the report data.
3. Click **Done** when you are finished viewing the report.

---

## LAN Client Performance

The LAN Client Performance Report provides information about LAN Clients, including read and write speed.

### Run the LAN Client Performance Report

1. On the **Reports** menu, click **LAN Client Performance**. The **Reports > LAN Client Performance** page appears. The **LAN Client Performance Report** provides the following information:

Parameter	Description
<b>File System</b>	The name of the file system supporting the clients.
<b>Gateway</b>	The name of the Gateway Server on the indicated file system.
<b>Client</b>	The IP address for the corresponding client interface listed in the Client Interface column.
<b>Client Interface</b>	The name of the LAN Client for the indicated file system and Gateway Server.
<b>Read Bytes/Sec</b>	The number of bytes read by the LAN Client.
<b>Write Bytes/Sec</b>	The number of bytes written by the LAN Client.

2. 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:

Type of Report	Description
<b>Policy Activity</b>	Displays statistics related to replication and deduplication. This report also provides statistics such as savings realized by replication, and the current progress of ongoing replications.
<b>Policy Summary</b>	Displays 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.
<b>Progress Report</b>	Displays the information of an active replication. For example, a replication policy that is currently running. It provides a tabular display of replication stages with various corresponding parameters and a status in percentage.

### Policy Activity

The **Replication/ Deduplication Policy Activity Report** displays deduplication and replication statistics such as the number of bytes deduplicated or replicated, and the percentage of savings realized by deduplication.

## Run the Replication/ Deduplication Policy Activity Report

1. On the **Reports** menu, click **Replication/ Deduplication > Policy Activity**. The **Reports > Replication/ Deduplication > Policy Activity** page appears. The **Replication/ Deduplication Policy Activity Report** provides the following information:

Total Managed Data	Description
<b>Source Directories</b>	The number of source directories from which StorNext checked for data to replicate.
<b>Replicated Bytes</b>	The number of bytes of data replicated.
<b>Replicated Tags</b>	The number of replication tags applied to files.
<b>Skipped Files</b>	The number of files not included in the replication process.
<b>Last Replication Complete</b>	The date and time the last replication was finished.

Deduplication Activity	Description
<b>Deduplicated Files</b>	The number of files deduplicated.
<b>Deduplicated Bytes</b>	The number of bytes of data deduplicated.
<b>Size on Disk</b>	The size (in bytes) of deduplicated data on disk.
<b>Replication Network Traffic Sent</b>	The number of bytes of deduplicated data sent over the network.
<b>Unique Files</b>	The number of unique files included in the deduplication process.
<b>Unique Bytes</b>	The number of unique bytes included in the deduplication process.
<b>Reclaimable Space</b>	The amount of reclaimable space realized by data deduplication.
<b>Replication Network Traffic Received</b>	The number of bytes of deduplicated data received over the network.
<b>Data Savings</b>	The percentage of data savings realized by data deduplication.

2. **(Optional)** Click **Refresh** to manually update (refresh) the report data.
3. To view the Policy Summary Report, click **Policy Summary**.

## Policy Summary

The **Replication/ Deduplication Policy Summary Report** displays 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.

## Run the Replication/ Deduplication Policy Summary Report

1. On the **Reports** menu, click **Replication/ Deduplication > Policy Summary**. The **Reports > Replication/ Deduplication > Policy Summary** page appears. The **Replication/ Deduplication Policy Summary Report** provides the following information:

Parameter	Description
<b>Policy</b>	The name of the replication storage policy.
<b>File System</b>	The name of the file system for which replication is enabled.
<b>Source Directory</b>	The name of the source directory from which information is replicated.

Replicated	Description
<b>Directories</b>	The number of replicated directories to date.
<b>Files</b>	The number of replicated files to date.
<b>Skipped</b>	The number of files skipped by the replication process.
<b>Bytes</b>	The total number of data bytes replicated to date.

Replication	Description
<b>Average Rate</b>	The approximate rate at which data was replicated from the source to the target.
<b>Estimated Completion</b>	The estimated time replication is currently scheduled to complete.
<b>Last Completed</b>	The date and time the last replication was finished.

Replication Network Traffic	Description
<b>Sent</b>	The amount of replicated data sent from the source.
<b>Received</b>	The amount of replicated data received on the target.

2. **(Optional)** Click **Refresh** to manually update (refresh) the report data.
3. To view details for a particular policy, select the desired policy and then click **Details**.
4. To view a report showing completed replication, click **Completion Report**.
5. To view the information of an active replication, click **Progress Report**. For example, a replication policy that is currently running.
6. To view the **Policy Activity** report, click **Policy Activity**.

# Progress Report

The **Replication/Deduplication Progress Report** displays the information of an active replication. For example, a replication policy that is currently running. It provides a tabular display of replication stages with various corresponding parameters and a status in percentage.

## Run the Replication/Deduplication Progress Report

1. On the **Reports** menu, click **Replication/Deduplication**, and then click **Policy Summary**. The **Policy Summary** page appears.
2. Select a replication policy from the list of policies. A spinning icon under the Running column implies that the replication policy is currently active.
3. Click **Progress Report...** to navigate to the progress report page. For details on what is displayed on the progress report page, see [Source Statistics below](#) and [Target Statistics on the next page](#) below.
4. To view a report showing replication policy summary information, click **Policy Summary**.
5. **(Optional)** Click **Refresh** to manually update (refresh) the report data.
6. To display the home page, click **Done**.
7. The progress at each target is displayed in a separate collapsible panel. Click **Minimize All Targets** to minimize all target windows.

## Source Statistics

Parameter	Description
<b>Source</b>	The source from which data is being replicated.
<b>Start Time</b>	The time at which replication started.
<b>Elapsed Time</b>	The amount of time elapsed since the replication started.
<b>Stream ID</b>	A unique ID that represents the replication process.
<b>Status</b>	The status of current replication process at the source.
<b>Replication Types</b>	This parameter displays whether the current replication is a <b>Data Replication</b> or a <b>Namespace Replication</b> .
<b>Stage</b>	This column displays the various stages of replication at the source.
<b>Start Time</b>	The time at which the processing for this stage started.
<b>Elapsed Time</b>	The amount of time elapsed since the processing from this stage started.

Parameter	Description
<b>Details</b>	This column displays the details of this stage. This includes details of various parameters corresponding to this stage.
<b>Status</b>	This column displays a visual representation of progress for this stage in percentage. If the processing for this stage has completed, the status is shown as <b>Completed</b> .

## Target Statistics

Parameter	Description
<b>Target</b>	The target to which data is being replicated.
<b>Start Time</b>	The time at which replication started.
<b>Elapsed Time</b>	The amount of time elapsed since the replication started.
<b>Stream ID</b>	A unique ID that represents the replication process.
<b>Status</b>	The status of current replication process at the target.
<b>Stage</b>	This column displays the various stages of replication at the source.
<b>Start Time</b>	The time at which the processing for this stage started.
<b>Elapsed Time</b>	The amount of time elapsed since the processing from this stage started.
<b>Details</b>	This column displays the details of this stage. This includes details of various parameters corresponding to this stage.
<b>Status</b>	This column displays a visual representation of progress for this stage in percentage. If the processing for this stage has completed, the status is shown as <b>Completed</b> . If the stage is skipped, the status is shown as <b>Skipped</b> .

---

# Distributed Data Mover

The **Distributed Data Mover Report** displays a list of details pertaining to the Distributed Data Mover feature.

## Run the Distributed Data Mover Report

1. On the **Reports** menu, click **Distributed Data Mover**. The **Distributed Data Mover Report** appears. The **Distributed Data Mover Report** provides the following information:

Parameter	Description
<b>Host</b>	The name of the machine on which the Data Mover is running.
<b>Request ID</b>	The identification number for the move-data request.
<b>Device Alias</b>	The alias of the destination device to which data is moved.
<b>Run Time</b>	The time the data movement occurred.
<b>Total Files</b>	The total number of files moved.
<b>Files Copied</b>	The total number of files copied in the move-data process.
<b>Files Failed</b>	The number of files that were not moved during the move-data process.

2. Click **Refresh** to manually update (refresh) the report data.
3. When you are finished viewing report information, click **Done**.

---

## Hardware Status

The **Hardware Status Report** displays 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:** The **Hardware Status** menu option is **only** available on Xcellis, Artico, Pro Foundation, and StorNext Metadata Appliances.

## Run the Hardware Status Report

**Note:** At any time, you can click **Refresh** to manually update the report data.

1. On the **Reports** menu, click **Hardware Status**. The **Hardware Status Report** appears. The first tab, **System Board**, displays the following information about the system board in both nodes:



Parameter	Description
<b>Component</b>	The name of the component on the system board, such as the CMOS battery, IPMI, fans, and so on.
<b>Type</b>	Indicates the category for the component, such as <b>Alarm</b> for the CMOS battery.
<b>Value</b>	If applicable, indicates the current value for the component, such as “27 degrees” for the <b>Ambient Temperature</b> .
<b>Status</b>	Indicates the current state of the component, such as <b>Normal</b> .

- When you are finished viewing system board information, click the **Network Ports** tab. The **Network Ports** page tab displays the following information related to the network ports on both nodes:

Parameter	Description
<b>Component</b>	The name of the port component, such as <b>eth1</b> for Ethernet port 1.
<b>Value</b>	Indicates the correct value for the component, such as <b>1000 Mb/s</b> for the Ethernet port.
<b>Status</b>	Indicates the current state of the component, such as <b>Up</b> or <b>Down</b> .

- When you are finished viewing network ports information, click the **Fibre Channel Ports** tab. The **Fibre Channel Ports** page appears. 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 page displays the following about your storage array:

Parameter	Description
<b>Source</b>	The array node or cluster.
<b>Array</b>	The name of the array.
<b>Type</b>	The array manufacturer.
<b>Status</b>	Indicates the current state of the component, such as <b>Attention</b> or <b>Normal</b> .
<b>Component Group</b>	The array component category, such as trays, controllers, volumes, drives, alarms and so on.

Parameter	Description
<b>Count</b>	The current number of each component group in the array (for example, “2 Volumes”).
<b>Status</b>	The current operational status of the component group, such as <b>Normal</b> . Information for the following headings appears after you select a component group.
<b>Component</b>	The name or location indicator of the selected component group (such as “SAS Phy 1”).
<b>Location</b>	The current location of the selected component group.
<b>Value</b>	The current value of the selected component group, such as <b>Up</b> or <b>Down</b> .
<b>Status</b>	The current operation status of the selected component group, such as <b>Normal</b> .

---

## Gateway Metrics

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.

**i Note:** The **Gateway Metrics Report** only displays Distributed LAN Client data.

The **Gateway Metrics Report** consists of five main sections:

Section	Description
<b>Two graphs at the top of the page</b>	Displays 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.
<b>Gateway Summary</b>	Displays a list of all gateways and statistics for each one.
<b>Client Summary</b>	Displays a list of all clients and statistics for each one.
<b>File System Summary</b>	Displays a list of all file systems and statistics for each one.
<b>Settings</b>	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 adding columns or removing columns as described in [Change Displayed Columns on page 387](#).

---

**i Note:** You can collapse or expand any section by clicking the double arrow icon at the right side of the section title bar.

## Run the Gateway Metrics Report

1. On the **Reports** menu, click **Gateway Metrics**. 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 [Change to Detail Mode below](#).
2. **(Optional)** Click **Refresh** to manually update (refresh) the report data.

You can change the information displayed in the graphs or view detailed information rather than a summary as described in the following sections.

## Change 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.
- **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.

---

**i Note:** On particularly large systems that generate a lot of metric information, if you select 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 displays 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.

## Change to Detail Mode

- 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.

## View 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.

**i Note:** The graphs display a maximum of 20 clients or 20 file systems per gateway. If your gateway has more than 20 clients or 20 file systems, only the first 20 are shown in the graphs.

## Change 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:

Parameter	Description
<b>Gateway (in Gateway section)</b>	The name of the gateway.
<b>Client (in Client section)</b>	The name of the client.
<b>File System (in File System section)</b>	The name of the file system.
<b>Average Overall Throughput</b>	The average overall data throughput including reading and writing, expressed in kilobytes per second.
<b>Read Count</b>	The total number of read (I/O) operations.

Parameter	Description
<b>Read Rate</b>	The number of reads (I/O) per second.
<b>Read Bytes</b>	The number of bytes read.
<b>Average Read Throughput</b>	The average data read throughput.
<b>Read Distribution</b>	The percentage of data read by the specific gateway, client or file system.
<b>Write Count</b>	The total number of write operations.
<b>Write Rate</b>	The number of writes per second.
<b>Write Bytes</b>	The total number of bytes written.
<b>Average Write Throughput</b>	The average write throughput per second.
<b>Write Distribution</b>	The percentage of data written by the specific gateway, client or file system.
<b>Utilization</b>	The currently selected Gateway Network Interface Card's maximum aggregate throughput.

## View and Change 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 page shows the following settings for all gateways:

Parameter	Description
<b>Gateway</b>	The name of the gateway.
<b>Metrics Enabled</b>	Shows whether metrics tracking is enabled ( <b>True</b> ) or disabled ( <b>False</b> ).
<b>Sample Interval (Secs)</b>	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.
<b>Last Connect</b>	The date and time when the last connection was established with the gateway.
<b>Last Update</b>	The date and time when the statistics and graphs on the Gateway Metrics page were last updated. Gateway metrics must be enabled ( <b>True</b> ) for a particular gateway to be included in the update.

**i Note:** You can hide or show columns in the Gateway Metrics Setup section just as you can in the other sections.

## 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.

---

# StorNext Metrics

The StorNext Metrics reports provide performance data logging and visual reporting and graphing features for StorNext systems. The StorNext Metrics reports are a visual reporting tool. This tool combines comprehensive performance data logging with powerful visual reporting and analysis tools to help you identify potential problems and optimize system operations.

This topic contains the following sections:

Section	Description
<b>Introduction to the StorNext Metrics Reports</b>	Provides an overview of the features of StorNext Metrics reports.
<b>StorNext Metrics Navigation</b>	Describes how to access and work with the Web-based user interface of StorNext Metrics reports.
<b>Reports and Graphs</b>	Describes how to view and interpret the available performance reports.
<b>Work With Time Ranges</b>	Describes how to move the time range backward and forward in time, and make the time range longer or shorter
<b>Work With Graphs</b>	Describes how to work with the various types of graphs..
<b>Interpret Performance Data</b>	Describes how to interpret a particular type of performance data.

## Introduction to the StorNext Metrics Reports

## About the StorNext Metrics Reports

The StorNext Metrics reports are a visual reporting tool. This tool combines comprehensive performance data logging with powerful visual reporting and analysis tools to help you identify potential problems and optimize system operation.

## Extension to StorNext Remote Management

With the StorNext Metrics reports, you can view an array of performance statistics for a StorNext system and see how those statistics change over time. This lets you identify trends or determine when a problem began.

By showing you how various operations affect performance, the StorNext Metrics reports also helps you optimize the network ecosystem and business procedures for data management policies, file system disk usage, and replication.

## On Demand Reports

The StorNext Metrics reports continually works in the background to log performance data. To view logged data, use the StorNext Metrics graphical reports. Reports are available on demand through a Web-based interface. You can check up-to-the-minute system status or view data for any time period since data logging began.

The StorNext Metrics reports let you view and work with a wealth of performance and system statistics, such as Ethernet I/O, Fibre Channel I/O, CPU load, File System (FS), and Memory. Each report includes two or more graphs. Use the report tools to zoom in on a graph to see just the time period you want to see, or zoom out to see data for a longer time period.

No matter what time period you select, all of the graphs in the report stay in sync. In addition, the StorNext Metrics reports maintain the current time period when you select a new report. This lets you compare performance data between graphs in the same report or between different reports. For example, you can see how CPU load is affected during deduplication or space reclamation activities.

## Historical Data

StorNext Metrics reports maintain a maximum of 30 days of metrics data in the database. The maximum number of days of metrics data you can view is the last 3 days (click the **Last 3 Days** radio button preset time range). To view metrics data for the last 4 to 30 days, select a specific date using the **Select Time Period** calendar widget.

StorNext Metrics reports' historical record lets you compare current performance to past performance. It also lets you see the effect of any recent changes to system and network configuration or business processes.

## Logging Database

StorNext Metrics reports record performance data in the logging database. The database resides on the StorNext system where StorNext Metrics reports are is running.

## StorNext Metrics Navigation

### Access the StorNext Metrics Reports

To access the StorNext Metrics reports, follow the procedure below:

- From the StorNext GUI, on the **Reports** menu, click **StorNext Metrics**. The StorNext Metrics reports page displays. The user interface is also referred to as the report window.

### Use the Report Window

The report window displays the performance graphs for the currently selected report. When you first access the StorNext Metrics reports, the **Ethernet I/O** report displays.

## Reports and Graphs

### Overview

StorNext Metrics reports maintain the currently selected time range when you select a new report. For example, if you are currently viewing the most recent day of memory usage for the Memory report, StorNext Metrics reports displays data for that same time range when you select a new report. This makes it easy to compare different performance statistics for the same time range.

---

**i Note:** For more information about time ranges, see [Work With Time Ranges on page 397](#).

This section provides information to help you interpret the reports available in StorNext Metrics.

Each report available in StorNext Metrics is made up of two or more graphs. Some graphs appear in more than one report.

The table below lists the graphs included in each report. The reports are designated as (L) for layered graphs or (S) for stacked graphs. This distinction does not apply to graphs that report only a single variable. For information about interpreting each report, see **Report Descriptions**.

Report	Graphs
<b>Ethernet I/O Report</b>	Ethernet Activity (S) ethn Activity
<b>Fibre Channel I/O Report</b>	Fibre Channel Activity (S)FCn Activity
<b>CPU load Report</b>	CPU Load Average CPU stats in % (S)
<b>Memory Report</b>	Memory usage (always base 1024) (S) Swap usage (always base 1024) (S)



Report	Graphs
<b>File System Reports</b>	Space Inodes Connections CPU Load Memory Usage

## Report Descriptions

This section describes the graphs included in the following reports available in StorNext Metrics.

### Ethernet I/O Report

The Ethernet I/O report displays detailed information about the amount of data passing through the Ethernet ports in the system. The report contains the following graphs:

#### Ethernet Activity

The Ethernet Activity (Aggregate Activity) graph displays the amount of data passing through all of the Ethernet ports in the system.

Use the Ethernet Activity graph to monitor writes to and reads from the system using the Ethernet ports. The graph displays each port in a different color.

The symmetry between the four ports indicate the Ethernet ports are bonded (not segmented) and traffic is balanced across the ports.

- Write activity (above the zero line) indicates target replication to the system (if configured), web management activity, heartbeat traffic, and metadata traffic.
- Read activity (below the zero line) indicates source replication from the system (if configured), web management activity, heartbeat traffic, or metadata traffic.

#### ethn Activity

The ethn Activity (Comparative Activity) graph displays the amount of data passing through Ethernet port n. A graph appears for each Ethernet port in the system, for example, **eth0**, **eth1**, **eth2**, and **eth3**.

Use the ethn Activity graph to monitor writes to and reads from the system using Ethernet port n.

- Write activity (above the zero line) indicates target replication to the system (if configured), web management activity, heartbeat traffic, and metadata traffic.
- Read activity (below the zero line) indicates source replication from the system (if configured), web management activity, heartbeat traffic, or metadata traffic.

View the Ethernet I/O report when you need to monitor writes to and reads from the system using the Ethernet ports. For example, if you have replication configured on the system, you should view this report daily to make sure you see replication traffic occurring at the correct times.

If you don't have replication configured, you will see minor heartbeat and web management traffic and then occasional spikes for metadata requests, allocations, and deletes.

To access the Ethernet I/O report, select **Reports > StorNext Metrics > Ethernet I/O**.

## Fibre Channel I/O Report

The Fibre Channel I/O report displays detailed information about the amount of data passing through the Fibre Channel ports in the system. The report contains the following graphs:

### Fibre Channel Activity

The Fibre Channel Activity (Aggregate Activity) graph displays the amount of data passing through all of the Fibre Channel ports in the system.

Use the Fibre Channel Activity graph to monitor writes to and reads from the system using the Fibre Channel ports.

- The graph displays each port in a different color.
- Fibre Channel write activity (above the zero line) occurs during backups.
- A regular backup schedule results in repeating patterns.
- Symmetrical read and write activity (that is, mirrored patterns above and below the zero line) indicate Tertiary Storage Manager (TSM) tape reclamation.

### FCn Activity

The FCn Activity graph displays the amount of data passing through Fibre Channel port n. A graph appears for each Fibre Channel port in the system, for example, port 4, port 5, and so on.

Use the FCn Activity graph to monitor writes to and reads from the system using Fibre Channel port n.

- Fibre Channel write activity (above the zero line) occurs during backups.
- A regular backup schedule results in repeating patterns.

## CPU Load Report

The CPU Load report displays information about the usage of CPU resources in the system. The report contains the following graphs:

To access the CPU load report, select **Reports > StorNext Metrics > CPU load**.

## CPU Load Average

The CPU Load Average graph displays the one minute load average for the system.

Use the CPU Load Average graph to determine if the system has adequate CPU resources.

- The load average represents the average number of processes, in a one minute time period, that were running on a CPU or that were waiting to run on a CPU.
- A load average higher than the number of CPU cores in the system indicates that the system is CPU limited.

For example, StorNext Metadata Appliances have four CPUs. In this case, a load average of greater than four means that some processes had to wait for an available CPU before running. In contrast, a load average of less than four means no processes had to wait for a CPU.

## CPU stats in %

The CPU stats in % graph displays the relative CPU usage for seven categories of processes (see the table below).

Process Category	Description
<b>iowait</b>	The CPU is waiting for an I/O device to respond (for example, the system is waiting on a disk).
<b>irq</b>	The CPU is handling an interrupt request related to I/O (for example, network, Fibre Channel, disk, keyboard, or serial port activity).
<b>softirq</b>	The CPU is handling a high level I/O task (for example, timer interrupts or packets in the TCP/IP stack).
<b>system</b>	The CPU is handling a kernel process (for example, filesystem operations related to the StorNext or blocklet filesystems).
<b>nice</b>	The CPU is handling processes that have lower priority (for example, background processes).

Process Category	Description
<b>user</b>	The CPU is handling processes that are not owned by the kernel (for example, deduplication as well as some space management and replication activities).
<b>idle (not shown)</b>	The CPU is not handling one of the other process categories.

Use the CPU stats in % graph to see how CPU resources are allocated among different categories of processes. The amount of CPU activity consumed by each category of process is expressed as a percentage. The percentages (including the value for idle, which is not shown in the graph) total to 100%.

If a system has a high CPU load average (see CPU Load Average on page 399), then consider the following guidelines:

- A high percentage of system and user activity indicates the system is CPU limited. Add more CPUs to improve system performance.
- A high percentage of iowait activity indicates the system is I/O limited. Add more disks or arrays to improve system performance.

## Memory Report

The Memory report displays information about memory usage. You can view this report to make sure the cache settings are configured to maximize system performance. The report contains the following graphs:

### Memory usage

The Memory usage graph monitors the amount of physical memory (RAM) used during system operation.

- **Note:** For newly configured systems, you will see mostly free memory (light green). As the system increases the number of file systems and clients, monitor this graph often to make sure that there is as little free memory available as possible. This means the cache settings are configured to maximize performance.

In the Memory usage graph:

- Values are in GB. The graph always displays values in base 1024.
- For standard memory (4 KB) pages, the graph displays the amount of memory that is free and used. The graph also displays the amount of memory used for caching. Memory used for caching can be easily freed up, therefore it usually can be treated as available even though it is not free.
- For huge memory (2 MB) pages, the graph displays the amount of memory that is free and used.

## Swap usage

The Swap usage graph monitors the amount of virtual memory used during system operation.

In the Swap usage graph:

- Values are in GB. The graph always displays values in base 1024.
- The graph displays the amount of the disk swap file that is free and used.

---

**i Note:** In the examples, some of the data in the Memory graphs may not be accurate.

To access the Memory report, select **Reports > StorNext Metrics > Memory**.

## File System Reports

To access the **File System** reports, select **Reports > StorNext Metrics > Filesystem**. The **File System** reports provide file system statistics. Each **File System** report displays a single file system server process running on the StorNext MDC Node. The **File System** reports display detailed statistics about the configured file systems in the StorNext MDC. Each **File System** report contains the following graphs:

### Space

The **File System** Space graph displays the number of bytes of data space currently in use or available for use by file system clients. This graph displays you exactly when file system clients started to quickly fill space in the file system.

This graph is a stacked graph, meaning the graph superimposes data for two or more variables on top of each another. A different color is assigned to each variable, so you can see how the values for each variable differ over time.

- **Bytes Free:** Displays the number of bytes of data space available for use by StorNext file system clients. As the level of free data space gets low, expanding the data space for the given file system should be considered.
- **Bytes Used:** Displays the number of bytes of data space currently in use by StorNext file system clients.

### Inodes

The **File System** inodes graph displays the number of inodes that are currently available or are in use by the selected file system.

A large number of inodes in use in the file system indicates that either the file system contains a lot of files or that it contains a lot of fragmented files.

- **Inodes free:** Displays the number of free inodes available. These inodes are available for use in allocating files in the selected StorNext file system. This value should not reach zero unless all of the StorNext metadata space has been consumed. Inodes are allocated dynamically.

- **Inodes used:** Displays the number of inodes allocated. These inodes are in use by the StorNext file system. This graph will show the growth in the number of files used in the StorNext file system.

## Connections

The **File System** connections graph displays the number of StorNext DLC and SAN clients that are connected to the selected file system.

Use this graph to view the history of your client connections. The number of connections will remain steady if the client computers are left running all the time.

- **Proxy Connections:** Displays the number of StorNext DLC clients connected to this StorNext file system.
- **Connections:** Displays the number of StorNext SAN clients connected to this StorNext file system.

## CPU Load

The **File System** CPU Load graph displays the percentage of CPU resources consumed by the **File System** process for the selected file system.

## Memory Usage

The **File System** Memory Usage graph displays the total amount of physical and virtual memory in use by the selected file system.

If the file system's performance degrades, then check this graph for an increase of physical memory (size) in use. This could indicate that runaway processes are doing lots of busy work (thrashing) or that caching is set too high.

A rapid increase in physical memory usage (for example, the file system was using around 200 MB of memory suddenly shot up to 2 GB) may indicate a client or caching issue. If this occurs, you should check the cvlogs to see what happened at the time of the increase.

- **Size:** Indicates the total amount of physical memory in use by the by the file system. High levels of memory consumption by the StorNext FS process could lead to performance if virtual memory swapping once contention for memory resources is seen.
- **Vsize:** Indicates the total amount of virtual memory in use by the file system. The virtual memory size will display the total process footprint for the file system.

## Work With Time Ranges

A time range is like a window through which you view performance data. Each report displays performance data for the time range you choose. All graphs in a report display data for the same time range. By default, StorNext Metrics reports display data for the most recent hour of logging. You can move the time range backward and forward in time, and you can make the time range longer or shorter. When you change the time range, the report automatically adjusts the resolution of performance data. For example, the resolution is finer (more granular) for shorter time ranges and is coarser (less granular) for longer time ranges.

---

**i Note:** No matter how long the time range is, the report scales all graphs in the report so that the time range uses the entire width of each graph.

## Change the Time Range

To view performance data for a different time range, use one of the following methods:

### Use the Timeline

To move the time range forward or backward in time, use the selection handles below a graph. The timeline displays the current time range used for the report.

- To move the time range backward or forward in time, drag the timeline to the left or right.
- To make the time range longer or shorter, drag the left and right selection handles.

### Select a Preset Time Range

To quickly display performance data for a different time range, use the time range presets on the timeline.

To view logged data using a preset time range, perform one of the following:

- Click **Day**, and then select the day from the calendar widget, or enter a date.
- Click **Last 3 Days**.
- Click **Last 24 Hours**.
- Click Last Hour.

When you apply a preset, StorNext Metrics resizes the time range while maintaining the center of the time range.

### Move Forward and Backward

ove the time range forward or backward in time using the navigation buttons on the left or right of the button bar. StorNext Metrics reports shift the time range while maintaining the length of the time range.

Button	Description
< or >	Moves the time range back or forward an amount equal to one quarter of the current time range.

---

**i Note:** The graph displays the starting date and time and the ending date and time of the current time range, as well as the total length of the time range.

## Work With Graphs

In StorNext Metrics, each report is made up of one or more graphs. Each graph displays a particular type of performance data for the current time range. For example, the Memory report includes the following graphs:

**Memory Usage**, and **Swap Usage**. The horizontal axis of each graph represents time and displays the current time range. The vertical axis varies depending on the graph. It is often a capacity or data amount, but can also be a calculated value such as a ratio, average, or percentage. See the following sections for more information about graphs:

## Gaps in Graphs

A white gap in a graph indicates an absence of logging data for a period of time. This can occur for the following reasons:

- A system reboot occurred.
- No StorNext Metrics report logging took place because the system was busy.
- StorNext Metrics reports logging was turned off.

## Stacked and Layered Graphs

StorNext Metrics reports often display data for multiple variables on the same graph. This lets you see the interaction between different variables.

StorNext Metrics reports use two different methods for placing multiple variables on the same graph:

### Layered Graphs

StorNext Metrics reports use layered graphs to compare related variables. A layered graph superimposes data for two or more variables on top of one another. StorNext Metrics reports assign a different color to each variable, so you can see how the values for each variable differ over time.

For example, in the Filesystem Space graph, StorNext Metrics reports display a separate value line for the variables Bytes used and Bytes free.

- 
- i Note:** StorNext Metrics reports always display the smaller variable in front of the larger variable. Because of this, shifts in the color pattern in a graph can occur if the variable that was smaller becomes larger at some point in time.

### Stacked Graphs

StorNext Metrics reports use stacked graphs to display aggregate performance. A stacked graph adds together values for two or more variables to arrive at a total value. StorNext Metrics reports assign a different color to each variable, so you can see the contribution that each variable makes to the total.

For example, in the Ethernet Activity graph, values for each Ethernet port are added together to reach a total value for each point in the time range.

### The Zero Line

StorNext Metrics reports use graphs with a zero line to show when the StorNext MDC is being written to or being read from.



- Positive values (above the line) represent data being written to the StorNext MDC.
- Negative values (below the line) represent data being read from the StorNext system.

By using a zero line, StorNext Metrics reports show data reads and writes on the same graph, for example, on the Ethernet Activity graph.

## Interpret Performance Data

The power of StorNext Metrics reports is that it lets you compare different types of performance data for the same time range. This lets you see patterns and trends and helps you identify relationships between events. Keep in mind the following general concepts as you work with graphs:

## Correlate Information Across Graphs

When you view a report, try to correlate information in one graph with information in the other graphs. Remember that all graphs in a report display the same time range and always remain in sync. That means an event that happens in the center of one graph can be correlated with an event that happens in the center of another graph in the same (or in a different) report. In other words, if you can draw a straight vertical line between events in two graphs, then the events happened at the same time.

## Look For Interactions Between Events

As you work in StorNext Metrics reports, look for interactions between events in different graphs. While correlation is not the same as causation, if you consistently see that events in one graph happen at the same time as events in another graph, there is a strong possibility that the two types of events are related.

## Understand the Effects of Time Resolution

StorNext Metrics reports use aggregation to convert the resolution of the database to the resolution of the graph. This means that, in many cases, each pixel in the graph is an aggregate of multiple data points in the database. Depending on how many data points are aggregated to create each pixel in the graph, the resulting value can change.

The underlying data does not change. The difference in amplitude is due to the different number of data points StorNext Metrics reports aggregate when calculating the value for each pixel in the graph. Be aware of this effect as you work with graphs and time ranges in StorNext Metrics reports.

- 
- i Note:** StorNext Metrics reports use the finest resolution of data available in the database. Finer-grained data is available for more recent time ranges as opposed to time ranges further in the past. This affects the number of data points StorNext Metrics reports aggregate when displaying a graph, and in turn can affect amplitude.



# Chapter 11: Lattus Object Storage

---

StorNext Storage Manager (SNSM) has a storage destination for copies of managed files to go along with Tape and Storage Disk. The storage destination is known as Lattus Object Storage. StorNext also supports both AXR protocol and Simple Storage Service (S3) protocol with Multipart Upload for Object Storage.

**i Note:** The information in this section is applicable for releases up to StorNext 5 release 5.3.x (in other words, prior to StorNext 5 release 5.4.0). Refer to [Object Storage on page 429](#) for the new information, beginning with StorNext 5 release 5.4.0. Also, this section is specific to StorNext Storage Manager Lattus Object 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>.

This chapter contains the following topics:

Audience .....	402
Overview .....	402
Object Storage Media versus Storage Manager Media .....	402
Converting an AXR Namespace to an S3 Bucket .....	403
Object Storage Features in the StorNext GUI .....	404
Configuring Object Storage .....	410
Setting Up Object Storage Destinations .....	414
HTTPS Support for Object Storage .....	421
HTTPS Support for Q-Cloud .....	424
Changes to Existing CLI Commands .....	424

Other Changes and Considerations .....	426
Object Storage Segment Size .....	426

---

## Audience

This topic is targeted at users who will be configuring the new Object Storage 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 Object Storage destination is a set of external object storage devices that can be used for reliable long term data storage. Object Storage devices function at a basic level in the same way as physical tape media or Storage Disk. There is no software limit on the number of Object Storage devices or namespaces that you can add. The actual number of namespaces that can be configured is dependent on expected usage and performance requirements. The devices that make up the Object Storage are referred to as Object Storage devices.

A Object Storage namespace is used like a Storage Disk or a Tape Media. In fact the namespaces are also referred to as Object Storage Media. Since Lattus supports two REST API protocols: AXR and S3, two corresponding Object Storage media types: **Lattus** and **S3** are provided.

There can be many namespaces in a single Object Storage device. A namespace is first created in the Object Storage device, and later, made known to StorNext SM. StorNext SM moves data to a Object Storage 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 Object Storage.

---

## Object Storage Media versus Storage Manager Media

Here are a few comparisons between Object Storage devices and media (namespaces) as compared to other SM media:

- A Object Storage media belongs to no policy class. That is, multiple classes can have data on the same media. (This matches sdisk)

- A Object Storage media belongs to no policy class. That is, multiple classes can have data on the same media. (This matches sdisk)
- A Object Storage media belongs to no policy class. That is, multiple classes can have data on the same media. (This matches sdisk)
  - 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 Object Storage media (namespace), then a Object Storage 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.

---

## Converting an AXR Namespace to an S3 Bucket

With Lattus 3.5.1, it is possible to convert AXR namespaces to S3 buckets and make them accessible via the S3 interface. StorNext 5 provides the capability to convert the media type from AXR to S3.

### Convert One or More AXR Namespaces to S3 Buckets

1. Using the StorNext MDC (CLI only).

**i Note:** Ensure there is no store/retrieve operation occurring on the same AXR namespace which you plan to convert. If there is store/retrieve operation occurring on a particular namespace, wait for the operation to complete.

- a. Execute the command `fsubjcfg` and retrieve the **media ID** of the AXR namespace to be converted.
  - b. Stop TSM.
2. Refer to the documentation in the section titled **Converting an AXR Namespace to an S3 Bucket** located within the *Lattus Service Reference Guide* (Part Number 6-67798-07). Log into the Lattus controller to perform the conversion.
  3. After the Lattus conversion is complete from **Step 2**, execute the following commands using the StorNext MDC CLI:
    - a. Add a S3 connection endpoint via the StorNext GUI or the `fsubjcfg` command line. At least one S3 iopath is needed to access these S3 buckets under the same appliance:

```
#/usr/adic/TSM/exec/fsobjcfg -a -o iopath_alias -i connection_endpoint -e  
http -t S3 -n controller_node_alias
```

- b. Change the AXR namespace name and media type:

```
# fsobjcfg -m -b NameSpace -t S3 -U <S3_bucket_username> -P <S3_bucket_  
password> -X -f <Media-ID>
```

using the NameSpace value from **Step 2** and the Media-ID value from **Step 1a**.

- c. Modify the file `/usr/adic/TSM/config/filesize.config` and change **LATTUS** to **S3** for the **media ID** corresponding to the converted namespace. Alternatively, you can achieve the same result using the StorNext GUI by changing the fspolicy's steering information from Lattus to S3.

---

**i Note:** To use the StorNext GUI, TSM should be available.


- d. Start TSM.
- e. Verify the store and retrieve are working as expected with converted media type.
- f. Use the command `fsfileinfo -u` to verify the object URL reflects (S3) for every file stored before or after the conversion.

---

## Object Storage Features in the StorNext GUI

On the **Home** page of the GUI, the **Object Storage**, **I/O Path** status and **Used Space** indicator are available for monitoring purposes (see [Figure 47 on the next page](#) below).

Figure 47: Home Page



On the **Configuration > Licenses** page of the GUI, the **Object Storage** license row is available for you to access Lattus Object Storage features (see [Figure 48 on the next page](#) below).

Figure 48: Configuration > Licenses Page

The screenshot shows the Quantum StorNext M660 GUI. The top navigation bar includes Configuration, Tools, Service, Reports, and Help. The main content area is titled "Configuration > Licenses".

**Primary System Information:**

- System ID: 3048618D6C
- Serial Number: SV1110BVF30430
- All Licenses Expire At: 2017-10-31 23:59:59 MDT

Status	Feature	Limit	Expires	Description
✓	SAN Client	100	Never	You must have a SAN client license for each client connected to StorNext over a SAN.
✓	LAN Client	200	Never	You must have a LAN Client license for each LAN client you use with StorNext (in addition to any SAN clients).
✓	Storage Manager Subscription (FlexStor)	500 TB	2017-10-31 23:59:59 MDT	A capacity-based subscription license that provides full access to StorNext's Storage Manager functionality for storing and retrieving files on tape, disk, object store platforms or managed cloud services.
✓	Replication	Unlimited	Never	A Replication license is required if you want to use StorNext's Data Replication feature.
✓	Deduplication	600 TB	Never	A Deduplication license is required if you want to use StorNext's Data Deduplication (blockpool) feature.
✓	Vaulting	Unlimited	Never	A Vaulting license provides the ability to move seldom-used media to a manual archive vault, freeing room for media in the managed archives.
✓	Distributed Data Mover (DDM)	300	Never	A license is required for using Distributed Data Mover if you plan to use additional machines besides the primary server to distribute data streams.
✓	Failover (HA)	Unlimited	Never	A Failover (High Availability) license is required if you plan to use StorNext's HA failover features.
✓	Maintenance		2017-09-30 23:59:59 MDT	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.
✓	Gateway Licensed MDC	Unlimited	Never	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.
✓	Certified Disk	700 TB	Never	A Certified Disk license is required to use disks/LUNS which have been certified by Quantum.
✓	Uncertified Disk	800 TB	Never	An Uncertified Disk license is required to use disks/LUNS which have not been certified by Quantum.
✓	Encryption	Unlimited	Never	An Encryption license is required to create encryption keys and use them in storage manager policies for client-side encryption.
✓	NAS	Unlimited	Never	A NAS license enables you to create SMB and NFS shares.
✓	Alternate Store Location	Unlimited	Never	An Alternate Store Location license is required for the automatic copying of files from a main instance of StorNext to a remote instance of StorNext.
✓	Dynamic Application Environment	1	Never	A Dynamic Application Environment license allows the use of guest operating systems on the appliance. Each guest operating system consumes one license count.

Buttons: Import..., Upload DCC..., Download DLS..., Refresh

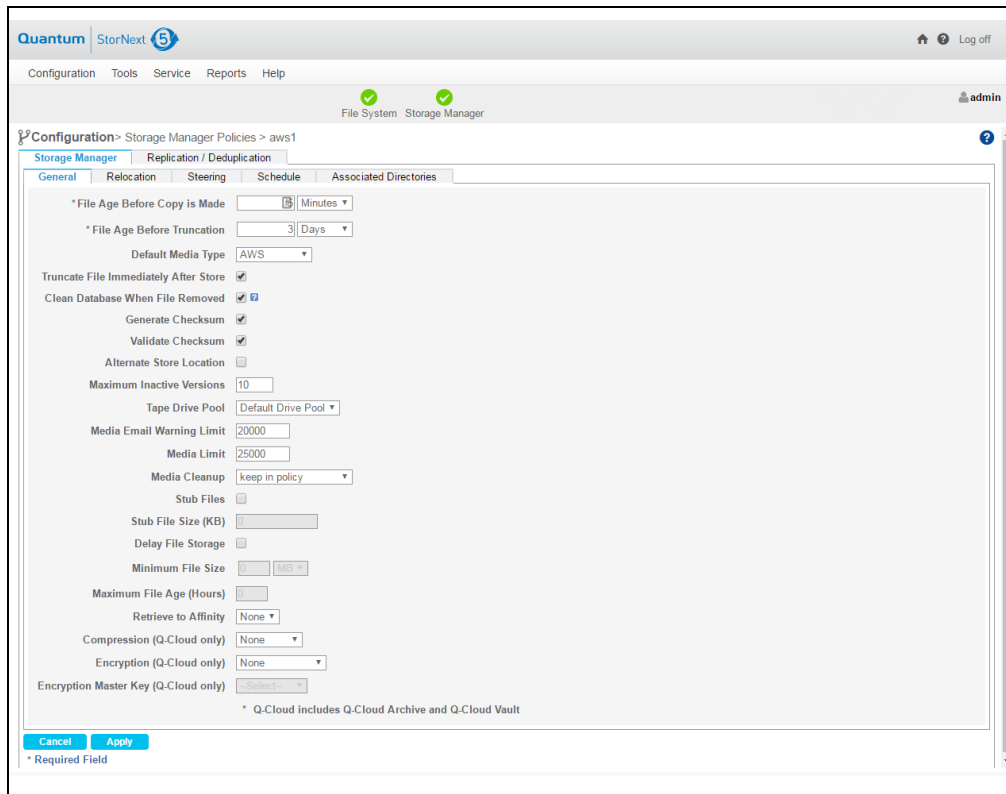
Rows: 16

Connected Licensed Gateways 0

\*CAUTION: Importing a license file overwrites the existing license file. Licenses not present in the imported license file are removed from StorNext.  
\*When importing a license file for HA systems, the imported license file must contain license keys for both nodes.  
\*DLS stands for Disk Licensing Status. The DLS file contains information about the licensing status of the disks.  
\*DCC stands for Disk Certification Catalog. The DCC file contains information about the catalog of Quantum's disk certifications. You can download the latest catalog [here](#).

On the **Configuration > Storage Manager Policies** page of the GUI, support for Object Storage media type when configuring a SM policy is available (see [Figure 49 on the next page](#) below).


Figure 49: Configuration > Storage Manager Policies > Edit Page



On the **Tools > Storage Manager > Distributed Data Mover** page of the GUI, the **Object Storage I/O** path to DDM configuration is available (see [Figure 50 on the next page](#) below).

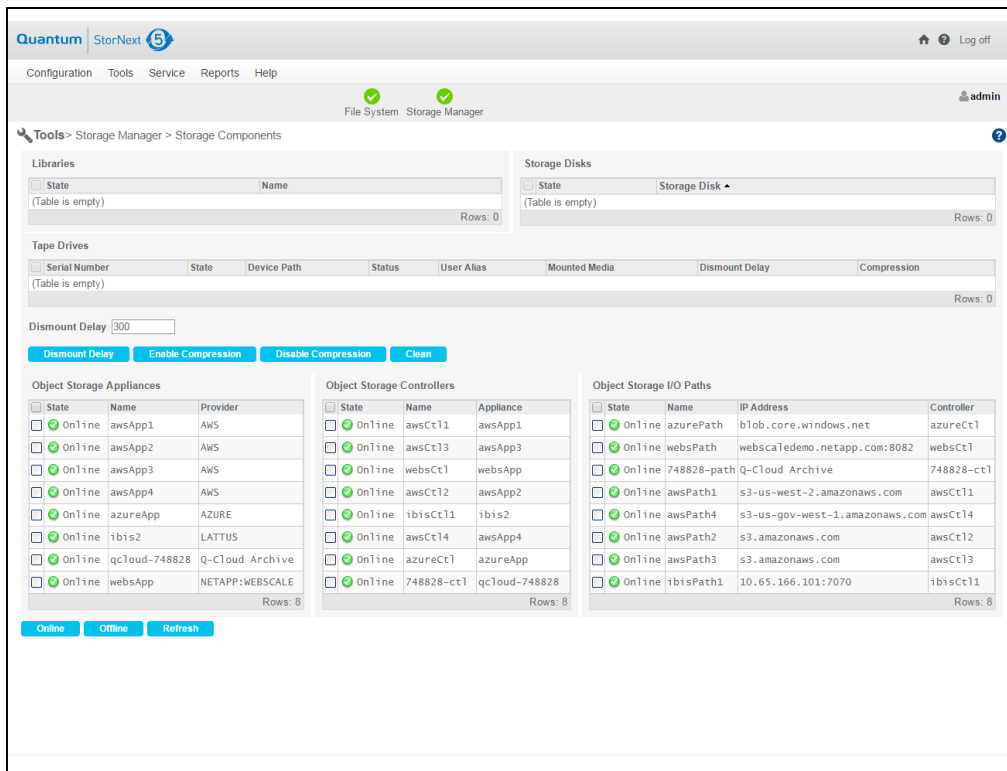


Figure 50: Tools > Storage Manager > Distributed Data Mover Page



On the **Tools > Storage Manager > Storage Components** page of the GUI, you can toggle **Lattus**, **Lattus Controllers**, and **Lattus I/O Paths** online/offline (see [Figure 51 below](#) below).

Figure 51: Tools > Storage Manager > Storage Components Page



On the **Tools > Storage Manager > Media Actions** page of the GUI, the **Object Storage** media type is available for you to remove, purge, assign policy, attributes and clean media operations (see [Figure 52 below](#) and [Figure 53 on the next page below](#)).

**Figure 52: Tools > Storage Manager > Media Actions (Media Selection) Page**

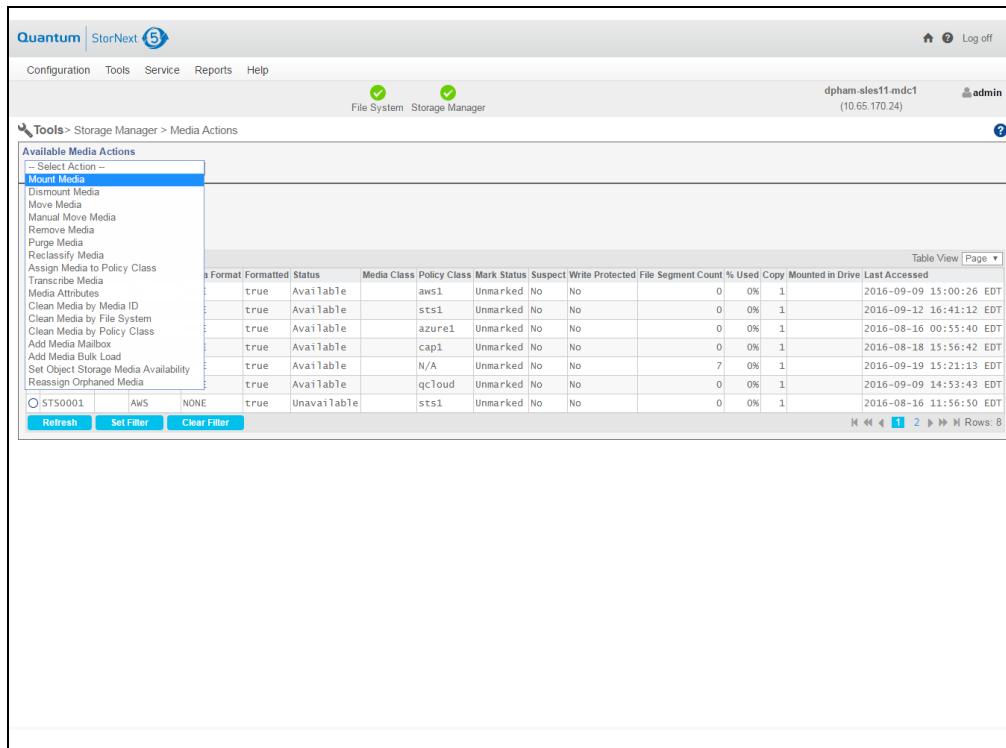
The screenshot displays the 'Media Actions' page in the StorNext GUI. At the top, there is a navigation bar with 'Quantum StorNext 5' and a 'Log off' button. Below this is a menu with 'Configuration', 'Tools', 'Service', and 'Reports'. The main breadcrumb trail reads 'Tools > Storage Manager > Media Actions'. There are two green checkmarks above the breadcrumb, and a user profile 'admin' is visible in the top right.

The page features a dropdown menu for 'Available Media Actions' with the text '-- Select Action --'. Below this is a 'Media ID Filter' input field. The main content is a table titled 'Media' with the following columns: Media ID, Library, Media Type, Media Format, Formatted, Status, Media Class, Policy Class, Mark Status, Suspect, Write Protected, File Segment Count, % Used, Copy, Mounted in Drive, and Last Accessed. The table contains 8 rows of data, each with a radio button in the first column.

Media ID	Library	Media Type	Media Format	Formatted	Status	Media Class	Policy Class	Mark Status	Suspect	Write Protected	File Segment Count	% Used	Copy	Mounted in Drive	Last Accessed
<input type="radio"/> AWS0001		AWS	NONE	true	Available		aws1	Unmarked	No	No	0	0%	1		2016-09-09 15:00:26 EDT
<input type="radio"/> AWS0010		AWS	NONE	true	Available		sts1	Unmarked	No	No	0	0%	1		2016-09-12 16:41:12 EDT
<input type="radio"/> AZUR0001		AZURE	NONE	true	Available		azure1	Unmarked	No	No	0	0%	1		2016-08-16 00:55:40 EDT
<input type="radio"/> CAP0001		AWS	NONE	true	Available		cap1	Unmarked	No	No	0	0%	1		2016-08-18 15:56:42 EDT
<input type="radio"/> LAT0003		S3	NONE	true	Available		N/A	Unmarked	No	No	7	0%	1		2016-09-19 15:21:13 EDT
<input type="radio"/> QCV10003		S3	NONE	true	Available		qcloud	Unmarked	No	No	0	0%	1		2016-09-09 14:53:43 EDT
<input type="radio"/> STS0001		AWS	NONE	true	Unavailable		sts1	Unmarked	No	No	0	0%	1		2016-08-16 11:56:50 EDT

At the bottom of the table, there are buttons for 'Refresh', 'Set Filter', and 'Clear Filter'. On the far right, there is a pagination control showing '1' of '2' pages and 'Rows: 8'.

Figure 53: Tools > Storage Manager > Media Actions (Available Media Actions list) Page



## Configuring 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 Object Storage. The Object Storage devices and media must be configured before they can be used by Storage Manager.

When configuring 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 Object Storage device in Storage Manager, the namespaces, etc. you plan to use must reside in an existing Object Storage appliance prior to configuring in Storage Manager.

## The fsobjcfg Command

The command used for configuring Storage Manager to make use of a Object Storage appliance is **fsobjcfg**. Via this command, the components that make up Storage Manager's view of 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 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 -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 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.

### Route Backup Files to Object Storage Media and into Backup-exclusive Namespaces

1. On the **Configuration** menu, click **Storage Destinations**, and then click the **Object Storage** tab.
2. Click **New...**
  - In the **Namespaces** section, click **Add** to create a Object Storage namespace for each copy number that is required for backup to 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.

**i 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 below](#) 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 Object Storage Destinations

The **Configuration** menu's **Storage Destinations** option enables you to view, add, edit or delete Object Storage destinations.

For information on troubleshooting Object Storage and Cloud errors, see [Debugging StorNext for Object Storage Systems and Cloud Providers on page 638](#).

## View Object Storage Destinations

Follow this procedure to view a list of currently configured Object Storage destinations.

1. Click **Storage Destinations** from the **Configuration** menu.
2. Click the **Object Storage** tab. Information for any previously configured Object Storage destination is shown. For each configured destination, the page 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 Object Storage destination whose information you want to view.
4. Click **View...**
5. When you are finished viewing library information, click **Done**.

## Add a New Object Storage Destination

Follow this procedure to add a new Object Storage destination.

**i 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 291](#). To manage SSL certificates, click **Object Storage Certificates** on the **Tools** menu. For additional information, see [Basic Secure Sockets Layer \(SSL\) Guidelines on page 497](#) in the *StorNext 5 User's Guide*, 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 **Object Storage** tab.
3. Click **New...**
4. Enter the appropriate value into the following parameters:

**i Note:** Parameters marked with an asterisk (\*) are required.

Parameter	Description
*Name	Enter the name of the new Object Storage destination
*Provider	Select <b>LATTUS</b> , <b>AWS</b> , <b>AZURE</b> , <b>ANY:S3COMPAT</b> , <b>IBM:CLEVERSAFE</b> , <b>StorageGRID</b> , <b>Webscale</b> , <b>HGST: ACTIVEARCHIVE</b> , <b>EMC:ECS</b> , <b>HDS:HCP</b> , <b>DDN:WOS</b> , <b>FUJITSU:CD10000</b> , or <b>SCALITY:RING</b> .



Parameter	Description
<b>*Manager Host</b>	Enter the host address for the Object Storage manager host.
<b>*Manager Port</b>	Enter a decimal integer to specify the port number of the Object Storage manager GUI interface. The default port number is <b>80</b> .
<b>Manager Protocol</b>	Select the <b>http</b> or <b>https</b> protocol.  <b>i Note:</b> If you plan to use HTTPS, you must <b>create</b> or <b>import</b> a security certificate prior to creating a <b>Lattus Object Storage Destination</b> . This applies to each Lattus controller configured for HTTPS. If you only intend to use HTTP, certificates are not needed. To <b>create</b> or <b>import</b> a Lattus Object Storage security certificate, see <a href="#">Lattus Certificates on page 291</a> . To manage SSL certificates, click <b>Object Storage Certificates</b> on the <b>Tools</b> menu. For additional information, see <a href="#">Basic Secure Sockets Layer (SSL) Guidelines on page 497</a> in the <i>StorNext 5 User's Guide</i> , as it outlines some standard information about using private and public certificates.
<b>Authentication</b>	Select if authentication is required for this configuration.
<b>User Name</b>	Select a global username to be used for namespace permission, for this configuration. This parameter is mandatory if <b>Authentication</b> is set to " <b>Enabled</b> ".
<b>Password</b>	Select a global password to be used for namespace permissions, for this configuration. This parameter is mandatory if <b>Authentication</b> is set to " <b>Enabled</b> ".

5. In the **Controllers** section, click **Add** or **Add Controller**, and then specify the following criteria:

**i 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 291](#). To manage SSL certificates, click **Object Storage Certificates** on the **Tools** menu. For additional information, see [Basic Secure Sockets Layer \(SSL\) Guidelines on page 497](#) in the *StorNext 5 User's Guide*, as it outlines some standard information about using private and public certificates.

Parameter	Description
<b>Name</b>	Enter the name of the controller.

Parameter	Description
<b>Object Access Protocol</b>	<p>Select the protocol to be used for Object Storage or Cloud Cold Storage object access. By default, the protocol is set to <b>http</b>.</p> <hr/> <p><b>i Note:</b> If you plan to use HTTPS, you must <b>create</b> or <b>import</b> a security certificate prior to creating a <b>Lattus Object Storage Destination</b>. This applies to each Lattus controller configured for HTTPS. If you only intend to use HTTP, certificates are not needed. To <b>create</b> or <b>import</b> a Lattus Object Storage security certificate, see <a href="#">Lattus Certificates on page 291</a>. To manage SSL certificates, click <b>Object Storage Certificates</b> on the <b>Tools</b> menu. For additional information, see <a href="#">Basic Secure Sockets Layer (SSL) Guidelines on page 497</a> in the <i>StorNext 5 User's Guide</i>, as it outlines some standard information about using private and public certificates.</p>
<b>Port(s)</b>	<p>Enter port numbers using a comma-separated string (for example, 8080,8081,8082) to specify the port number(s) of the Object Storage or Cloud Cold Storage manager GUI interface. A maximum of 64 characters, including commas, is allowed.</p> <hr/> <p><b>i Note:</b> For improved performance, it is recommended to define two ports for each controller.</p>
<b>Max Streams</b>	<p>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 <b>Max Streams</b> drop-down list.</p>

6. In the **I/O Paths** section, click **Add** or **Add I/O Path**, and then specify the following criteria:

Parameter	Description
<b>Name</b>	<p>Enter the unique name of the I/O path. If you do not have unique names, the "<b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b>" error message appears.</p>
<b>Controller Name</b>	<p>Select the name of the controller associated with the new I/O path.</p>
<b>Media Type</b>	<p>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.</p> <ul style="list-style-type: none"> <li>For <b>LATTUS</b>, the available options are <b>AXR</b> and <b>S3</b>.</li> <li>For <b>AWS</b>, select <b>AWS</b>.</li> <li>For <b>AZURE</b>, select <b>AZURE</b>.</li> <li>For <b>ANY:S3COMPAT</b>, <b>IBM:CLEVERSAFE</b>, <b>StorageGRID Webscale</b>, <b>HGST:ACTIVEARCHIVE</b>, <b>EMC:ECS</b>, <b>HDS:HCP</b>, <b>DDN:WOS</b>, <b>FUJITSU:CD10000</b>, and <b>SCALITY:RING</b>, the available option is <b>S3COMPAT</b>.</li> </ul>

Parameter	Description
<b>URL Style</b>	<p>There are two ways to format the URL:</p> <ul style="list-style-type: none"> <li>• PATH</li> <li>• VHOST</li> </ul> <p>This parameter defines which style of URL to use.</p>
<b>Object Access Protocol</b>	<p>Select the protocol to be used for Object Storage object access. By default, the protocol is set to <b>http</b>.</p> <p><b>i Note:</b> If you plan to use HTTPS, you must <b>create</b> or <b>import</b> a security certificate prior to creating a <b>Lattus Object Storage Destination</b>. This applies to each Lattus controller configured for HTTPS. If you only intend to use HTTP, certificates are not needed. To <b>create</b> or <b>import</b> a Lattus Object Storage security certificate, see <a href="#">Lattus Certificates on page 291</a>. To manage SSL certificates, click <b>Object Storage Certificates</b> on the <b>Tools</b> menu. For additional information, see <a href="#">Basic Secure Sockets Layer (SSL) Guidelines on page 497</a> in the <i>StorNext 5 User's Guide</i>, as it outlines some standard information about using private and public certificates.</p>
<b>Host [Port]</b>	<p>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.</p>

7. In the **Containers** section, perform one of the following:

- 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**.
- Click **Add** or **Add Container**, and then specify the following criteria:

Parameter	Description
<b>Container</b>	Select ( <b>Scan</b> mode) or enter ( <b>Manual</b> mode) the appropriate container for this configuration.
<b>Media ID</b>	Enter the StorNext Media ID associated with the selected container. The <b>Media ID</b> must be unique.

Parameter	Description
<b>Media Type</b>	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. <ul style="list-style-type: none"><li>• For <b>LATTUS</b>, the available options are <b>AXR</b> and <b>S3</b>.</li><li>• For <b>AWS</b>, select <b>AWS</b>.</li><li>• For <b>AZURE</b>, select <b>AZURE</b>.</li><li>• For <b>ANY:S3COMPAT</b>, <b>IBM:CLEVERSAFE</b>, <b>StorageGRID Webscale</b>, <b>HGST:ACTIVEARCHIVE</b>, <b>EMC:ECS</b>, <b>HDS:HCP</b>, <b>DDN:WOS</b>, <b>FUJITSU:CD10000</b>, and <b>SCALITY:RING</b>, the available option is <b>S3COMPAT</b>.</li></ul>
<b>Authentication</b>	Select if authentication needs be enabled for this container. If the authentication is disabled, the global username and password will be used, if applicable.
<b>User Name</b>	Select a username to be used to access this container. This parameter is mandatory if Authentication is set to “ <b>Enabled</b> ”. This selection overrides the global permissions settings.
<b>Password</b>	Select a password to be used to access the container. This parameter is mandatory if Authentication is set to “ <b>Enabled</b> ”. This selection overrides the global permissions settings.
<b>Copy Number</b>	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.

**i 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 Object Storage destinations.

**i Note:** The containers on Lattus-M 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.

## Edit a Object Storage Destination

Follow this procedure to edit an existing Object Storage destination.

1. If you have not already done so, click **Storage Destinations** from the **Configuration** menu.
2. Click the **Object Storage** tab.
3. Select the 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.

**i 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 Object Storage destination was successfully modified, click **OK**.

## Delete a Object Storage Destination

Follow this procedure to delete an existing Object Storage destination.

1. If you have not already done so, click **Storage Destinations** from the **Configuration** menu.
2. Click the **Object Storage** tab.
3. Select the 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**.

## Perform Other Object Storage Destination Actions

Follow this procedure to launch the Object Storage manager GUI application.

1. If you have not already done so, click **Storage Destinations** from the **Configuration** menu.
2. Click the **Object Storage** tab.
3. Select the Object Storage destination.
4. Click **Launch Manager**. A new browser window appears, and displays the Object Storage manager GUI application login page. If you entered a **User Name** and **Password** when you created the selected Object Storage destination, then the credentials are used as your login.

---

**i Note:** If you are using Safari as your browser, you may have to enable pop-ups. See [Enable Pop-ups in Safari below](#).

## 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.

## Change the Current State of Object Storage Destinations, Controllers, and I/O Paths

You can also change the current state of Object Storage destinations, controllers, and I/O paths. To change the state, select the Object Storage destination, and then choose one of these options from the **Select Action** drop-down list:

Parameter	Description
<b>Online</b>	Select this option to set the Object Storage destination online.
<b>Offline</b>	Select this option to take the Object Storage destination offline.
<b>Controllers Online</b>	Select this option to set the controllers online.
<b>Controllers Offline</b>	Select this option to take the controllers offline.
<b>I/O Paths Online</b>	Select this option to set the I/O Paths online.
<b>I/O Paths Offline</b>	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 ([Change the Current State of Object Storage Destinations, Controllers, and I/O Paths on page 472](#)). 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 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 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.

---

**i 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 on page 292](#) and [Update an Expired CA Root Certificate on page 302](#).

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.

## Configure 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`.

**i Note:** StorNext only supports certificates in `.pem` format. For additional information, see [Basic Secure Sockets Layer \(SSL\) Guidelines on page 497](#), as it outlines some standard information about using private and public certificates.

## Examples

### 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
```

### Update 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
```

### Add 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

Q-Cloud 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 on page 292](#) and [Update an Expired CA Root Certificate on page 302](#).

---

## Changes to Existing CLI Commands

Some existing SM commands had to be updated for use with the new Object Storage destination. Most of the updates to these commands are related to the object IDs that are now assigned to files that are stored to Object Storage. It is the object ID that is used to identify a file segment in the Object Storage. With this object ID and the namespace (determined via the media ID) a file can be accessed directly in the Object Storage. 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.

**i 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 the new options added for reading a file from Object Storage media. See the `fsmedread` man page in the *MAN Pages Reference Guide* for additional details.

## The `fsmedwrite` Command

This command can be used to read a file from disk and write it to Object Storage media. See the `fsmedwrite` man page in the *MAN Pages Reference Guide* for additional details.

## The `fsmedscan` Command

This command was updated to fail if an attempt is made to run against a Object Storage media. Due to the nature of the media and files stored there, no scanning of the media can be done to determine contents.

## The `fsobjcfg` Command

The `-v` option now supports a Q-Cloud device. To specify the device is a Q-Cloud device, use the option `-v QVC1`.

## The `fsaddclass` Command

The `-e` option is added to support Q-Cloud server side encryption. The valid values are either `0` or `1`. `0` means no encryption and `1` means use AES256 encryption.

## 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 Object Storage Object Store, and HTTPS connections.
- StorNext provides support for an object size of 16 TiB (1024<sup>4</sup> bytes).
  - Support for 16 TiB objects (refer to the AmpliStor 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 [Object Storage Segment Size below](#) for additional information.

---

## Object Storage Segment Size

When a file is stored to Lattus 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 Lattus storage.

Segment size should be configured if very large files exist, since the Lattus 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 `objs.conf` applies to all namespaces that are associated with the same Lattus config ID. While the `snpolicy` user-defined policy also has the policy parameter `objs_seg_size`. If `objs_seg_size` is defined in a policy, its value is used as the segment size regardless if `max_seg_size` is defined in Lattus config. This offers the flexibility to specify different segment sizes for different policies on the same Lattus config.

---

**i 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 Lattus 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 Lattus 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 Object Storage 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: Object Storage

---

StorNext Storage Manager (SNSM) now has a new storage destination for copies of managed files, in addition to Tape and Storage Disk destinations. The new destination is known as Object Storage. StorNext 5 release 5.4 supports multiple Object Storage providers and media types, that can be categorized as follows:

Category	Description
Object Storage providers and media types which typically involve the use of <i>on-premises</i> object storage systems that are managed by you.	StorNext offers the following Object Storage providers and media types for this category: <ul style="list-style-type: none"><li>• <b>Quantum Lattus (Lattus Object Storage)</b> supports either the <b>LATTUS AXR</b> protocol configured using the <b>Quantum Lattus</b> media type or the <b>AWS S3</b> compatible protocol configured using the <b>S3</b> media type.</li><li>• <b>IBM Cleversafe, StorageGRID Webscale</b> and <b>Scality RING</b> use the <b>AWS S3</b> compatible protocol and are configured with the <b>S3COMPAT</b> media type.</li></ul>
Object Storage providers and media types which typically involve the use of <i>off-premises</i> object storage systems that you purchase subscriptions for. <b>i Note:</b> Refer to the <b>Storage Manager Subscription (FlexTier™)</b> license type in the <i>StorNext 5 Licensing Guide</i> .	StorNext offers the following Object Storage providers and media types for this category: <ul style="list-style-type: none"><li>• <b>Amazon Web Service (AWS)</b> uses the <b>AWS S3</b> protocol and is configured with the <b>AWS</b> media type.</li><li>• <b>Microsoft Azure (AZURE)</b> uses the Microsoft Azure Blob protocol and is configured with the <b>AZURE</b> media type.</li><li>• <b>QCV1</b> and <b>QVV1</b> which are the Q-Cloud Archive and Vault products described in <a href="#">Chapter 13: Q Cloud</a>.</li></ul>

This chapter contains the following topics:

Audience .....	430
Overview .....	430
Object Storage Media versus Other Storage Manager Media .....	431
Object Storage Features in the StorNext GUI .....	431
System Parameters .....	437
Configuring Object Storage .....	438
Setting Up Object Storage Destinations .....	443
Change the Current State of Object Storage Destinations, Controllers, and I/O Paths .....	472
HTTPS Support for Object Storage .....	473
HTTPS Support for Q-Cloud .....	475
Changes to Existing CLI Commands .....	475
Other Changes and Considerations .....	477
Object Storage Segment Size .....	478

---

## Audience

This section is targeted at users who configure an Object Storage destination or make use of managed files that have been stored to that destination.

---

## Overview

When Storage Manager is managing your data, it creates copies of those managed files on Storage Manager media. When you create a Storage Manager Policy (see [Adding a Storage Manager Policy on page 99](#)), you define the **Default Media Type**.

**i Note:** Devices that make up an Object Storage destination are referred to as *Object Storage devices*, or *Object Storage systems*.

An Object Storage destination is a set of external Object Storage devices used for reliable long term data storage.

At the basic level, the Object Storage devices function in the same way as physical tape media or Storage Disk.

- There is no software limit on the number of Object Storage devices or name-spaces you can add. The actual number of name-spaces configured are dependent on expected usage and performance requirements.
- An Object Storage name-space is used like a Storage Disk or a Tape Media. In fact, the name-space is also referred to as an Object Storage Media, a bucket, or a container depending on the Object Storage providers or protocols.
- There can be many name-spaces in a single Object Storage system. A name-space is first created on the Object Storage system, and later, made known to StorNext Storage Manager. StorNext Storage Manager moves data to an Object Storage name-space for long-term retention in addition to (or instead of) tape and Storage Disk.

---

## Object Storage Media versus Other Storage Manager Media

Below are a few comparisons between Object Storage devices and media (namespaces) as compared to other Storage Manager media:

- An Object Storage media belongs to no policy class as the default. 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 Object Storage media when it is configured. When files are stored, only copies of the indicated number go to a specific media.
  - This copy number can be changed by an administrator when the media is blank.
  - By definition, if files have multiple copies stored to Object Storage media, then an Object Storage media (namespace) per copy number must be configured.
  - If there are multiple media configured for the same copy number, then at store time the media with the most available space is selected for use.
  - If a set of media become full, new ones are configured and assigned whatever copy numbers are needed or desired. This is assuming there is another set of storage nodes with space available.

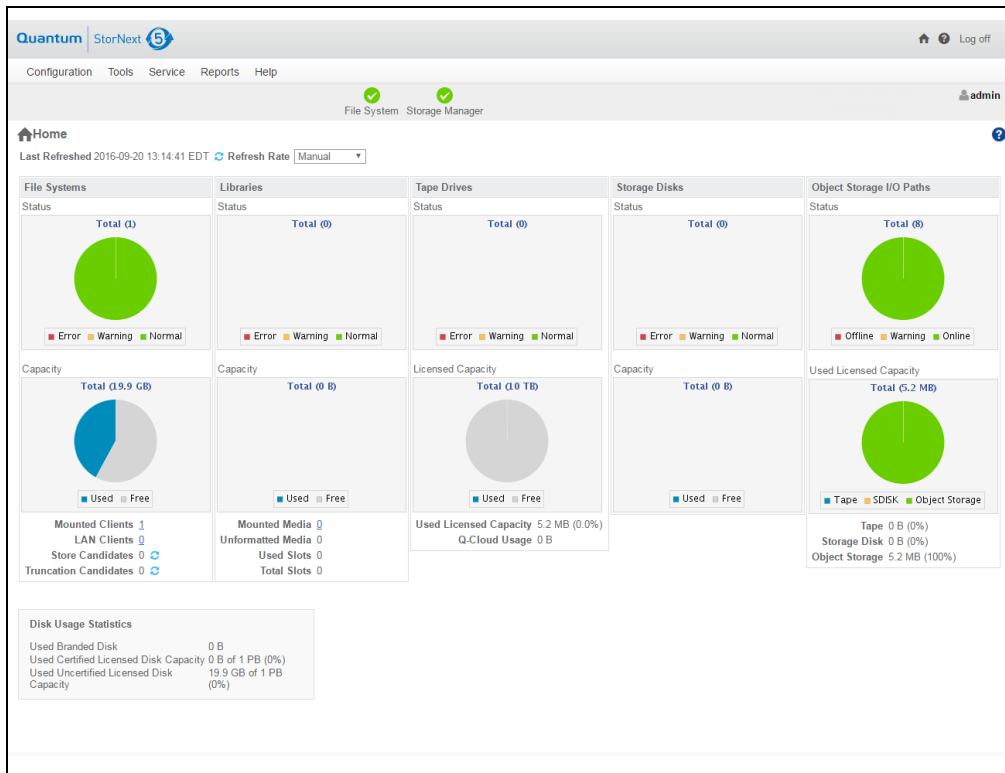
---

## Object Storage Features in the StorNext GUI

On the **Home** page of the GUI, the **Object Storage**, **I/O Path** status and **Used Space** indicator are available for monitoring purposes (see [Figure 54 on the next page](#) below).



Figure 54: Home Page



On the **Configuration > Licenses** page of the GUI, the **Storage Manager Subscription** license row is available for you to access Object Storage features (see [Figure 55 on the next page](#) below).

Figure 55: Configuration > Licenses Page

The screenshot shows the Quantum StorNext M660 GUI. At the top, there are navigation tabs: Configuration, Tools, Service, Reports, and Help. Below these, there are system status indicators: File System (green check), Storage Manager (green check), Tickets (206 open), and Admin Alerts (31). The user is logged in as 'admin' (SV1110BVF30430) with IP 10.20.241.7. The main content area is titled 'Configuration > Licenses' and shows the 'Primary System' information: System ID 3048618D6C, Serial Number SV1110BVF30430, and All Licenses Expire At 2017-10-31 23:59:59 MDT.

Status	Feature	Limit	Expires	Description
✓	SAN Client	100	Never	You must have a SAN client license for each client connected to StorNext over a SAN.
✓	LAN Client	200	Never	You must have a LAN Client license for each LAN client you use with StorNext (in addition to any SAN clients).
✓	Storage Manager Subscription (FlexStor)	500 TB	2017-10-31 23:59:59 MDT	A capacity-based subscription license that provides full access to StorNext's Storage Manager functionality for storing and retrieving files on tape, disk, object store platforms or managed cloud services.
✓	Replication	Unlimited	Never	A Replication license is required if you want to use StorNext's Data Replication feature.
✓	Deduplication	600 TB	Never	A Deduplication license is required if you want to use StorNext's Data Deduplication (blockpool) feature.
✓	Vaulting	Unlimited	Never	A Vaulting license provides the ability to move seldom-used media to a manual archive vault, freeing room for media in the managed archives.
✓	Distributed Data Mover (DDM)	300	Never	A license is required for using Distributed Data Mover if you plan to use additional machines besides the primary server to distribute data streams.
✓	Failover (HA)	Unlimited	Never	A Failover (High Availability) license is required if you plan to use StorNext's HA failover features.
✓	Maintenance		2017-09-30 23:59:59 MDT	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.
✓	Gateway Licensed MDC	Unlimited	Never	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.
✓	Certified Disk	700 TB	Never	A Certified Disk license is required to use disks/LUNS which have been certified by Quantum.
✓	Uncertified Disk	800 TB	Never	An Uncertified Disk license is required to use disks/LUNS which have not been certified by Quantum.
✓	Encryption	Unlimited	Never	An Encryption license is required to create encryption keys and use them in storage manager policies for client-side encryption.
✓	NAS	Unlimited	Never	A NAS license enables you to create SMB and NFS shares.
✓	Alternate Store Location	Unlimited	Never	An Alternate Store Location license is required for the automatic copying of files from a main instance of StorNext to a remote instance of StorNext.
✓	Dynamic Application Environment	1	Never	A Dynamic Application Environment license allows the use of guest operating systems on the appliance. Each guest operating system consumes one license count.

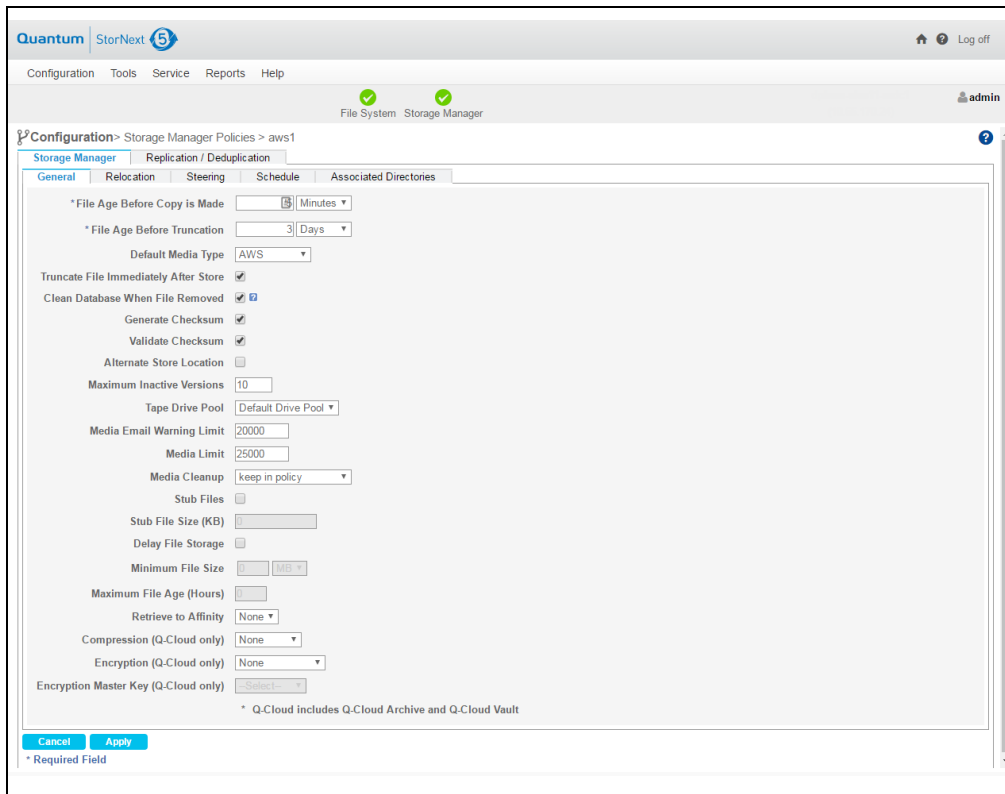
At the bottom of the table, there are buttons for 'Import...', 'Upload DCC...', 'Download DLS...', and 'Refresh'. Below the table, there are several cautionary notes:
 

- \*CAUTION: Importing a license file overwrites the existing license file. Licenses not present in the imported license file are removed from StorNext.
- \*When importing a license file for HA systems, the imported license file must contain license keys for both nodes.
- \*DLS stands for Disk Licensing Status. The DLS file contains information about the licensing status of the disks.
- \*DCC stands for Disk Certification Catalog. The DCC file contains information about the catalog of Quantum's disk certifications. You can download the latest catalog [here](#).

 The bottom right corner of the table area indicates 'Rows: 16'.

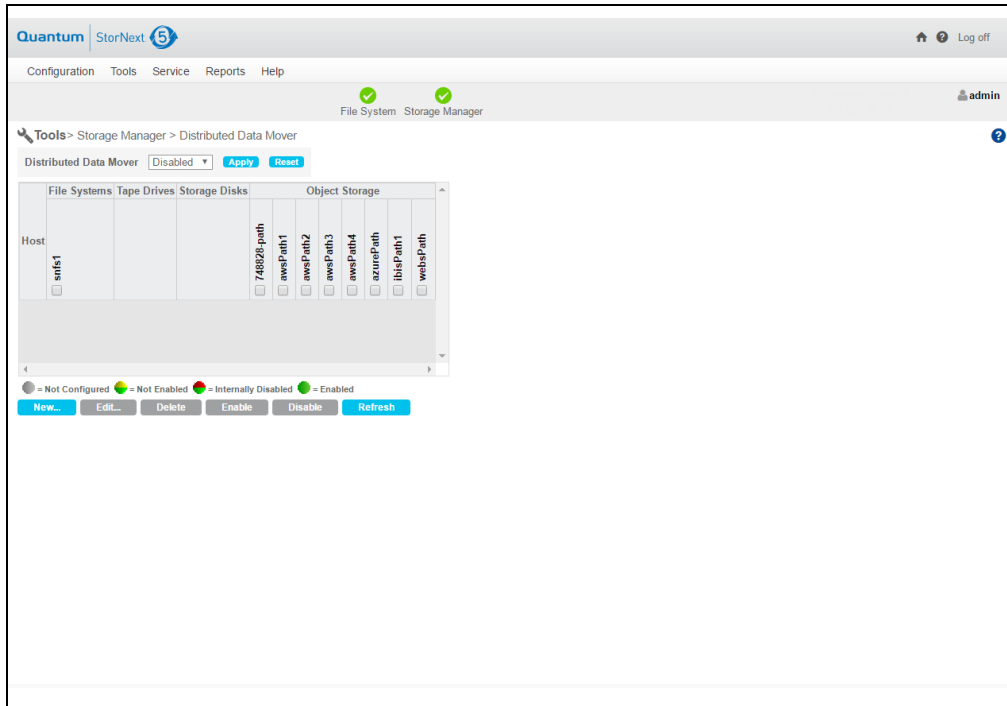
On the **Configuration > Storage Manager Policies** page of the GUI, support for Object Storage media type when configuring a SM policy is available (see [Figure 56 on the next page](#) below).

Figure 56: Configuration > Storage Manager Policies > Edit Page




On the **Tools > Storage Manager > Distributed Data Mover** page of the GUI, the **Object Storage I/O** path to DDM configuration is available (see [Figure 57 on the next page](#) below).

Figure 57: Tools > Storage Manager > Distributed Data Mover Page



On the **Tools > Storage Manager > Storage Components** page of the GUI, you can toggle **Object Storage Appliances, Controllers, and I/O Paths** online/offline (see [Figure 58 below](#) below).

Figure 58: Tools > Storage Manager > Storage Components Page



On the **Tools > Storage Manager > Media Actions** page of the GUI, the **Object Storage** media type is available for you to remove, purge, assign policy, attributes and clean media operations (see [Figure 59 below](#) and [Figure 60 on the next page below](#)).

**Figure 59: Tools > Storage Manager > Media Actions (Media Selection) Page**

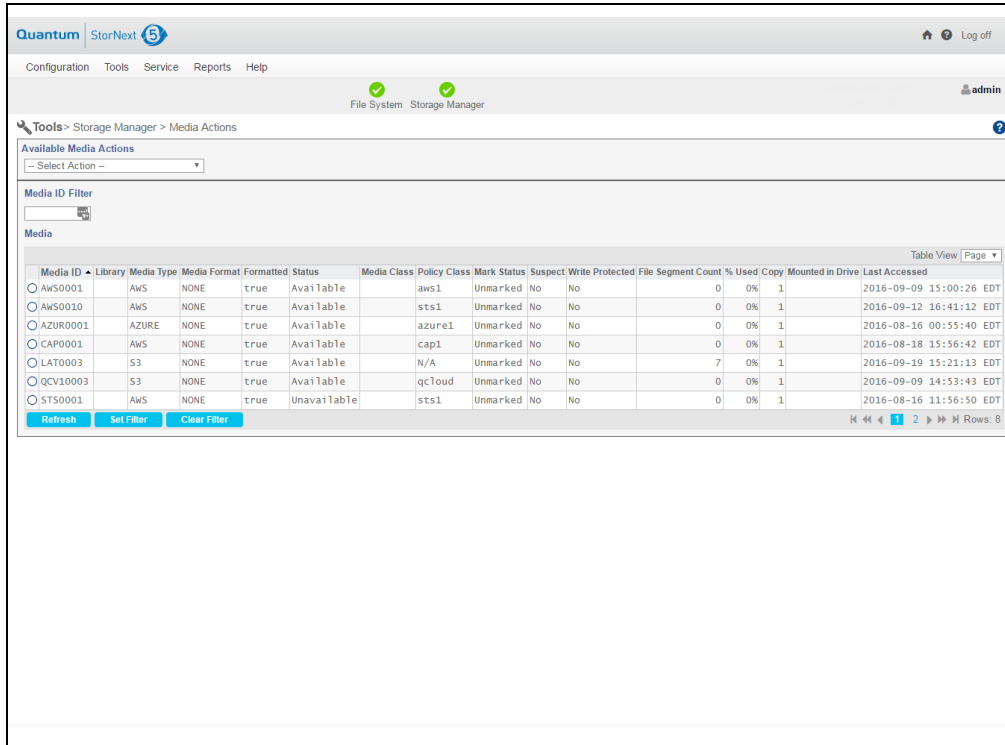
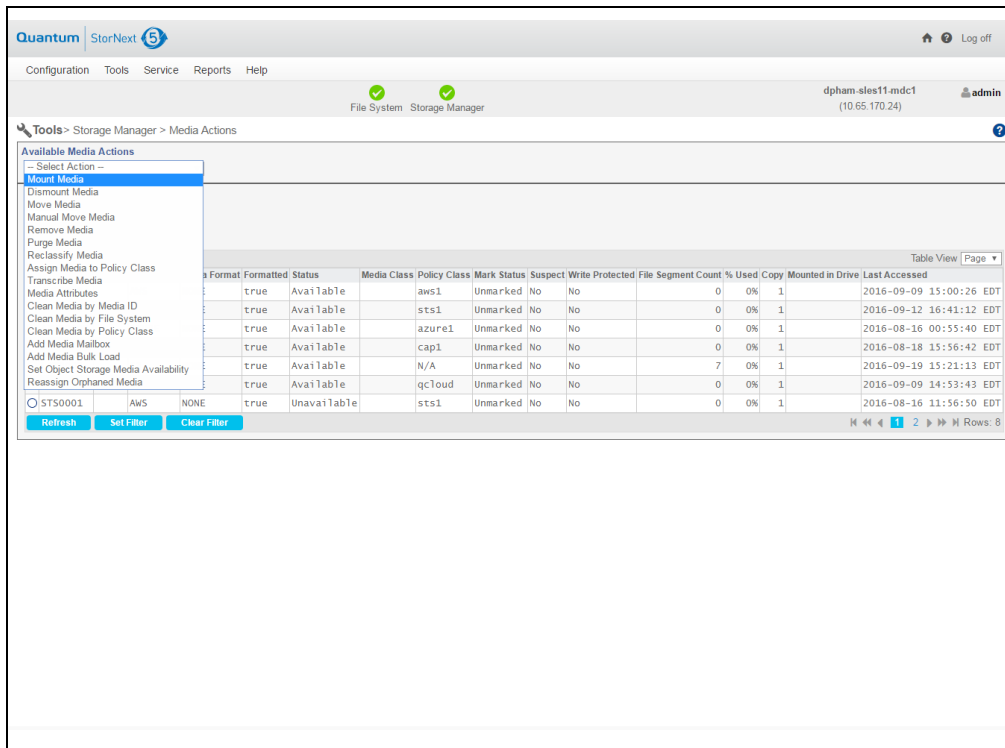


Figure 60: Tools > Storage Manager > Media Actions (Available Media Actions list) Page




## System Parameters

The **Tools > Storage Manager > System Parameters** page allows you to set and modify a sub-set of StorNext system parameters for Object Storage.

### Information on the System Parameters Page

Parameter	Description
<b>Name</b>	Displays the name of the system parameter.  For additional information about system parameters, see the <code>fs_sysparm.README</code> file under the installation directory <code>/usr/adic/TSM/config/</code> .

Parameter	Description
<b>Value</b>	Displays the value of the system parameter. This also serves as a free-form text entry field so that you may update a system parameter value.  <b>Caution:</b> When editing values, ensure that the value you enter for a given system parameter is valid. An invalid value can interfere with proper functioning of Storage Manager. Updating a system parameter value restarts Storage Manager.
<b>Adjusted</b>	Displays <b>Yes</b> , or <b>No</b> , which signifies whether the default Storage Manager system parameter value has been overridden.
<b>Apply</b>	Click to apply the system parameter in the <b>Value</b> field for a given system parameter <b>Name</b> (row).
<b>Reset</b>	Click to clear all the values in the <b>Value</b> field for all system parameters in the table.
<b>Done</b>	Click to exit the <b>System Parameters</b> page.

## Edit or Update a System Parameter Value

1. On the **Tools** menu, click **Storage Manager**, and then click **System Parameters**. The **Tools > Storage Manager > System Parameters** page appears.
2. Select a system parameter under the **Name** column (click a row), and then input a valid system parameter into the text field under the **Value** column. If necessary, click **Reset** to clear all the values in the **Value** field for all system parameters in the table.
3. Click **Apply** to submit your system parameter changes. A confirmation dialog appears.
4. Click **Yes** to confirm your changes, or click **No** to cancel your changes and return to the **System Parameters** page.

## Exit the System Parameters Page

Click **Done** to exit the **System Parameters** page.

---

# Configuring Object Storage

The Object Storage devices and media must be configured before they can be used by Storage Manager. There are two ways to configure Object Storage:

- Use the StorNext GUI, see [Setting Up Object Storage Destinations on page 443](#).
- Use the CLI, see [The fsobjcfg Command on the next page](#).

When configuring 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.

---

**i Note:** Before you create an Object Storage device in Storage Manager, the namespaces/buckets/containers you plan to use must already exist in an Object Storage appliance or system.

## The fsobjcfg Command

The **fsobjcfg** command is used to configure Storage Manager's view of an Object Storage system. The attributes include the host address, TCP/IP port number, network protocol, I/O paths and various namespace attributes and credentials. These together with an object identifier, form the required URL to create/write/read/delete objects in the 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:



```
Object Storage Report
-----
fsobjcfg [-l] [-F text | xml | json]

Object Storage Appliance
-----
fsobjcfg -a -i host_address -p port_number [-e http | https]
          [-U username -P password] [-v provider] [-B]
          appliance_alias
fsobjcfg -m [-i host_address] [-p port_number] [-e http | https]
          [-U username [-P password]] [-B]
          appliance_alias
fsobjcfg -d appliance_alias

Object Storage Controller Node
-----
fsobjcfg -a -n controller_alias [-s streams] [-B] appliance_alias
fsobjcfg -m -s streams [-B] -n controller_alias
fsobjcfg -d -n controller_alias

Object Storage IO Path
-----
fsobjcfg -a -o iopath_alias -i connection_endpoint [-e http | https]
          [-t mediatype] [-u PATH | VHOST] [-B] -n controller_alias
fsobjcfg -m -o iopath_alias [-i connection_endpoint] [-e http | https]
          [-t mediatype] [-u PATH | VHOST] [-B] -n controller_alias
fsobjcfg -d -o iopath_alias -n controller_alias

Object Storage Namespace
-----
fsobjcfg -a -b namespace [-c copy] [-C class] [-f mediaid] [-t mediatype]
          [-U username -P password] [-S signing_type] [-O storageclass]
          [-Y authentication_type] [-B] appliance_alias
fsobjcfg -a -b namespace -t AWS -Y STS_PUBLIC | STS_GOVCLOUD -R role
          -U username -P password [-c copy] [-C class] [-f mediaid]
          [-D role_duration] [-O storageclass] [-S signing_type]
          [-Z authentication_endpoint] [-B] appliance_alias
fsobjcfg -a -b namespace -t AWS -Y CAP -R role -A CAP_agency -M CAP_mission
          [-c copy] [-C class] [-f mediaid] [-D role_duration]
          [-O storageclass] [-S signing_type] [-B] appliance_alias
fsobjcfg -m -f mediaid [-b namespace] [-c copy] [-C class] [-t mediatype]
          [-U username] [-P password] [-O storageclass] [-S signing_type]
          [-Y authentication_type] [-B] [-X]
fsobjcfg -m -f mediaid [-t AWS] [-Y STS_PUBLIC | STS_GOVCLOUD] [-b namespace]
          [-c copy] [-C class] [-U username] [-P password] [-R role]
          [-D role_duration] [-O storageclass] [-S signing_type]
          [-Z authentication_endpoint] [-B]
fsobjcfg -m -f mediaid [-t AWS] [-Y CAP] [-b namespace] [-c copy] [-C class]
          [-A CAP_agency] [-M CAP_mission] [-R role] [-D role_duration]
          [-O storageclass] [-S signing_type] [-B]
fsobjcfg -d -f mediaid
fsobjcfg -r -f mediaid
```

## How to Route Backup Files to Object Storage for Easier Recovery

If you use Object Storage media for system backups, use the procedure to create a special backup namespace and assign it to the backup policy. Use the GUI and follow the procedure below to route backup files to an Object Storage media and into backup-exclusive namespaces/containers.

**Note:** Quantum recommends you create a specific namespace/bucket/container for system backup data.

1. On the **Configuration** menu, click **Storage Destinations**, and then click the **Object Storage** tab.
2. Select an existing Object Storage system that has been previously configured and Click **Edit**.
  - In the **Containers** section, click **Add** to create an Object Storage container for each copy number that is required for backup to Object Storage. Complete all required options to configure the namespace. In the **Policy Class** option, select "**\_adic\_backup**" from the drop down list for the option. The container is now assigned exclusively for backup and not shared with any other SM policies.
3. On the **Configuration** menu, click **Storage 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, select the media type of the container configured specifically for backup in **Step 2**. Also, for the **Media Cleanup** option, select "**keep in policy**" so that the container stays assigned to the backup policy after a complete policy cleanup and the container gets emptied.
  - c. In the **Steering** tab, for each copy to be configured, select the media type of the container which you intend to use.

## How to Route File Copies to a Specific Object Storage Namespace

By default, namespaces can be used by any policy class configured with the same media type as the namespace. If desired, the namespace can be restricted to the exclusive use of a specific policy class by specifying that policy in the namespace configuration. The steps to accomplish this are as follows:

1. On the **Configuration** menu, click **Storage Policies**, and then click the **Storage Manager** tab.
  - a. Click **New** to create a new policy, or click on an existing policy and click **Edit** to edit the policy.
  - b. In the **General** tab, for the **Default Media Type** option, select the media type of the container to be configured specifically for the policy. Also, for the **Media Cleanup** option, select **keep in policy** if the container needs to remain with the policy after it is emptied.
  - c. In the **Steering** tab, for each copy to be configured, select the media type of the container which you intend to use.
2. On the **Configuration** menu, click **Storage Destinations**, and then click the **Object Storage** tab.

3. Select an existing Object Storage system that has been previously configured and click **Edit**.
  - In the **Containers** section, click **Add** to create an Object Storage container for each copy number that is required for the policy configured in **Step 1** for storing file copies to Object Storage. Complete all required options to configure the namespace. When configuring the **Policy Class** option, select the policy configured in **Step 1** from the drop-down list for the option. The container is now assigned exclusively to the policy configured in **Step 1**, and not shared with any other Storage Manager policies.

**i Note:** If the same namespace is to be used for both system (`_adic_backup`) and user backups, do not set the the policy class association.

You can also use the `fsobjcfg` command's option `-C` or the `fschmedstate` command's option `-c` to assign an Object Storage media to a policy class, as long as the media is still blank. See the man pages of these two commands in the *MAN Pages Reference Guide* for additional details.

## The MAX\_STORE\_SIZE System Parameter

The system parameter, `MAX_STORE_SIZE`, set in file `/usr/adic/TSM/config/fs_sysparm_override`, specifies the maximum size of a file that will be recognized as a candidate for automatic storage by `fspolicy`. Storage of files larger than this size can be carried out manually using the `fstore` command.

StorNext provides support for a maximum object size of 16 TiB (bytes) when storing to LATTUS media, 5 TiB when storing to AWS S3 compatible media or Q-Cloud media, and 195 GiB when storing to AZURE media. Since the default value for system parameter `MAX_STORE_SIZE` is set at 2 TB, this system parameter must be reset to a higher value for StorNext `fspolicy` to archive the largest object sizes supported by **LATTUS AXR** and **AWS S3** protocols.

The `MAX_STORE_SIZE` is specified in gigabytes, unless qualified by a suffix specifying one of the multipliers detailed in [Table 4 below](#).

**Table 4:** Suported 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 Object Storage Destinations

Just as with tape and storage disk, you can configure object storage destinations using the **Storage Destinations > Object Storage** option under the **Configuration** menu. See [Configuring Object Storage on page 438](#).

Depending on the provider and media type you are configuring, follow the appropriate procedure below to configure your Object Storage Destination:

- [Setting up Lattus Object Storage Destinations on a StorNext Configuration below](#)
- [Setting up S3COMPAT Object Storage Destinations on a StorNext Configuration on page 450](#)
- [Setting up AWS Object Storage Destinations on a StorNext configuration on page 455](#)
- [Setting up Azure Object Storage Destinations on a StorNext Configuration on page 468](#)

For information on troubleshooting Object Storage and Cloud errors, see [Debugging StorNext for Object Storage Systems and Cloud Providers on page 638](#).

## Setting up Lattus Object Storage Destinations on a StorNext Configuration

### View Lattus Object Storage Destinations

Follow this procedure to view a list of currently configured Lattus Object Storage destinations.

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab. Information for any previously configured Lattus Object Storage destination is shown as entries that have **Quantum Lattus** listed as the **Provider**. For each configured destination, the page displays the **Name**, **Provider**, **Appliance 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 destination information, click **Done**.

### Add a New Lattus Object Storage Destination

Follow this procedure to add a new Lattus Object Storage destination.

**i 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 291](#). To manage SSL certificates, click **Object Storage Certificates** on the **Tools** menu. For additional information, see [Basic Secure Sockets Layer \(SSL\) Guidelines on page 497](#) in the *StorNext 5 User's Guide*, as it outlines some standard information about using private and public certificates.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Click **New....** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue....**
4. Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new Lattus Object Storage destination
<b>Provider</b>	Select <b>Quantum Lattus</b> from the <b>Provider</b> list.
<b>Manager Host</b>	Enter the host address for the Lattus Object Storage manager host.
<b>Manager Port</b>	Enter a decimal integer to specify the port number of the Lattus Object Storage manager GUI interface. The default port number is <b>80</b> .
<b>Manager Protocol</b>	Select the <b>http</b> or <b>https</b> protocol.  <b>i Note:</b> If you plan to use HTTPS, you must <b>create</b> or <b>import</b> a security certificate prior to creating a <b>Lattus Object Storage Destination</b> . This applies to each Lattus controller configured for HTTPS. If you only intend to use HTTP, certificates are not needed. To <b>create</b> or <b>import</b> a Lattus Object Storage security certificate, see <a href="#">Lattus Certificates on page 291</a> . To manage SSL certificates, click <b>Object Storage Certificates</b> on the <b>Tools</b> menu. For additional information, see <a href="#">Basic Secure Sockets Layer (SSL) Guidelines on page 497</a> in the <i>StorNext 5 User's Guide</i> , as it outlines some standard information about using private and public certificates.
<b>Authentication</b>	Select if authentication is required for this configuration.
<b>User Name</b>	Select a global user name to be used for namespace permission for this configuration. This parameter is mandatory if <b>Authentication</b> is set to " <b>Enabled</b> ".
<b>Password</b>	Select a global password to be used for namespace permissions for this configuration. This parameter is mandatory if <b>Authentication</b> is set to " <b>Enabled</b> ".

5. In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

6. In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the " <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> " error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.
<b>Media Type</b> ( <i>New in StorNext 5 release 5.4</i> )	This parameter specifies the object storage media type assigned to an I/O Path that is associated with a specific Object Storage API. The available values for provider <b>Quantum Lattus</b> are <b>AXR</b> and <b>S3</b> .
<b>URL Style</b>	There are two ways to format the URL: <ul style="list-style-type: none"><li>• PATH</li><li>• VHOST</li></ul> This parameter defines which style of URL to use.
<b>Object Access Protocol</b>	Select the protocol to be used for Object Storage object access. By default, the protocol is set to <b>http</b> .  <b>i Note:</b> If you plan to use HTTPS, you must <b>create</b> or <b>import</b> a security certificate prior to creating a <b>Lattus Object Storage Destination</b> . This applies to each Lattus controller configured for HTTPS. If you only intend to use HTTP, certificates are not needed. To <b>create</b> or <b>import</b> a Lattus Object Storage security certificate, see <a href="#">Lattus Certificates on page 291</a> . To manage SSL certificates, click <b>Object Storage Certificates</b> on the <b>Tools</b> menu. For additional information, see <a href="#">Basic Secure Sockets Layer (SSL) Guidelines on page 497</a> in the <i>StorNext 5 User's Guide</i> , as it outlines some standard information about using private and public certificates.
<b>Host</b> <b>[[:Port]]</b>	Enter the connection endpoint address that contains the 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 a user name and password are required, either use the credential specified for the manager host or check the **Use different credentials** box and enter the username 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** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Select ( <b>Scan</b> mode) or enter ( <b>Manual</b> mode) the appropriate container for this configuration.
<b>Media ID</b>	Enter the StorNext Media ID associated with the selected container. The <b>Media ID</b> must be unique.
<b>Media Type</b> <i>(New in StorNext 5 release 5.4)</i>	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. The available values for provider <b>Quantum Lattus</b> are <b>AXR</b> and <b>S3</b> .
<b>Storage Class</b>	This parameter should be left at <b>none</b> as it is not applicable to Lattus media.
<b>Signing Type</b>	For Lattus S3 media, use the default value of <b>V2</b> . This parameter is not applicable to AXR media.
<b>Authentication Type</b>	This parameter specifies the authentication type for the container being configured. An authentication type is required for Lattus S3 media but not for AXR. The available values for provider <b>Quantum Lattus</b> are <b>NONE</b> and <b>STANDARD</b> . The <b>STANDARD</b> type authenticates with a user name and password for Object Storage access.
<b>User Name</b>	Enter a user name to be used to access this container. This parameter is mandatory if Authentication is set to " <b>Enabled</b> ". This selection overrides the global permissions settings.
<b>Password</b>	Enter a password to be used to access the container. This parameter is mandatory if Authentication is set to " <b>Enabled</b> ". This selection overrides the global permissions settings.
<b>Copy Number</b>	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.
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details

**Note:** If no data has been written to a controller, I/O path, or container, you can click **Delete** to

- 
- i** remove the item, and then click **Apply** to save the changes
  - 8. **(Optional)** Repeat **Step 7** to add additional containers to the same Lattus Object Storage Destination.
  - 9. Click **Apply** to save your changes, or **Cancel** to exit without saving.
  - 10. **(Optional)** Repeat **Step 3** through **Step 9** to add additional Lattus Object Storage destinations.
- 

**i Note:** The containers on Lattus-M share the same I/O paths and storage capacity. There is no advantage to be gained by defining multiple containers for the same policy class and copy number. Storage Manager selects the first available container that meets the policy class criteria for the store operation.

## Edit a Lattus Object Storage Destination

Follow this procedure to edit an existing Lattus Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **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.

---

**i 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**.

## Delete a Lattus Object Storage Destination

Follow this procedure to delete an existing Lattus Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **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 Object Storage destination was successfully deleted, click **OK**.



## Perform 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, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the Lattus Object Storage destination.
4. Click **Launch Manager**. A new browser window appears, and displays the Object Storage Manager GUI application login page. 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.

---

**i Note:** If you are using Safari as your browser, you may have to enable pop-ups. See [Enable Pop-ups in Safari below](#).

## 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.

## 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 ([Change the Current State of Object Storage Destinations, Controllers, and I/O Paths on page 421](#)). 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.

## Converting an AXR Namespace to an S3 Bucket

With Lattus 3.5.1, it is possible to convert AXR namespaces to S3 buckets and make them accessible via the S3 interface. StorNext 5 provides the capability to convert the media type from AXR to S3.

### Convert One or More AXR Namespaces to S3 Buckets

1. Using the StorNext MDC (CLI only).

---

**i Note:** Ensure there is no store/retrieve operation occurring on the same AXR namespace which you plan to convert. If there is store/retrieve operation occurring on a particular namespace, wait for the operation to complete.

- a. Execute the command `fsobjcfg` and retrieve the **media ID** of the AXR namespace to be converted.
  - b. Stop TSM.
2. Refer to the documentation in the section titled **Converting an AXR Namespace to an S3 Bucket** located within the *Lattus Service Reference Guide* (Part Number 6-67798-07). Log into the Lattus controller to perform the conversion.
  3. After the Lattus conversion is complete from **Step 2**, execute the following commands using the StorNext MDC CLI:

```
#/usr/adic/TSM/exec/fsobjcfg -a -o iopath_alias -i connection_endpoint -e  
http -t S3 -n controller_node_alias
```

- b. Change the AXR namespace name and media type:

```
# fsobjcfg -m -b NameSpace -t S3 -U <S3_bucket_username> -P <S3_bucket_  
password> -X -f <Media-ID>
```

using the NameSpace value from **Step 2** and the Media-ID value from **Step 1a**.

- c. Modify the file `/usr/adic/TSM/config/filesize.config` and change **LATTUS** to **S3** for the **media ID** corresponding to the converted namespace. Alternatively, you can achieve the same result using the StorNext GUI by changing the fspolicy's steering information from Lattus to S3.

---

**i Note:** To use the StorNext GUI, TSM should be available.

- d. Start TSM.
- e. Verify the store and retrieve are working as expected with converted media type.
- f. Use the command `fsfileinfo -u` to verify the object URL reflects (S3) for every file stored before or after the conversion.

## Setting up S3COMPAT Object Storage Destinations on a StorNext Configuration

Storage Manager supports a rich set of features for AWS S3 compatible cloud storage providers including:

- HTTP and HTTPS access to AWS S3 compatible buckets
- AWS Signature Version 2 (V2) and Version 4 (V4)
- AWS Standard authentication that makes use of either your AWS Identity and Access Management (IAM) Access Key Id and Secret Access Key or your user name and password

The following S3 compatible providers are supported by Storage Manager for the S3COMPAT media type:

- **IBM Cleversafe**
- **StorageGRID Webscale**
- **Scality RING**

When you configure the Object Storage Destination for one of these providers, select the provider that matches your actual Object Storage system.

### View S3COMPAT Object Storage Destinations

Follow this procedure to view a list of currently configured S3COMPAT Object Storage destinations.

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab. Providers using the **S3COMPAT** media type are listed along with other object storage providers in the **Object Storage** tab. For each configured destination, the screen displays the **Name**, **Provider**, **Appliance 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 **S3COMPAT** Object Storage destination whose information you want to view.
4. Click **View...**
5. When you are finished viewing the destination information, click **Done**.

### Add a New S3COMPAT Object Storage Destination

Follow this procedure to add a new **S3COMPAT** Object Storage destination.

1. Ensure that the S3 buckets that you are configuring with Storage Manager, have been created on your Object Storage system and that you know the names, connection endpoint, Access Key Id, and Secret Access Key to your buckets.

**i Note:** Note: If you plan to use **HTTPS**, you may have to create or import a security certificate prior to creating a **S3COMPAT** Object Storage destination. Follow the documentation from your S3 Object Storage system's vendor to set this up on your Object Storage system. Also, see [HTTPS Default CA ROOT Certificate File or Path on page 292](#) for additional information on how to configure your customized CA PEM files on your StorNext system.

2. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
3. Click the **Object Storage** tab.
4. Click **New...** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue...**
5. Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new S3COMPAT Object Storage destination
<b>Provider</b>	Select one of the <b>S3COMPAT</b> providers from the drop-down list that matches your Object Storage system. <ul style="list-style-type: none"><li>• <b>IBM Cleversafe</b></li><li>• <b>StorageGRID Webscale</b></li><li>• <b>Scality RING</b></li></ul>
<b>Manager Host</b>	Enter the host address for the <b>S3COMPAT</b> Object Storage manager host.
<b>Manager Port</b>	Enter a decimal integer to specify the port number of the <b>S3COMPAT</b> Object Storage Host's port. The default port number is <b>80</b> .
<b>Manager Protocol</b>	Select the <b>HTTP</b> or <b>HTTPS</b> protocol. <p><b>i Note:</b> Note: If you plan to use <b>HTTPS</b>, you may have to create or import a security certificate prior to creating a <b>S3COMPAT</b> Object Storage destination. Follow the documentation from your S3 Object Storage system's vendor to set this up on your Object Storage system. Also, see <a href="#">HTTPS Default CA ROOT Certificate File or Path on page 292</a> for additional information on how to configure your customized CA PEM files on your StorNext system.</p>
<b>Authentication</b>	This parameter <b>does not</b> apply to the <b>S3COMPAT</b> Object Storage Destination. Leave it at the default <b>Disabled</b> .

Parameter	Description
<b>User Name</b>	This parameter <b>does not</b> apply to the <b>S3COMPAT</b> Object Storage Destination. Leave it at the default blank.
<b>Password</b>	This parameter <b>does not</b> apply to the <b>S3COMPAT</b> Object Storage Destination. Leave it at the default blank.

6. In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

7. In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

**i Note:** Using the Host name to configure the **I/O Path** for the provider **Scality RING**.

If you are not using an IP address as an endpoint to Scality, the default Host name endpoint for Scality is to emulate Amazon AWS S3 (for example, **s3.amazonaws.com**). This configuration is provided in the `config.json` file on the Scality host. Using any other host name as an endpoint (for example, **my-scality.host.com**) does not work even if the name resolves to the correct IP address. Scality rejects the request with the error message `HTTP/1.1 400 Bad Request`. If you want to use the hostname configured by default on Scality, you can configure your server to use **s3.amazonaws.com** as the endpoint. However, ensure that the name resolves to the IP address of your Scality host.

If your Scality host's IP address is `10.65.191.2`, you can resolve **s3.amazonaws.com** by having the following entry in your `/etc/hosts` file:

```
10.65.191.2 s3.amazonaws.com
```

If you want to use a DNS name that is not the default on Scality, modify the `config.json` file on the scality host:

```
# cat config.json | grep my-scality  
"localregion": ["my-scality.host.com"]
```

And have an entry in your `/etc/hosts` file that resolves it correctly:

```
10.65.191.2 my-scality.host.com
```

Quantum recommends you consult your Scality vendor for further guidance.

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the " <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> " error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.
<b>Media Type</b> ( <i>New in StorNext 5 release 5.4</i> )	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. Select <b>S3COMPAT</b> from the drop-down list.
<b>URL Style</b>	There are two ways to format the URL: <ul style="list-style-type: none"><li>• PATH</li><li>• VHOST</li></ul> This parameter defines which style of URL to use.
<b>Object Access Protocol</b>	Select the protocol to be used for <b>S3COMPAT</b> Object Storage object access. By default, the protocol is set to <b>http</b> .
<b>Host [:Port]</b>	Enter the connection endpoint address that contains the 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 (for additional information, see <a href="#">Using the Host name to configure the I/O Path for the provider Scality RING. on the previous page</a> )

8. In the **Containers** section, perform one of the following:

- 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**, check the **Use different credentials** box and enter the username (the **Access Key ID**) and password (the **Secret Access Key**). You are then presented with a pre-populated list of available containers when you add a container. 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**.
- Click **Add** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Select ( <b>Scan</b> mode), or enter ( <b>Manual</b> mode) the name of your S3 bucket.

Parameter	Description
<b>Media ID</b>	Enter the StorNext <b>Media ID</b> associated with the selected container. The <b>Media ID</b> must be unique.
<b>Media Type</b>	Select <b>S3COMPAT</b> from the drop-down list.
<b>Storage Class</b>	This parameter should be left at <b>none</b> as it is not applicable to <b>S3COMPAT</b> media
<b>Signing Type</b>	This parameter specifies the signing type for the requests sent to the <b>S3COMPAT</b> Object Storage server. Available values include <b>V2</b> and <b>V4</b> . In order to use <b>V4</b> , the server must support <b>V4</b> signing for both AWS full payload and chunked uploading. <b>i Note:</b> When configuring the signing type for containers from the provider <b>Scality RING</b> , set this parameter at <b>V2</b> , until Scality supports <b>V4</b> chunked uploading.
<b>Authentication Type</b>	This parameter specifies the authentication type for the container being configured. An authentication type is required for all <b>S3COMPAT</b> media. Use the default value of <b>STANDARD</b> , which authenticates with an <b>Access Key ID</b> and <b>Secret Access Key</b> .
<b>User Name</b>	Enter the <b>Access Key Id</b> for this container.
<b>Password</b>	Enter the <b>Secret Access Key</b> for this container.
<b>Copy Number</b>	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.
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details.

**i Note:** If no data has been written to a controller, I/O path, or container, click **Delete** to remove the item, and then click **Apply** to save the changes.

9. **(Optional)** Repeat **Step 8** to add more containers to the same **S3COMPAT** Object Storage Destination.
10. Click **Apply** to save your changes, or **Cancel** to exit without saving.
11. **(Optional)** Repeat **Step 4** through **Step 10** to add additional **S3COMPAT** Object Storage Destinations.

## Edit a S3COMPAT Object Storage Destination

Follow this procedure to edit an existing **S3COMPAT** Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.

3. Select the **S3COMPAT** 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 **S3COMPAT** Object Storage destination was successfully modified, click **OK**.

## Delete a S3COMPAT Object Storage Destination

Follow this procedure to delete an existing **S3COMPAT** Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the **S3COMPAT** 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 Object Storage destination was successfully deleted, click **OK**.

## Setting up AWS Object Storage Destinations on a StorNext configuration

Storage Manager supports a rich set of features for Amazon Web Service (AWS) Simple Storage Service (S3) Cloud Storage including:

- HTTP and HTTPS access to AWS S3 buckets, though HTTPS is the recommended protocol for access AWS S3 buckets
- AWS Signature Version 2 (V2) and Version 4 (V4)
- Multiple AWS authentication types including:
  - AWS Standard authentication that makes use of your AWS Identity and Access Management (IAM) Access Key Id and Secret Access Key, for AWS public cloud and GovCloud
  - AWS Security Token Service (STS) authentication for AWS public cloud and GovCloud
  - AWS Commercial Cloud Service (C2S) Access Portal (CAP) authentication for AWS FedCloud
- Multiple AWS storage classes including **standard**, **standard\_ia**, and **glacier**



## View AWS Object Storage Destinations

Follow this procedure to view a list of currently configured **AWS** Object Storage destinations.

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab. Information for any previously configured **AWS** Object Storage destinations is shown as entries that have **AWS** listed as the **Provider**. For each configured destination, the screen displays the **Name**, **Provider**, **Appliance 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 **AWS** Object Storage destination whose information you want to view.
4. Click **View...**
5. When you are finished viewing the destination information, click **Done**.

## Add a New AWS Object Storage Destination that uses STANDARD Authentication

Prior to configuring the AWS Object Storage Destinations with Storage Manager for AWS STANDARD authentication, complete the following two steps, which apply to both AWS public cloud and GovCloud:

1. Have your AWS IAM Access Key Id and Secret Access Key ready.
2. Ensure that the S3 buckets that you are configuring with Storage Manager have been created with AWS and that you know the names and the AWS region endpoint of your buckets.

See <http://docs.aws.amazon.com/general/latest/gr/rande.html> for details on AWS Regions and Endpoints for AWS public cloud.  
See <http://docs.aws.amazon.com/govcloud-us/latest/UserGuide/using-govcloud-endpoints.html> for details on AWS GovCloud (US) Endpoints.

Follow this procedure to add a new **AWS** Object Storage destination that uses STANDARD authentication.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Click **New...** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue...**
4. Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new <b>AWS</b> Object Storage destination
<b>Provider</b>	Select <b>AWS</b> from the <b>Provider</b> list.

Parameter	Description
<b>Manager Host</b>	Enter <b>s3.amazonaws.com</b> for the <b>AWS</b> Object Storage Manager Host's address.
<b>Manager Port</b>	Enter <b>443</b> for the <b>AWS</b> Object Storage Manager Host's port.
<b>Manager Protocol</b>	Select <b>HTTPS</b> .
<b>Authentication Type</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default <b>Disabled</b> .
<b>User Name</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.
<b>Password</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.

5. In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

6. In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the " <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> " error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.
<b>Media Type</b> <i>(New in StorNext 5 release 5.4)</i>	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. Select <b>AWS</b> from the drop-down list.
<b>URL Style</b>	Select <b>VHOST</b> from the drop-down list.
<b>Object Access Protocol</b>	This specifies the network protocol to be used for the host. Select <b>HTTPS</b> from the drop-down list.

Parameter	Description
<b>Host[:Port]</b>	Enter the AWS region endpoint for your S3 bucket. For example: <ul style="list-style-type: none"><li>• <b>s3-us-west-2.amazonaws.com</b> for an S3 bucket created in the US West (Oregon) region</li><li>• <b>s3.amazonaws.com</b> for an S3 bucket created in the US East (N. Virginia) region</li><li>• <b>s3-us-gov-west-1.amazonaws.com</b> for an S3 bucket created in AWS GovCloud</li></ul>

7. In the **Containers** section, perform the following:
  - a. Leave the **Container Selection** at the default **Manual**.
  - b. Click **Add** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Enter the name of your S3 bucket.
<b>Media ID</b>	Enter the StorNext <b>Media ID</b> associated with the selected container. The <b>Media ID</b> must be unique.
<b>Media Type</b>	Select <b>AWS</b> from the drop-down list.
<b>Storage Class</b>	This parameter specifies the storage class of your S3 bucket's lifecycle property. Make sure that this value always matches that configured for the S3 bucket itself. Available values include <b>standard</b> , <b>standard_ia</b> , and <b>glacier</b> . Most configurations use the default <b>standard</b> storage class.
<b>Signing Type</b>	This parameter specifies the signing type for the requests sent to the AWS Object Storage server. Available values include <b>V2</b> and <b>V4</b> (default).
<b>Authentication Type</b>	This parameter specifies the authentication type for the container being configured. Available values include: <ul style="list-style-type: none"><li>• <b>STANDARD</b></li><li>• <b>STS_PUBLIC</b></li><li>• <b>STS_GOVCLOUD</b></li><li>• <b>CAP</b></li></ul> Select the default <b>STANDARD</b> since you are configuring the container to use <b>STANDARD</b> authentication.
<b>User Name</b>	Enter the <b>Access Key Id</b> for this container.
<b>Password</b>	Enter the <b>Secret Access Key</b> for this container.
<b>Copy Number</b>	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.

Parameter	Description
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details.

- i Note:** If no data has been written to a controller, I/O path, or container, click **Delete** to remove the item, and then click **Apply** to save the changes.
- 8. **(Optional)** Repeat **Step 7** to add more containers to the same **AWS** Object Storage Destination.
- 9. Click **Apply** to save your changes, or **Cancel** to exit without saving.
- 10. **(Optional)** Repeat **Step 3** through **Step 9** to add additional **AWS** Object Storage Destinations.

## Add a New AWS Object Storage Destination that uses STS Authentication

Prior to configuring the AWS Object Storage Destinations with Storage Manager for AWS STS authentication, complete the following three steps, which apply to both AWS public cloud and GovCloud:

1. Have your AWS IAM Access Key Id and Secret Access Key ready.
2. Ensure that the S3 buckets, to be configured in Storage Manager, exist and you know both the name and the AWS region for each.

See <http://docs.aws.amazon.com/general/latest/gr/rande.html> for details on AWS Regions and Endpoints for AWS public cloud.

See <http://docs.aws.amazon.com/govcloud-us/latest/UserGuide/using-govcloud-endpoints.html> for details on AWS GovCloud (US) Endpoints.

3. Ensure that the IAM roles and their managed policies have been defined for your S3 buckets with AWS and that you know the roles' Amazon Resource Name (ARN). [Table 5 on the next page](#) provides an example of a role's managed policy that Storage Manager recommends.

**Table 5:** Example of a role's managed policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListAllMyBuckets",
      "Resource": "arn:aws:s3:::*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": "arn:aws:s3:::joe-bucket"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:PutObject",
        "s3:DeleteObject",
        "s3:AbortMultipartUpload",
        "s3:RestoreObject"
      ],
      "Resource": "arn:aws:s3:::joe-bucket/*"
    }
  ]
}
```

**i Note:** The Action "s3:RestoreObject" is only required if you are using Glacier Storage Class for your S3 bucket's life-cycle property.

Follow this procedure to add a new **AWS** Object Storage destination that uses **STS** authentication.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Click **New...** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue...**
4. Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new <b>AWS</b> Object Storage destination
<b>Provider</b>	Select <b>AWS</b> from the <b>Provider</b> list.
<b>Manager Host</b>	Enter <b>s3.amazonaws.com</b> for the <b>AWS</b> Object Storage Manager Host's address.
<b>Manager Port</b>	Enter <b>443</b> for the <b>AWS</b> Object Storage Manager Host's port.
<b>Manager Protocol</b>	Select <b>HTTPS</b> .
<b>Authentication Type</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default <b>Disabled</b> .
<b>User Name</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.
<b>Password</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.

5. In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

6. In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the " <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> " error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.
<b>Media Type</b> <i>(New in StorNext 5 release 5.4)</i>	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. Select <b>AWS</b> from the drop-down list.
<b>URL Style</b>	Select <b>VHOST</b> from the drop-down list.
<b>Object Access Protocol</b>	This specifies the network protocol to be used for the host. Select <b>HTTPS</b> from the drop-down list.

Parameter	Description
<b>Host[:Port]</b>	Enter the AWS region endpoint for your S3 bucket. For example: <ul style="list-style-type: none"><li>• <b>s3-us-west-2.amazonaws.com</b> for an S3 bucket created in the US West (Oregon) region</li><li>• <b>s3.amazonaws.com</b> for an S3 bucket created in the US East (N. Virginia) region</li><li>• <b>s3-us-gov-west-1.amazonaws.com</b> for an S3 bucket created in AWS GovCloud</li></ul>

7. In the **Containers** section, perform the following:
  - a. Leave the **Container Selection** at the default **Manual**.
  - b. Click **Add** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Enter the name of your S3 bucket.
<b>Media ID</b>	Enter the StorNext <b>Media ID</b> associated with the selected container. The <b>Media ID</b> must be unique.
<b>Media Type</b>	Select <b>AWS</b> from the drop-down list.
<b>Storage Class</b>	This parameter specifies the storage class of your S3 bucket's life-cycle property. Make sure that this value always matches that configured for the S3 bucket itself. Available values include <b>standard</b> , <b>standard_ia</b> , and <b>glacier</b> . Most configurations use the default <b>standard</b> storage class.
<b>Signing Type</b>	This parameter specifies the signing type for the requests sent to the AWS Object Storage server. Available values include <b>V2</b> and <b>V4</b> (default).
<b>Authentication Type</b>	This parameter specifies the authentication type for the container being configured. Available values include: <ul style="list-style-type: none"><li>• <b>STANDARD</b></li><li>• <b>STS_PUBLIC</b></li><li>• <b>STS_GOVCLOUD</b></li><li>• <b>CAP</b></li></ul> Since you are configuring your bucket to use <b>STS</b> authentication, select <b>STS_PUBLIC</b> if your bucket is in AWS public cloud. Select <b>STS_GOVCLOUD</b> if your bucket is in AWS GovCloud.
<b>Role</b>	This parameter specifies the Amazon Resource Name (ARN) of the IAM role to assume for obtaining temporary credentials. Enter the IAM role's ARN defined to access this container.

Parameter	Description
<b>Role Duration</b>	This parameter specifies the duration, in seconds, of the role session or temporary credentials. The value must be in the range <b>900</b> to <b>3600</b> . A default value of <b>3600</b> seconds is used if a role duration is not specified.
<b>Authentication Endpoint</b>	This parameter specifies an alternate authentication endpoint, which is used to override the default STS server for the AWS public or GovCloud region. Otherwise, leave it at the default blank.
<b>User Name</b>	Enter the <b>Access Key Id</b> for this container.
<b>Password</b>	Enter the <b>Secret Access Key</b> for this container.
<b>Copy Number</b>	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.
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details.

**i Note:** If no data has been written to a controller, I/O path, or container, click **Delete** to remove the item, and then click **Apply** to save the changes.

8. **(Optional)** Repeat **Step 7** to add more containers to the same **AWS** Object Storage Destination.
9. Click **Apply** to save your changes, or **Cancel** to exit without saving.
10. **(Optional)** Repeat **Step 3** through **Step 9** to add additional **AWS** Object Storage Destinations.

## Add a New AWS Object Storage Destination that uses CAP Authentication

If you use CAP authentication, complete the following six steps prior to configuring an object storage destination:

1. Ensure that the S3 buckets, to be used by Storage Manager, have been created with AWS FedCloud (C2S account) and that you know the name and the AWS region endpoint of each bucket.
2. Ensure that the IAM roles and their managed policies have been defined for your S3 buckets with AWS FedCloud (C2S account). See [Table 5 on page 460](#) for an example of a managed policy which includes the required capabilities of the role.
3. Ensure that you have the following information:
  - a. The IAM role associated with your C2S account.
  - b. The agency associated with your C2S account.
  - c. The mission associated with your C2S account.



- d. Your CAP server connection endpoint.
  - e. Your customized certificate authority (CA) file in PEM format.
  - f. Your X.509 client certificate in PEM format, and any applicable private key file or passphrase.
4. Add the following Storage Manager sysparms to the file `/usr/adic/TSM/config/fs_sysparm_override` on the StorNext system to enable Storage Manager to communicate with the CAP server:
- a. **FS\_OBJSTORAGE\_C2S\_CAP\_HOSTPORT** identifies the connection endpoint for the CAP server and can be configured as follows:

```
FS_OBJSTORAGE_C2S_CAP_HOSTPORT=cap-portal:port;
```

- b. **FS\_OBJSTORAGE\_CAPATH** identifies the directory in which the issuer's certificate authority (CA) can be found if it is not already included in the operating system's default trusted root certificate file.

---

**i Note:** The certificate should be in PEM format.

For example, the certificate can be copied to `/usr/cvfs/config/ssl` and configured as follows:

```
FS_OBJSTORAGE_CAPATH=/usr/cvfs/config/ssl;
```

---

**i Note:** If your customized CA PEM file contains more than one certificate, we recommend that you append the content of your customized CA PEM file to your operating system's default CA bundle and to **NOT** use sysparm **FS\_OBJSTORAGE\_CAPATH** to set the location of your customized CA PEM file. Otherwise, you could split your CA PEM file into multiple CA PEM files, each of which contains a single CA certificate, and use sysparm **FS\_OBJSTORAGE\_CAPATH** to set the location of your newly split single certificate CA PEM files. See [HTTPS Default CA ROOT Certificate File or Path on page 292](#) for additional information on how to configure your customized CA PEM files.

- c. **FS\_OBJSTORAGE\_CLIENTCERT** identifies the location of the X.509 client certificate installed on the system for the CAP server to authenticate.

---

**i Note:** The certificate should be in PEM format.

For example, the client certificate can be copied to `/usr/cvfs/config/ssl/client-cert-filepath`, and configured as follows:

```
FS_OBJSTORAGE_CLIENTCERT=/usr/cvfs/config/ssl/client-cert-filepath;
```

- d. **FS\_OBJSTORAGE\_CLIENTKEY** which sets the location of the client private key if the client private key

is kept separately from (for example, not included in) the client certificate file. This parameter can be configured as follows:

```
FS_OBJSTORAGE_CLIENTKEY=/usr/cvfs/config/ssl/client-key-filename;
```

- e. **FS\_OBJSTORAGE\_CLIENTKEY\_PASS** which specifies the passphrase used to protect the client private key and can be configured as follows:

```
FS_OBJSTORAGE_CLIENTKEY_PASS=passphrase;
```

- 5. Execute the following command to generate the hash for your certificates:

```
/opt/quantum/openssl/bin/c_rehash /usr/cvfs/config/ssl
```

- 6. Restart TSM to allow the system parameter changes to take effect.

**Note:** You can also use the GUI's [System Parameters on page 437](#) page to set the above system parameters.

Follow this procedure to add a new **AWS** Object Storage destination that uses **CAP** authentication.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Click **New....** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue....**
4. Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new <b>AWS</b> Object Storage destination
<b>Provider</b>	Select <b>AWS</b> from the <b>Provider</b> list.
<b>Manager Host</b>	Enter <b>s3.amazonaws.com</b> for the <b>AWS</b> Object Storage Manager Host's address.
<b>Manager Port</b>	Enter <b>443</b> for the <b>AWS</b> Object Storage Manager Host's port.
<b>Manager Protocol</b>	Select <b>HTTPS</b> .
<b>Authentication Type</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default <b>Disabled</b> .

Parameter	Description
<b>User Name</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.
<b>Password</b>	This parameter <b>does not</b> apply to <b>AWS</b> Object Storage Destination. Leave it at the default blank.

5. In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

6. In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the " <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> " error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.
<b>Media Type</b> <i>(New in StorNext 5 release 5.4)</i>	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. Select <b>AWS</b> from the drop-down list.
<b>URL Style</b>	Select <b>VHOST</b> from the drop-down list.
<b>Object Access Protocol</b>	This specifies the network protocol to be used for the host. Select <b>HTTPS</b> from the drop-down list.
<b>Host[:Port]</b>	Enter the AWS region endpoint for your S3 bucket.

7. In the **Containers** section, perform the following:

- a. Leave the **Container Selection** at the default **Manual**.
- b. Click **Add** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Enter the name of your S3 bucket.
<b>Media ID</b>	Enter the StorNext <b>Media ID</b> associated with the selected container. The <b>Media ID</b> must be unique.

Parameter	Description
<b>Media Type</b>	Select <b>AWS</b> from the drop-down list.
<b>Storage Class</b>	This parameter specifies the storage class of your S3 bucket's lifecycle property. Make sure that this value always matches that configured for the S3 bucket itself. Available values include <b>standard</b> , <b>standard_ia</b> , and <b>glacier</b> . Most configurations use the default <b>standard</b> storage class.
<b>Signing Type</b>	This parameter specifies the signing type for the requests sent to the AWS Object Storage server. Available values include <b>V2</b> and <b>V4</b> (default).
<b>Authentication Type</b>	<p>This parameter specifies the authentication type for the container being configured. Available values include:</p> <ul style="list-style-type: none"><li>• <b>STANDARD</b></li><li>• <b>STS_PUBLIC</b></li><li>• <b>STS_GOVCLOUD</b></li><li>• <b>CAP</b></li></ul> <p>Select <b>CAP</b> since you are configuring the container to use <b>CAP</b> authentication.</p>
<b>Role</b>	This parameter specifies the IAM role associated with the target C2S account for obtaining temporary credentials. Enter the IAM role associated with your C2S account to access this container.
<b>Role Duration</b>	This parameter specifies the duration, in seconds, of the role session or temporary credentials. The value must be in the range <b>900</b> to <b>3600</b> . A default value of <b>3600</b> seconds is used if a role duration is not specified.
<b>CAP Agency</b>	This parameter specifies the CAP agency associated with the target C2S account for obtaining temporary credentials. Enter the agency associated with your C2S account for this container.
<b>CAP Mission</b>	This parameter specifies the CAP mission associated with the target C2S account for obtaining temporary credentials. Enter the mission associated with your C2S account for this container.
<b>Copy Number</b>	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.
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details.

- 
- i Note:** If no data has been written to a controller, I/O path, or container, click **Delete** to remove the item, and then click **Apply** to save the changes.
  - 8. **(Optional)** Repeat **Step 7** to add more containers to the same AWS Object Storage Destination.
  - 9. Click **Apply** to save your changes, or **Cancel** to exit without saving.
  - 10. **(Optional)** Repeat **Step 3** through **Step 9** to add additional **AWS** Object Storage Destinations.

## Edit an AWS Object Storage Destination

Follow this procedure to edit an existing **AWS** Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the **AWS** 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.

---

**i 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 **AWS** Object Storage destination was successfully modified, click **OK**.

## Delete an AWS Object Storage Destination

Follow this procedure to delete an existing **AWS** Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the **AWS** 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 Object Storage destination was successfully deleted, click **OK**.

# Setting up Azure Object Storage Destinations on a StorNext Configuration

Storage Manager has Object Storage support for Microsoft Azure Cloud Services that includes:

- HTTP and HTTPS access to Microsoft Azure containers, though HTTPS is the recommended protocol for access Azure containers
- Append blob storage service

## View Azure Object Storage Destinations

Follow this procedure to view a list of currently configured Azure Object Storage destinations.

1. On the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab. Information for any previously configured **Azure** Object Storage destinations is shown as entries that have **Microsoft Azure** listed as the **Provider**. For each configured destination, the screen displays the **Name**, **Provider**, **Appliance 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 **Azure** Object Storage destination whose information you want to view.
4. Click **View...**
5. When you are finished viewing the destination information, click **Done**.

## Add a New Azure Object Storage Destination

Follow this procedure to add a new **Azure** Object Storage destination.

1. Ensure that the **Azure** containers that you are configuring with Storage Manager, have been created with Microsoft Azure Cloud Services; and that you know the **Storage Account Name**, **Storage Access Key**, and names of your **Azure** containers.
2. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
3. Click the **Object Storage** tab.
4. Click **New...** The page is updated and displays various configuration prerequisites. If any of the configuration prerequisites are required, click **Cancel** and configure them before returning to this page. Otherwise, click **Continue...**
5. Enter the appropriate value into the following parameters:

Parameter	Description
<b>Name</b>	Enter the name of the new <b>Azure</b> Object Storage destination
<b>Provider</b>	Select <b>Microsoft Azure</b> from the <b>Provider</b> list.
<b>Manager Host</b>	Enter <b>portal.azure.com</b> for the <b>Azure</b> Object Storage Manager Host's address.
<b>Manager Port</b>	Enter <b>443</b> for the <b>Azure</b> Object Storage Manager Host's port.
<b>Manager Protocol</b>	Select <b>HTTPS</b> .

Parameter	Description
<b>Authentication Type</b>	This parameter <b>does not</b> apply to <b>Azure</b> Object Storage Destination. Leave it at the default <b>Disabled</b> .
<b>User Name</b>	This parameter <b>does not</b> apply to <b>Azure</b> Object Storage Destination. Leave it at the default blank.
<b>Password</b>	This parameter <b>does not</b> apply to <b>Azure</b> Object Storage Destination. Leave it at the default blank.

6. In the **Controllers** section, click **Add** and then specify the following to add a controller:

Parameter	Description
<b>Name</b>	Enter the name of the controller.
<b>Max Streams</b>	By default, the maximum number of concurrent I/O streams per controller is <b>48</b> . This can be changed by selecting the desired value from the <b>Max Streams</b> drop-down list.

7. In the **I/O Paths** section, click **Add** and then specify the following to add an I/O path:

Parameter	Description
<b>Name</b>	Enter the unique name of the I/O path. If you do not have unique names, the " <b>Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed.</b> " error message appears.
<b>Controller Name</b>	Select the name of the controller associated with the new I/O path.
<b>Media Type</b> <i>(New in StorNext 5 release 5.4)</i>	This parameter specifies the object storage media type assigned to a namespace that is associated with a specific Object Storage API. Select <b>Azure</b> from the drop-down list.
<b>URL Style</b>	Leave the URL style at the default <b>PATH</b> .
<b>Object Access Protocol</b>	This specifies the network protocol to be used for the host. Select <b>HTTPS</b> from the drop-down list.
<b>Host[:Port]</b>	Enter <b>blob.core.windows.net</b> for the I/O path's connection endpoint to your Azure containers.

8. 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.
  - i. If you select **Scan**, check the **Use different credentials** box and enter the username (the **Azure Storage Account Name**) and a password (the **Azure Storage Access Key**). The

page displays a pre-populated list of available containers when you add container.

ii. If you select **Manual**, the page displays a text box to manually enter the name of the container.

b. Click **Add** and then specify the following to add a container:

Parameter	Description
<b>Container</b>	Select ( <b>Scan</b> mode) or enter ( <b>Manual</b> mode) the name of your <b>Azure</b> container.
<b>Media ID</b>	Enter the StorNext <b>Media ID</b> associated with the selected container. The <b>Media ID</b> must be unique.
<b>Media Type</b>	Select <b>Azure</b> from the drop-down list.
<b>Storage Class</b>	This parameter specifies the storage class for the <b>Azure</b> Object Storage media. Leave it at the default <b>azure_append_blob</b> .
<b>Signing Type</b>	This parameter specifies the signing type for the requests sent to the <b>Azure</b> Object Storage server. Leave it at the default <b>azure</b> .
<b>Authentication Type</b>	This parameter specifies the authentication type for the container being configured. Use the default value of <b>STANDARD</b> , which authenticates with the Storage Account Name and Storage Access Key.
<b>Account Name</b>	Enter the <b>Azure Storage Account Name</b> for this container.
<b>Key</b>	Enter the <b>Azure Storage Access Key</b> for this container.
<b>Copy Number</b>	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.
<b>Policy Class</b>	This parameter specifies the policy class that has the exclusive use of the container being configured. If left at <b>System Blank</b> , no policy class association is set for the container and the container can be used by all policy classes. To configure this parameter, select one of the pre-defined policy classes from the drop-down list for the <b>Policy Class</b> option. See <a href="#">How to Route File Copies to a Specific Object Storage Namespace on page 441</a> for additional details.

**i Note:** If no data has been written to a controller, I/O path, or container, click **Delete** to remove the item, and then click **Apply** to save the changes.

9. **(Optional)** Repeat **Step 8** to add more containers to the same **Azure** Object Storage Destinations.
10. Click **Apply** to save your changes, or **Cancel** to exit without saving.
11. **(Optional)** Repeat **Step 4** through **Step 10** to add additional **Azure** Object Storage Destinations.

## Edit an Azure Object Storage Destination

Follow this procedure to edit an existing **Azure** Object Storage destination.



1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the **Azure** 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 **Azure** Object Storage destination was successfully modified, click **OK**.

## Delete an Azure Object Storage Destination

Follow this procedure to delete an existing **Azure** Object Storage destination.

1. If you have not already done so, on the **Configuration** menu, click **Storage Destinations**.
2. Click the **Object Storage** tab.
3. Select the **Azure** 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 Object Storage destination was successfully deleted, click **OK**.

---

# Change the Current State of Object Storage Destinations, Controllers, and I/O Paths

You can also change the current state of existing Object Storage destinations, controllers, and I/O paths. To change the state, select the Object Storage destination, and then choose one of these options from the **Select Action** drop-down list:

Parameter	Description
<b>Online</b>	Select this option to set the Object Storage destination online.
<b>Offline</b>	Select this option to take the Object Storage destination offline.
<b>Controllers Online</b>	Select this option to set the controllers online.

Parameter	Description
<b>Controllers Offline</b>	Select this option to take the controllers offline.
<b>I/O Paths Online</b>	Select this option to set the I/O Paths online.
<b>I/O Paths Offline</b>	Select this option to take the I/O Paths offline.

---

## HTTPS Support for 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 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 on page 292](#) and [Update an Expired CA Root Certificate on page 302](#).

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.

## Configure 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`.

---

**i Note:** StorNext only supports certificates in `.pem` format. For additional information, see [Basic Secure Sockets Layer \(SSL\) Guidelines on page 497](#), as it outlines some standard information about using private and public certificates.

## Examples

### 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
```

### Update 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
```

### Add 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

Q-Cloud 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 on page 292](#) and [Update an Expired CA Root Certificate on page 302](#).

---

## Changes to Existing CLI Commands

Some existing Storage Manager commands are updated for use with the new Object Storage destination. Most of the updates to these commands are related to the object IDs that are now assigned to files that are stored to Object Storage. It is the object ID that is used to identify a file segment in the Object Storage. With this object ID and the namespace (determined via the media ID) a file can be accessed directly in the Object Storage appliance. Furthermore, as Storage Manager expands its support for new Object Storage providers and protocols, the commands are updated to support new features, such as new media types, storage classes, signing methods, or authentication schemes. 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.

The options `-e` and `-q` were also added to this command to display encryption and compression details on Object Storage copies that used these features.

---

**i 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 the new options added for reading a file from an Object Storage media. Various storage class, signing method, and authentication options are provided depending on the media type being used. See the `fsmedread` man page in the *MAN Pages Reference Guide* for additional details.

## The fsmedwrite Command

This command can be used to read a file from disk and write it to an Object Storage media. Various signing method and authentication options are provided depending on the media type being used. See the `fsmedwrite` man page in the *MAN Pages Reference Guide* for additional details.

## The fsmedscan Command

This command was updated to fail if an attempt is made to run against an Object Storage media. Due to the nature of the media and files stored there, no scanning of the media can be done to determine contents.

## The fsobjcfg Command

This command is used to configure and report Object Storage components and media for StorNext. The command supports a wide range of Object Storage providers, media types, network and authentication protocols. See the `fsobjcfg` man page in the *MAN Pages Reference Guide* for additional details.

## The fsaddclass and fsmodclass Commands

These commands have the `-e` and `-q` options added to support encryption and compression for Q-Cloud Archive™ and Vault™ products.

## 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 is 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 a maximum object size of 16 TiB when storing to LATTUS media, 5 TiB when storing to AWS S3 compatible media or Q-Cloud media, and 195 GiB when storing to AZURE media.
- Support for 16 TiB objects on LATTUS media (refer to the AmpliStor 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.
  - 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.

# Object Storage Segment Size

When a file is stored to Object 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 up into multiple segments and each segment is stored as an object in Object Storage. Segment size should be configured if very large files exist, since different Object Storage providers or protocols could place a limit on the maximum object size it can store. For example, Lattus 3.0.0 has a limit of 16 TiB for each object. If no segment size is configured, uploading a file whose size is larger than 16 TiB to Lattus Object Storage (version 3.0.0) will fail. Similarly, AWS S3 Multipart Upload and Microsoft Azure Blob protocols have a maximum object size of 5 TiB and 195 GiB, respectively.

Storage Manager uses the `MED_SEG_OVER_LATTUS` system parameter to control the segment size for files targeted for storage to LATTUS media. The default size is set at 128 GiB (137,438,953,472 bytes). Storage Manager will segment files greater than `MED_SEG_OVER_LATTUS`, when storing to LATTUS media.

The default segment size values for other Object Storage media types are set as follows:

- 128 GiB (137,438,953,472 bytes) for S3 media using `MED_SEG_OVER_S3` system parameter, which applies to Lattus S3 and Q-Cloud Archive S3 media.
- 5 TiB (5,497,558,138,880 bytes) for Q-Cloud Vault, S3COMPAT, and AWS media using `MED_SEG_OVER_QVAULT`, `MED_SEG_OVER_S3COMPAT`, and `MED_SEG_OVER_AWS` system parameters respectively.
- 195 GiB (209,715,200,000 bytes) for AZURE media using `MED_SEG_OVER_AZURE` system parameter.

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 or copy information. 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 Object Storage. 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 segment size could also reduce the overall storage and retrieval time, because Storage Manager can store and retrieve segments concurrently using multiple data moving streams.



# Chapter 13: Q Cloud

---

This chapter contains the following topics:

Overview of Q-Cloud .....	479
Setting Up Q-Cloud .....	481
Q-Cloud Tuning Considerations .....	486

---

## Overview of Q-Cloud

**Q-Cloud** is a cloud storage tier for StorNext 5 customers who need access to additional or off-premises storage capacity. Q-Cloud provides StorNext 5 users with easy and cost-effective on-demand access to cloud storage for increased flexibility, data protection, and availability. Q-Cloud is integrated with StorNext so there's no additional hardware or software, separate applications or programming. StorNext is known for its automatic policy-based data management; Q-Cloud extends that to the cloud for reliable, always-available storage and maximum flexibility.

Q-Cloud offers the following benefits:

### Ease of Use

- Q-Cloud provides StorNext users with the easiest and most efficient access to cloud storage.
- Tightly integrated into StorNext, Q-Cloud is available on demand to every StorNext customer.



- StorNext users already know how to use it because it's based on the same policy based data management they use for other storage tiers.
- No additional software, configuration or installation required.
- No programming or scripting required.
- No tedious file-by-file operations.
- No gateways to manage.

## Flexibility

- Available on-demand to instantly scale storage capacity.
- Users can respond to workload or business practice changes immediately.
- Free up valuable primary disk capacity for more demanding tasks.
- Simultaneously archive files to multiple storage tiers including object storage, tape and cloud.
- Utilize off-site and multi-copy data protection as needed.
- Start using the cloud at any time.
- Use only as much as you need.

## Reliability and Security

- Create automatic archive policies for data protection.
- Utilize off-site and multi-copy data protection as needed.
- Server side encryption protects data at rest.
- Data distribution and redundancy.

## Cost Effectiveness

- Create automatic archive policies to move files from higher cost storage tiers and free up capacity.
- Utilize off-site and multi-copy data protection as needed.
- Zero dollar upfront investment.
- No additional power, cooling or personnel expenses.
- Pay as you go and pay only for usage billing model.
- Preset billing alerts to help control usage and costs.

---

# Setting Up Q-Cloud

The **Q-Cloud** tab on the **Configuration > Storage Destinations** page enables you to perform actions pertaining to Q-Cloud resources. The **Q-Cloud** page displays a table showing the Q-Cloud resources that have been registered/activated for use in StorNext.

For information on troubleshooting Object Storage and Cloud errors, see [Debugging StorNext for Object Storage Systems and Cloud Providers on page 638](#).

## Q-Cloud Retrieval

- File retrieval from Q-Cloud Vault can take up to 5 hours.
- Retrieval requests to StorNext for files in Q-Cloud Vault are issued immediately, similar to other devices supported by Storage Manager.
- If an application requests multiple files (issues these requests one file at a time) and then waits for each retrieval to complete before issuing the next file to be retrieved, each file can take up to 5 hours. Since the requests are sequential (it waits for the previous retrieve to complete for two files), file retrieval can take up to 10 hours to complete. However, if the retrievals are submitted to StorNext concurrently, it can take up to 5 hours for two files to complete.
- While each file retrieval waits for completion (up to 5 hours), the retrieval process consumes and retains a certain amount of system resources for that length of time until completion.

---

 **Caution:** A large number of file retrieval requests issued concurrently can severely impact system performance.

## Q-Cloud and Partial File Retrieval (PFR)

- The **Partial File Retrieval** feature is supported with Q-Cloud Archive targets, with the exception of configurations where client-side encryption or compression is used.
- The **Partial File Retrieval** feature is not supported with Q-Cloud Vault targets.

For additional information on **Partial File Retrieval**, see the *StorNext Partial File Retrieval User's Guide*.

## Firewall Rules and IP Ranges Specific to Q-Cloud

If the StorNext configuration has restricted network access because of firewall rules, the firewall configuration may need to be updated to use Q-Cloud.

## Configure the Firewall if the Software Supports Named Entries

1. Open HTTPS access (port **443**) to `s3-us-west-2.amazonaws.com` (note, this is an example for the

**US-West-2** region). The appropriate region should be used in its place.

2. For access to the Q-Cloud Controller, open HTTPS access (port **443**) to `api-qcloud.quantum.com`. This allows you to configure and validate your Q-Cloud installation.

## Configure the Firewall if the Software Does Not Support Named Entries

1. Locate the IP addresses for the appropriate AWS region in <https://ip-ranges.amazonaws.com/ip-ranges.json>. Allow access to these addresses. The list may change several times per week.


**i Note:** Amazon rotates IP addresses through a range in order to provide access to the AWS S3 service. Each region of AWS has a specific set of IP ranges that are described at <https://aws.amazon.com/blogs/aws/aws-ip-ranges-json/>, and listed at <https://ip-ranges.amazonaws.com/ip-ranges.json>.

2. Determine the IP address of `api-qcloud.quantum.com`, and open HTTPS access to this address.

Alternatively, open all outgoing HTTPS traffic from the StorNext installation to the Internet.

## Parameters and Buttons on the the Q-Cloud Page

Parameter/Button	Description
<b>Media ID</b>	Displays the <b>Media ID</b> representing a purchased resource. It relates the Q-Cloud bucket to a StorNext media.
<b>Product Key</b>	Displays the <b>Product Key</b> associated with the StorNext media.
<b>Provider</b>	Displays the underlying storage provider.
<b>State</b>	Displays the operational state of the storage media.
<b>Max Streams</b>	The <b>Max Streams</b> value defines the number of concurrent I/O operations that can write concurrently. By default, the maximum number of streams is set to <b>48</b> , or you can select a different value from the <b>Max Streams</b> drop-down list when you click <b>Preferences...</b> See <a href="#">Configure a Media ID Preference on page 485</a> .
<b>Copy Number</b>	Displays the copy number assigned to the Q-Cloud media. If files do not exist on the media, you can select copy number <b>1</b> , <b>2</b> , <b>3</b> , or <b>4</b> from the list when you click <b>Preferences...</b> See <a href="#">Configure a Media ID Preference on page 485</a> .
<b>Policy Class</b>	Displays the name of the policy.
<b>Evaluation Key</b>	Displays <b>Yes</b> or <b>No</b> . If <b>Yes</b> , the product key is an evaluation key. If <b>No</b> , the product key is not an evaluation key.
<b>Expiration Date</b>	Displays the expiration date of the selected product key.

Parameter/Button	Description
<b>Overused</b>	This parameter applies to evaluation keys and indicates if you have exceeded the limits of the evaluation. For example, too much data stored in Q-Cloud.
<b>Region</b>	Displays the region the Q-Cloud media belongs to.
<b>File Count</b>	Displays the number of files currently stored on the media.
<b>Manage Keys...</b>	Click to display the <b>Manage Keys</b> page where you can add or remove additional <b>Product Keys</b> . For additional procedures, see <a href="#">Add a Product Key on the next page</a> and <a href="#">Delete a Product Key on the next page</a> below.
<b>Preferences...</b>	Click to configure your preferences for the currently selected product key. See <a href="#">Configure a Media ID Preference on page 485</a> .
<b>Check Connectivity</b>	Click to check the connectivity to the Q-Cloud media. If connectivity fails, contact Quantum Technical Support.
<b>Refresh</b>	Click to update the <b>Q-Cloud</b> page with current information.
<b>Select Action (Online)</b>	Select <b>Online</b> manually change the media state to online.
<b>Select Action (Offline)</b>	Select <b>Offline</b> to manually change the media state to offline.
<b>Select Action (Remove All Q-Cloud Storage)</b>	Select <b>Remove All Q-Cloud Storage</b> to purge and factory reset the system.  <b>Note:</b> The option is only available if the Q-Cloud media do not contain files.

## Parameters and Buttons on the Manage Keys Page

On the **Configurations > Storage Destinations > Q-Cloud** page, click **Manage Keys...** to display the **Manage Keys** page.

For additional procedures, see [Add a Product Key on the next page](#) and [Delete a Product Key on the next page](#) below.

Parameter/Button	Description
<b>Access ID</b>	Enter the customer access ID provided by Quantum.
<b>Product Keys</b>	Enter the customer product key provided by Quantum.
<b>Delete</b>	Click to remove the selected <b>Product Key</b> from the list.  You must enter an <b>Access ID</b> in order to delete <b>Product Keys</b> .

Parameter/Button	Description
<b>Add</b>	Click to add a new <b>Product Key</b> to the list after you enter it into the <b>Product Keys</b> field. You must enter an <b>Access ID</b> in order to add <b>Product Keys</b> .
<b>Apply</b>	Click to commit the edited list of <b>Product Keys</b> and/or <b>Access ID</b> .
<b>Reset</b>	Click to restore the current list.
<b>Cancel</b>	Click to cancel the editing session and return to the <b>Q-Cloud</b> page.

## Add a Product Key

1. In the **Access ID** field, enter the customer access ID provided to you by Quantum. Upon entering the **Access ID** once and successfully committing a configuration, the **Access ID** field remains populated when you return to the **Manage Keys** page.
2. In the **Product Keys** field (adjacent to the **Add** button), enter the customer product key provided to you by Quantum, select a **Policy Class** to assign the Q-Cloud media directly to it, and then click **Add**. The **Product Keys** list (above the field) displays the value just added. You must enter an **Access ID** in order to add **Product Keys**. Repeat this step to add additional product keys.

**i Note:** If a **Policy Class** is not selected, the media is assigned to **System Blank** (in other words, the default blank pool).

3. Click **Apply**. The **Product Key** appears on the main **Q-Cloud** page. Upon clicking **Apply**, the actual process to commit the configuration is run as a background task. The GUI returns to the **Q-Cloud** page immediately if there are no locally detected errors. In some cases, the background process may require some time to execute and the **Q-Cloud** page will not immediately reflect your changes. As with other similar processes, navigate to the [Jobs on page 365](#) page to verify the outcome of the background task. In the case of a delayed success, click **Refresh** on the **Q-Cloud** page to display the current state of the system. In the case of a failure, the **Q-Cloud** page may not update; navigate to the [Jobs on page 365](#) page to determine the cause of the error.
4. **(Optional)** Click **Reset** to restore the current list.
5. **(Optional)** Click **Cancel** to cancel the editing session and return to the **Q-Cloud** page.

## Delete a Product Key

1. In the **Access ID** field, enter the customer access ID provided to you by Quantum.
2. In the **Product Keys** list, click a product key, and then click **Delete** to remove the selected product key from the **Product Keys** list. You must enter an **Access ID** in order to delete **Product Keys**. Repeat this step to delete additional product keys.
3. Click **Apply**. The **Product Key** is removed on the **Q-Cloud** page. Upon clicking **Apply**, the actual process to commit the configuration is run as a background task. The GUI returns to the **Q-Cloud** page immediately if there are no locally detected errors. In some cases, the background process may require some time to execute and the **Q-Cloud** page will not immediately reflect your changes. As with other

similar processes, navigate to the [Jobs on page 365](#) page to verify the outcome of the background task. In the case of a delayed success, click **Refresh** on the **Q-Cloud** page to display the current state of the system. In the case of a failure, the **Q-Cloud** page may not update; navigate to the [Jobs on page 365](#) page to determine the cause of the error.

4. **(Optional)** Click **Reset** to restore the current list.
5. **(Optional)** Click **Cancel** to cancel the editing session and return to the **Q-Cloud** page.

## Configure a Media ID Preference

1. On the **Configurations > Storage Destinations > Q-Cloud** page, click a radio button to select a media ID you want to configure (corresponding to a media ID located under the **Media ID** column).
2. Click **Preferences...** to display the **Preferences** page. The table below provides a description for each parameter/button on the **Preferences** page.

Parameter/Button	Description
<b>Product Key</b>	Displays the product key corresponding to the <b>Media ID</b> you selected on the Q-Cloud page from Step 1.
<b>Max Streams</b>	The <b>Max Streams</b> value defines the number of concurrent I/O operations that can write concurrently. By default, the maximum number of streams is set to <b>48</b> , or you can select a different value from the <b>Max Streams</b> drop-down list.
<b>Copy Number</b>	Displays the copy number assigned to the <b>Q-Cloud</b> media. Select copy number <b>1, 2, 3,</b> or <b>4</b> from the list if files do not exist on the media.
<b>Update</b>	Click to update your preferences for the currently selected <b>Product Key</b> .
<b>Cancel</b>	Click to cancel the editing session and return to the <b>Q-Cloud</b> page.

3. In the **Max Streams** drop-down list, select the number of streams. For **Copy Number**, the only option available is 1.
4. Click **Update** to proceed and display the confirmation dialog requesting you confirm your preferences. Perform one of the following:
  - Click **Yes** to confirm your selected preferences. A confirmation dialog appears; click **OK**. The **Q-Cloud** page is displayed with your updated preferences.
  - Click **No** to disregard your selected preferences and return to previous page.
5. **(Optional)** Click **Cancel** to cancel the editing session and return to the **Q-Cloud** page.

## Replace a Product Key

Follow the procedure below to replace your Q-Cloud Product Key. For example, if your Q-Cloud Product Key is compromised and you are reissued a new Product Key from Quantum.

1. On the **Configurations > Storage Destinations > Q-Cloud** page, click a radio button to select a media ID you want to configure (corresponding to a media ID located under the **Media ID** column).

2. Click **Preferences...** to display the **Preferences** page (see [Configure a Media ID Preference on the previous page](#)).
3. Note the **Product Key** corresponding the **Media ID** you selected in **Step 1**.
4. Click **Cancel** to return to the **Q-Cloud** page.
5. Click **Manage Keys.....**
6. In the **Product Keys** list, click the obsolete product key (from **Step 3**) and then click **Delete**.
7. In the **Product Keys** field (adjacent to the **Add** button), enter the **new** product key provided to you by Quantum, and then click **Add**. The **Product Keys** list (above the field) displays the value just added.
8. Click **Apply**. The new **Product Key** appears on the **Q-Cloud** page.

---

**i Note:** If you click **Apply** after deleting the obsolete product key (**Step 6**) without adding the **new** product key, you will receive a **putConfig** error message.

## Enable Encryption on the Q-Cloud Data

1. On the **Configuration** menu, click **Storage Policies**.
2. Click the radio button corresponding to the Q-Cloud policy class you want to enable encryption on, and then click **Edit....**
3. Click the **General** tab.
4. Click the **Encryption (Q-Cloud only)** list, and then select **Server AES256**.
5. Click **Apply** to save your changes to the policy class.
6. **(Optional)** Click **Cancel** to discard your changes and return to the **Storage Manager Policies** page.

## Backup and Restore Using Q-Cloud

Contact Quantum Technical Support for assistance with disaster recovery to and from Q-Cloud.

---

# Q-Cloud Tuning Considerations

StorNext's capability to push data may exceed the bandwidth capability of your network for Q-Cloud. This may consume your entire WAN bandwidth and can cause packet lost leading to disconnections between StorNext and Q-Cloud destinations.

Review the following sections and consider the tuning suggestions.

## General Considerations

- By default, Storage Manager policies retain 10 inactive versions of files, in addition to the current (active) version. As this will consume up to 11 times the storage of the active version, most users will want to save 0 inactive versions.
- By default, Storage Manager policies retain files in the database, and on media, for up to seven years. To get cost savings quickly after deleting unwanted files, check the **Clean Database When File Removed** checkbox in the policy.
- To free up space, delete files and execute the command `fsclean` to clean the database and Q-Cloud.
- Q-Cloud Vault is intended for the coldest archive data: it offers a very low per-TB storage cost, but carries charges for data out, retrieve, and deletion (or change) of data in less than 90 days. Use Q-Cloud Vault for unchanging files that will very rarely be accessed. Time to first byte is up to 5 hours with Q-Cloud Vault, and retrieves must be manually throttled to control retrieve fees. For additional information, see [Q-Cloud Retrieval on page 481](#).
- Q-Cloud modification to policy class settings may be advisable if many files change frequently such that, in addition to lowering **Maximum Inactive Versions**, to have the Storage Manager policy **Maximum File Age (Hours)** set sufficiently high to minimize storing files that are likely to change soon. Also, for Q-Cloud Vault, writing many versions of files may also imply many deletions of old files, which may lead to deletion or change fees. Of course, higher **Maximum File Age (Hours)** settings comes with risks if something happens to the file before it is stored.
- Q-Cloud Archive is intended for data that changes more frequently than Q-Cloud Vault: it has a higher per-TB storage cost, but does not have retrieve fees or early deletion fees.

## TCP/IP Tuning

You can manage the TCP/IP traffic by tuning TCP buffers. For example, where we have seen server side disconnections, reducing the size of `tcp_rmem` or `tcp_wmem` has improved connection reliability. Quantum does not have any specific settings that are known to work for all situations as it is very dependent on your networking environment.

## Network Interface (NIC) Bonding

Using network interface bonding can improve both network bandwidth and reliability. If you have more than one network interface on the host to Q-Cloud destinations, you may want to consider using NIC bonding.

## Number of Concurrent IO Streams

The current default concurrent streams set for each Q-cloud device is 48 concurrent streams which may be too high for your WAN. If this is too high, you can lower this stream count. The optimal stream count depends on your network bandwidth and reliability.



## Quality of Service Configuration

You may want to consider isolating Q-Cloud traffic where you can configure QOS for the overall WAN bandwidth allowed to be consumed. StorNext does not contain any configurable parameters to manage the total bandwidth it consumes.

## Q-Cloud Time Stamp Requirement

A valid time on your StorNext system that is within 15 minutes of Q-Cloud is required. If not, the request will fail with the RequestTimeTooSkewed error code. Quantum recommends you configure NTP on your system.

## Configure Storage Manager to Perform Continuous Retries

If the network is unreliable, and the system is experiencing a large amount of connectivity or client/server connection resets, Storage Manager may take the iopath to the Q-Cloud destination off-line. When this happens you will have to manually execute the command `fschstate` to enable the path on-line.

To configure Storage Manager to perform continuous retries, you can set the following parameters in the `/usr/adic/TSM/config/fs_sysparm_override` file:

- `FS_THRESHOLD_INC_NUM`: Amount to increment each time a drive device is used with errors encountered.
- `FS_THRESHOLD_DEC_NUM`: Amount to decrement each time a drive device is used without errors.
- `FS_DRIVE_ERR_THRESHOLD`: Threshold value, which when equaled or surpassed, drives will be taken off-line.

The default values are:

- `FS_THRESHOLD_INC_NUM=5;`
- `FS_THRESHOLD_DEC_NUM=1;`
- `FS_DRIVE_ERR_THRESHOLD=20;`

When the `FS_DRIVE_ERR_THRESHOLD` is met, Storage Manager will take the path off-line. To configure Storage Manager to never take a path off-line, you can set `FS_THRESHOLD_INC_NUM` to `0`.

---

**Note:** These are Global parameters that affect all Storage Manager devices, not just Q-Cloud paths.

## Compression or Encryption Considerations

The **Compression** or **Encryption** features add additional load on a metadata controller's resources. Quantum recommends the **Stream Count** for Q-Cloud storage destinations be configured for two to four

streams. Increasing the **Stream Count** with **Encryption** and **Compression** increases the CPU load on the metadata controller. For additional information on **Compression** and **Encryption**, see [Client Side Compression and Encryption on page 114](#).



# Appendix A: Operating Guidelines

---

This appendix contains the following topics pertinent to operating StorNext, as well as some operating guidelines and limitations you should consider:

Gateway Server/Client Network and Memory Tuning .....	491
Configuring LDAP .....	492
Setting Up Restrictive ACLs .....	493
Default Single-Path I/O Retry Behavior .....	493
Event Handles for fsm.exe on a Windows Metadata Server .....	494
FSBlockSize, Metadata Disk Size, and JournalSize Settings .....	494
Disk Naming Requirements .....	495
Changing StorNext's Default Session Timeout Interval .....	496
Configuring a Data Partition for Use with Spectra Logic T-series Tape Storage Libraries .....	497
Basic Secure Sockets Layer (SSL) Guidelines .....	497
Name Limitations .....	500
Ports Used By StorNext .....	501
Log Rolling and Disk Space Health Check .....	503
General Operating Guidelines and Limitations .....	505

---

# 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 below](#)
- [Gateway Server Memory Tuning on the next page](#)

---

**i Note:** For additional information about LAN client and server performance tuning, see the *StorNext 5 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 x # of NICs on the Gateway Server used for Distributed LAN traffic x server buffer count x 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. Refer to the `dpserver` and `sndpsc fg` man pages for information about viewing and modifying Distributed LAN buffer settings on Linux. For this case:

Required Memory = 1GB + (4 x 2 x 8 x 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 supports Light Directory Access Protocol, or LDAP (RFC 2307). This feature allows you 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.

## 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:

- `UnixDirectoryCreationModeOnWindows Default 0755`
- `UnixFileCreationModeOnWindows Default 0644`

StorNext allows one set of values for all users of each file system.

---

**i 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

By default, StorNext continuously retries I/O operations until they succeed, regardless of the number of I/O paths. However, you can override the behavior by configuring the **I/O Retry Time** feature.

For additional information about the **I/O Retry Time** feature, refer to 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.

Beginning with StorNext 5, all file systems use a File System Block Size (`FsBlockSize` [FSB]) of 4KB. This is the optimal value and is no longer tunable. Any file systems created with versions prior to StorNext 5 will be automatically converted to use 4KB the first time the file system is started with StorNext 5. While internally file systems which have been upgraded from StorNext 4.x will use a 4KB block size, StorNext tools will continue to display the original `FsBlockSize` values. This is done to ensure that StorNext 5 can continue to support prior versions of StorNext clients.

### Metadata Disk Size Setting

When sizing volumes for metadata, provision an additional 2GB of metadata per million files and directories in the file system.

## JournalSize Setting

Beginning with StorNext 5, the recommended setting for `JournalSize` is 64Mbytes.

Increasing the `JournalSize` beyond 64Mbytes may be beneficial for workloads where many large size directories are being created, or removed at the same time. For example, workloads dealing with 100 thousand files in a directory and several directories at once will see improved throughput with a larger journal.

The downside of a larger journal size is potentially longer FSM startup and failover times.

Using a value less than 64Mbytes may improve failover time but reduce file system performance. Values less than 16Mbytes are not recommended.

---

**i Note:** Journal replay has been optimized with StorNext 5 so a 64Mbytes journal will often replay significantly faster with StorNext 5 than a 16Mbytes journal did with prior releases.

A file system created with a pre-5 version of StorNext may have been configured with a small `JournalSize`. This is true for file systems created on Windows MDCs where the old default size of the journal was 4Mbytes. Journals of this size will continue to function with StorNext 5, but will experience a performance benefit if the size is increased to 64Mbytes. This can be adjusted using the `cvupdatefs` utility. For more information, see the command `cvupdatefs` in the *StorNext MAN Pages Reference Guide*.

If a file system previously had been configured with a `JournalSize` larger than 64Mbytes, there is no reason to reduce it to 64Mbytes when upgrading to StorNext 5.

### Example (Linux)

```
<config configVersion="0" name="example" fsBlockSize="4096"  
journalSize="67108864">
```

### Example (Windows)

```
JournalSize 64M
```

---

## 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.

## Change the Timeout Interval

---

**i Note:** 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

## Configure the Data Partition

Refer to the specific product documentation to configure the libraries data partition.

---

## 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 499](#).

- 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 -ou 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 500](#).
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 **Object Storage Certificates**. The **Tools > Object Storage Certificates** page appears.
  5. On the **Tools > Object Storage 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 > Object Storage Certificates** page displays, "Certificate public.pem uploaded successfully."

## Example of a server.pem File

```
-----BEGIN RSA PRIVATE KEY-----
izfrNTmQLnfsLzi2Wb9xPz2Qj9fQYgGeug3N2MkDuVHwpPcgkhHkJgCQuuvT+qZI
MbS2U6wTS24SZk5RunJIUkitrKkewMS28SLGfKDs1bBY1SPa5smAd3/q10ePi4ae
dU6YgWuDxzBAKEKVSUu6pA2H0dyQ9N4F1dI+F8w9J990zE93EgyNqZFBBa2L70h4
M7DrB0gJBWMDUMoxGnun5glLiCMo2JrHZ9RkMia1lS1sHMhELx2UA1P8I1+0Mav8
iMlHGyUW8EJy0paVf09MPpceEcVwDBeX0+G4UQ10551GTFt0SRjcD8U+GkCzka9W
/SFQrSge3Gh3SDa0w/4JEMAjWPDLiCglwh0rLI04VwU6AxzTCuCW3d1ZxQsU6VFQ
PqHA8ha0UATZIrP3886PBThVqALBk9p1Nqn51bXlLh13Zy9DZIVx4Z5Ioz/EGuzgR
d68VW5wybLjYE2r6Q9nHpitSZ4ZderwjIZRes67HdxYFw8unm4Wo6kuGnb5jSSag
vwBxKzAf30mn+J6IthTJKuDd13rKZGMcRpQQ6VstwhYt1TahQ/qfJUWPjPcU5ML
9LkgVwA8Ndi1wp1/sEPe+U1L16L6v09jUHcueWN7+zsUOE/cDSJyMd9x/ZL8QASA
ETd5dujVIq1INL2vJKr1o4T+i0RsnpfFiqFmBK1Fqww/SkzJeChdyEtpa/dJMrt2
8S86b6zEmkser+SDYgGketS2DZ4hb+vh2ujSxmS8Gkwrn+BfHMzkbtkio8lWbGw0l
eM1tfdFZ6wMTLkxRhBkBK4JiMiUMvpERYpib6a2L6iXTfH+3RUDS6A==
-----END RSA PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
MIICMzCCAzyAwIBAgIJALiPnVsvq8dsMA0GCSqGSIb3DQEBBQUAMFMxCzAJBgNV
BAYTA1VTMQwwCgYDVQQIEwNmb28xDDAKBgNVBACjTA2ZvbzEMMAoGA1UEChMDZm9v
MQwwCgYDVQQLEwNmb28xDDAKBgNVBAMTA2ZvbzAeFw0xMzAzMTkxNTQwMTlaFw0x
ODAzMTgxNTQwMTlaMFMCzAJBgNVBAYTA1VTMQwwCgYDVQQIEwNmb28xDDAKBgNV
BACjTA2ZvbzEMMAoGA1UEChMDZm9vMQwwCgYDVQQLEwNmb28xDDAKBgNVBAMTA2Zv
bzCBnzANBghkqhkiG9w0BAQEFAA0BjQAwgYkCgYEAzdGfxi9CNbMf1UUcvDQh7MYB
OveIHyc0E0KIbHjK5FkCBU4CiZrbfHagaW7ZEcn0tt3Evpb0Mxxc/ZQU2WN/s/wP
xph0pSfsfFsTKM4RhTWD2v4fgk+xZiKd1p0+L4hTtpwnEw0uXRVd0ki6muwV5y/P
+5FHUelDq+pgTcgzuK8CAwEAAAMPMA0wCwYDVR0PBAQDAgkMA0GCSqGSIb3DQEB
BQUAA4GBAjiDAAtY0mQQeuxWdzLRzXmjvdSuL9GoyT3BF/jSnpXz5/58dba8pWen
v3pj4P3w5Do0so0rzkZy2jEsEitlVM2mLSbQpMM+MUVQCQoiG6W9xuCFuxSrwPIS
pAqEAuV4DN0xQKKWmhVv+J0ptMWD25PnpXeq5sXzghfJns1JlQND
-----END CERTIFICATE-----
```

## Example of a `public.pem` File

```
-----BEGIN CERTIFICATE-----
MIICMzCCAzygAwIBAgIJALiPnVsvq8dsMA0GCSqGSIb3DQEBBQUAMFMxCzAJBgNV
BAYTA1VTMQwwCgYDVQQIEwNmb28xDDAKBgNVBACjTA2ZvbzEMMAoGA1UEChMDZm9v
MQwwCgYDVQQLEwNmb28xDDAKBgNVBAMTA2ZvbzAeFw0xMzAzMTkxNTQwMTlaFw0x
ODAzMTgxNTQwMTlaMFMCzAJBgNVBAYTA1VTMQwwCgYDVQQIEwNmb28xDDAKBgNV
BACjTA2ZvbzEMMAoGA1UEChMDZm9vMQwwCgYDVQQLEwNmb28xDDAKBgNVBAMTA2Zv
bzCBnzANBkqhkiG9w0BAQEFAAOBjQAwGyKCGYEAzdGfxi9CNbMf1UUCvDQh7MYB
OveIHyc0E0KIbHjK5FkCBU4CiZrbfHagaW7ZEcn0tt3Evpb0Mxxc/ZQU2WN/s/wP
xph0pSfsfFsTKM4RhTWD2v4fgk+xZiKd1p0+L4hTtpwnEw0uXRVd0ki6muwV5y/P
+5FHUe1dq+pgTcgzuK8CAwEAAAMPMA0wCwYDVR0PBAQDAgkMA0GCSqGSIb3DQEB
BQUAA4GBAjiDAAtY0mQQeuxWdzLRzXmjvdSuL9GoyT3BF/jSnpxz5/58dba8pWen
v3pj4P3w5Do0so0rzkZy2jEsEitlVM2mLSbQpMM+MUVQCQoiG6W9xuCFuxSrwPIS
pAqEAuV4DN0xQKKWmhVv+J0ptMWD25Pnpxeq5sXzghfJns1Jl1QND
-----END CERTIFICATE-----
```

## Name Limitations

[Table 6 below](#) contains words that cannot be used as names for any part of a StorNext system.

**Table 6:** Name Limitations

Active	Enabled	JournalSize	Regular
Affinity	Exclusive	KiloInteger	Rotate
AllocationStrategy	Exit	Letter	Rtios
AttrTokenSize	Extended	Log	RtiosReserve
B	FileName	Long	Rtmb
Bits	FI	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	JournalIcBufNum	Raw	Write
Dump	JournalIcBufSize	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.

Appendix A: Operating Guidelines  
Ports Used By StorNext

Port	Protocol	StorNext Use	Notes
81	TCP	GUI (Java), Web Services	User starts at port 81, redirected to 443
443	TCP	GUI (Java), Web Services	
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	
Various	TCP	fsm, fsmpm	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	fsmpm	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.

Port	Protocol	StorNext Use	Notes
20566	TCP	MySQL	Only used internally on an MDC.
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: 80%), 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
```

- i 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.

1. If this is a High Availability (HA) configuration, execute the following command on both MDCs:

```
# mkdir -p /stornext/snfs1/.SNSM_LOG_BACKUP/`/bin/hostname`
```

2. 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
```

3. 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
```

4. Execute the following command to edit the tdlm crontab on the active MDC:

```
# /usr/bin/crontab -e -u tdlm
```

5. 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 7 on the next page](#) lists updated information and guidelines for running StorNext, as well as known limitations.

**Table 7:** 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> <b>Note:</b> Experimentation may be required to determine the optimal value.</p>
Linux	StorNext File System	<p>StorNext File System does not support the Linux <code>sendfile()</code> system call. This issue causes Apache web servers to deliver blank pages when content resides on StorNext file systems.</p> <p>This issue also affects Samba servers running on Linux.</p> <p>The workaround is to disable <code>sendfile</code> usage by adding the following entry into the Apache configuration file <code>httpd.conf</code>:</p> <pre>EnableSendfile off</pre> <p>The workaround for Samba servers is to add the following line into the configuration file:</p> <pre>sendfile=no</pre>

Operating System	Feature or Category	Description
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 <b>slocate</b> 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><b>RedHat Enterprise Linux 4, 5, and 6</b></p> <p>Add “cvfs” to the list of file system types to be skipped. This is usually done by modifying the “PRUNEFS” line in the <code>/etc/updatedb.conf</code> file to read:</p> <pre>PRUNEFS="cvfs sysfs selinuxfs usbdevfs devpts NFS nfs nfs4 afs sfs proc smbfs cifs autofs auto iso9660 udf"</pre> <p><b>SUSE Linux Enterprise Server 10, and 11</b></p> <p>The optional “findutils-locate” package is used to build the <code>slocate</code> database. The default behavior is to disable building the database. If enabled, to prevent <code>cvfs</code> file systems from being scanned, add “cvfs” to the list of file system types to be skipped. This is usually done by modifying the “UPDATEDB_PRUNEFS” line in the <code>/etc/sysconfig/locate</code> file to read:</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"><li>• If the file system is running Xsan 2.1.1 or earlier, it should be a simple upgrade: just replace the MDC.</li><li>• 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.</li><li>• If the file system is running Xsan 2.2 or later with “NamedStreams Yes,” you must completely remake (reformat) the file system. For obvious reasons, you should do a complete backup before migrating.</li></ul>
	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. 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 section <a href="#">Ports Used By StorNext on page 501</a>.</p>

Operating System	Feature or Category	Description
Linux	HA	<p>Changing the <code>haFsType</code> parameter in a file system configuration file to one of the HA types, and then (re)starting its FSM enables HA-specific features that change the functionality of StorNext.</p> <p>When the <code>HaShared</code> or <code>HaManaged</code> types are configured, other changes must be completed by successfully running the <code>cnvt2ha.sh</code> script, which is indicated by the creation of the</p> <p><code>/usr/adic/install/.snsm_ha_configured</code> touch file (<code>\$SNSM_HA_CONFIGURED</code> environment variable). No conversion is done or necessary for SNFS only (<code>HaUnmanaged</code>) configurations.</p> <p>If the conversion is not successfully completed, the <code>HaManaged</code> FSMs will not start, and the <code>HaShared</code> FSM will cause an HA Reset when it is stopped.</p> <p>To remedy this situation, edit every FSM configuration file to set its <code>haFsType</code> parameter to <code>HaUnmonitored</code>, then run the following commands to avoid the HA Reset in this special case only:</p> <pre>touch /usr/cvfs/install/.vip_down_hint service cvfs stop</pre>
	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 <code>127.0.0.2</code> 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 <code>127.0.0.2</code> entry should be deleted from the <code>/etc/hosts</code> file.</p>


Operating System	Feature or Category	Description
Linux	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 (turboboost), 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 displays if a system is impacted by this issue.</p>
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><b>Note:</b> 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>
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 <a href="http://support.microsoft.com/kb/263253">http://support.microsoft.com/kb/263253</a>).</p> <p>To work around this issue, follow these steps:</p> <ol style="list-style-type: none"><li>1. Click Start, and then click <b>Run</b>.</li><li>2. In the Open box, type Regedit and then click <b>OK</b>.</li><li>3. On the Edit menu, click <b>Find</b>.</li><li>4. In the Find what box, type Snfs_XXX.dat and then click <b>Find Next</b>.</li><li>5. If the search result selects a string value called PackageName, continue with these steps. Otherwise, repeat steps 3-4.</li><li>6. Double-click the <b>PackageName</b> string value.</li><li>7. In the <b>Value</b> 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.)</li><li>8. On the Registry menu, click <b>Exit</b>.</li></ol>



Operating System	Feature or Category	Description
Windows	Recycle bin	<p>If you are using the StorNext client software with the Windows Operating System 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"><li>1. On the Windows client machine, right-click the <b>Recycle Bin</b> icon on the desktop and then click <b>Properties</b>.</li><li>2. Click <b>Global</b>.</li><li>3. Click <b>Configure drives independently</b>.</li><li>4. Click the <b>Local Disk</b> tab that corresponds to the mapped or directory-mounted file system.</li><li>5. Click the checkbox <b>Do not move files to the Recycle Bin. Remove files immediately when deleted</b>.</li><li>6. Click <b>Apply</b>, and then click <b>OK</b>.</li></ol>

Operating System	Feature or Category	Description
Windows	Recycle bin (cont.)	<p>(Disabling the Recycle Bin, Continued)</p> <p>For Windows Server 2008, Windows Server 2012, Windows Vista, Windows 7 and Windows 8 systems, you must disable the Recycle Bin on C: and the File system name:</p> <ol style="list-style-type: none"><li>1. On the Windows client machine, right-click the <b>Recycle Bin</b> icon on the desktop and then click <b>Properties</b>.</li><li>2. Click the <b>General</b> tab.</li><li>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.</li><li>4. Choose the option <b>Do not move files to the Recycle Bin. Remove files immediately when deleted</b>.</li><li>5. Click <b>Apply</b>.</li><li>6. Repeat Step 3 through Step 5 for each remaining directory-mounted file system.</li><li>7. When finished, click <b>OK</b>.</li></ol>



---

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.</p> <p>This can happen by one of the following library maintenance type activities:</p> <ol style="list-style-type: none"> <li>1. If a tape drive is replaced with a new tape drive.</li> <li>2. If two tape drives are swapped within the library.</li> <li>3. If new tape drive(s) are added into location(s) that causes ACSLS to assign different locations to the previously existing drives.</li> <li>4. If a tape drive(s) are removed from location(s) that causes ACSLS to assign different locations to the remaining drives.</li> </ol> <p><b>Workaround:</b>            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 <code>st</code> command to rewind tapes may result in data loss.</p> <p><b>Workaround(s):</b></p> <ol style="list-style-type: none"> <li>Do not run any tools that could possibly issue a rewind.           <ul style="list-style-type: none"> <li>It may not be possible to determine this ahead of time.</li> <li>Gathering hardware info is imperative.</li> </ul> </li> <li>Rename the <code>/dev/st*</code> devices per “Product Alert 16”. Product Alerts are available on the Quantum Service and Support website at <a href="http://www.quantum.com/ServiceandSupport">www.quantum.com/ServiceandSupport</a>.           <ul style="list-style-type: none"> <li>They get created on the next reboot.</li> </ul> </li> <li>Stop SNSM before running any tools.           <ul style="list-style-type: none"> <li>Stopping TSM could take 10-15 minutes if very busy.</li> <li>Not optimal to stop/start TSM all the time.</li> </ul> </li> </ol>
	Cluster-Wide Central Control	<p>The StorNext Cluster-Wide Central Control file (<code>nss_cctl.xml</code>) is used to enforce the cluster-wide security control on StorNext nodes (client nodes, fsm nodes, and nodes running <code>cvadmin</code>). This file resides under <code>/usr/cvfs/config</code> 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>
All	Xsan	<p>Quantum recommends disabling sleep mode on the Xsan client to prevent the client from going into a sleep state while files are open in the StorNext file system or data is in the file system cache.</p> <p><b>Note:</b> If the client is in a sleep state while a file is open, then it will prevent other clients from opening the same file.</p>
	Xsan	<p>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.</p>
All	Labels	<p>Disks with existing non-StorNext labels may not show up in the StorNext GUI in order to protect non-StorNext disks from being accidentally overwritten. If you need to label a disk that is not visible in the StorNext GUI, use the <code>cvlabel</code> command to label the disk or use <code>cvlabel -U</code> to remove the existing label from the disks. (Refer to the <code>cvlabel</code> man pages for instructions on labeling and unlabeled drives.)</p> <p><b>Caution:</b> Modifying the label on an active non-StorNext disk can make the disk unusable. Proceed with caution.</p>

Operating System	Feature or Category	Description
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: <code>service cvfs restart</code></p> <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: <code>service snhamgr restart</code></p>
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"> <li>1. Remove the <code>/usr/cvfs/config/ha_peer</code> file from the secondary system.</li> <li>2. Reset the StorNext processes on the secondary system by running <code>/etc/init.d/stornext_web restart</code>.</li> </ol>

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"> <li>1. Complete the configuration changes</li> <li>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</li> <li>3. Restart CVFS on the secondary by issuing "service cvfs restart"</li> <li>4. Issue the command "cvadmin -e fsmlist" on the secondary MDC, and make sure that the output displays the FSM as "State: Blocked (waiting for MDC to become HA primary)"</li> </ol>
	HA	<p>Use caution when configuring the netmask for the HA Virtual Interface (VIP). 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><b>i Note:</b> 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>
All	Quotas	<p>Quotas can only be enabled or disabled by modifying the Quotas parameter of the file system configuration file. The CLI <b>snquota -L -Ffile-system</b> command informs you whether the file system has quotas enabled.</p> <p><b>i Note:</b> For a given file system, you can enable quotas in the StorNext GUI within the <b>Advanced Parameters &gt; Features</b> tab. Click <b>Quotas</b> to toggle the quotas function. For additional information on how to manage quotas, see <a href="#">Manage Quotas on page 148</a>.</p>

Operating System	Feature or Category	Description
All	fsretrieve	<p>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.</p> <p>Instead, use the recursive retrieve command. When you use this command the files under a directory are retrieved in batches, and more sorting is done to put files in tape order for increased performance. Run recursive retrieve by entering <code>% fsretrieve -R .</code></p>
All	Stripe group expansion	<p>As of StorNext 5, stripe group expansion, also known as bandwidth expansion, is not supported. The functionality works but has been deprecated and should not be used.</p> <p>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>
All	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"><li>• When there is one IP address associated with a NIC interface, the interface name alone is a sufficient identifier</li><li>• If there are multiple IP addresses associated with a NIC interface, one IP address is required in addition to the interface name</li><li>• On HA systems, the physical IP address is required if virtual IP is configured for the NIC interface. For additional information, see <a href="#">StorNext LAN Clients in HA Environments on page 348</a>.</li></ul>

Operating System	Feature or Category	Description
All	Truncation	By design, replication or deduplication must be completed before data files can be truncated if these files are associated with both a replication/dedup policy and a Storage Manager policy. Even if the Storage Manager policy is configured with the “Truncate Immediately” option, the truncation may not occur at store time unless the file has been replicated or deduplicated.
	DXi Virtual Tape Library Compatibility	Note the following recommendations and limitations for using DXi as a virtual tape library for StorNext: <ul style="list-style-type: none"> <li>• Recommended library emulation: “ADIC Scalar i2000”</li> <li>• Recommended tape drive emulation: “IBM LTO-x” or “HP LTO-x”</li> <li>• DDM (Distributed Data Mover): This feature is currently not supported due to lack of full SCSI3 support in DXi.</li> </ul>
	Affinities	When a file system with two affinities is to be managed by the Storage Manager, the GUI forces those affinities to be named tier1 and tier2. This will cause an issue if a site has an existing unmanaged file system with two affinities with different names and wants to change that file system to be managed. There is a process for converting a file system so it can be managed but it is non-trivial and time consuming. Please contact Quantum Support if this is desired. <p><b>i Note:</b> 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>
All	Converting file systems	StorNext does not currently support converting from a managed file system to an unmanaged file system.
	JournalSize Setting	For more information about JournalSize, refer to <a href="#">FSBlockSize, Metadata Disk Size, and JournalSize Settings on page 494</a> .
	Truncation	StorNext does not currently support running two truncation policies for different policy classes at the same time. Workaround: To avoid errors, do not run two truncation policies for different classes at the same time.
All	MDC	<b>⚠ Caution:</b> As the File System Manager (FSM) supports over 1000 clients (with more than 1000 file requests per client), the resource limits of your MDC may be exhausted with additional load from other processes. Exceeding the file descriptor limit will cause errors to your system. Quantum recommends you <b>not</b> run additional applications on the MDC.



Operating System	Feature or Category	Description
macOS	Trash Can	<p>If you are using the StorNext client software with the macOS operating system, disable the trashcan in the StorNext file system mount point.</p> <p><b>i Note:</b> If you do not disable the trashcan in the StorNext file system mount point, you might experience problems may which include failures when attempting to move folders to the trash in a managed file system.</p>

## To disable the trashcan, replace the directory with an empty file, as follows:

1. Empty the current **.Trashes** folder by running the command:

```
rm -rf <MNT>/ .Trashes
```

2. Create a file in place of the **.Trashes** folder by running the command:

```
touch <MNT>/ .Trashes
```

3. Remove permissions for the new **.Trashes** file by running the command:

```
chmod 0 <MNT>/ .Trashes
```



# Appendix B: Additional Replication and Deduplication Information

---

This appendix contains the following topics and provides detailed information about how replication and deduplication work, and about the underlying processes.

Replication Configuration File .....	522
Replication Terminology and Conventions .....	522
Copies in Replication Versus Copies and Versions in Storage Manager .....	522
Replication Target Directories .....	524
StorNext snpolicyd Policies .....	527
Replication Copies = 2 (Detail) .....	530
More About Replication Target Directories .....	532
Deduplication Overview .....	534
Deduplication and Truncation .....	536
Replication, Deduplication and Truncation .....	537
Replication, Deduplication and Storage Manager .....	538
The snpolicyd Debug Log .....	546

# 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 topic 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 topic 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 topic 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.

---

**i 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 snpolicyd storage policy.

The "Copies to Keep on Target" selection allows values of 1 through 16, and also a special case called in-place. We will not discuss the in-place selection in this section.

First, let's consider the case where "Copies to Keep on Target" is 2. Each time a replication occurs a new target directory is created. This target directory might have the same name as the previous target directory, but it is actually a new directory. The new directory reflects files added, deleted, and changed since the previous replication.

It is important to understand that in this example the target is a *new* directory. This has implications that might not be immediately obvious. For one thing, it means we cannot use the target directory in exactly the same way as we might use the source directory. Following is an explanation and examples.

## Example: Copies on Target = 2

In this example, we replicate source directory `/stornext/snfs1/photos`, a directory in file system `/stornext/snfs1`, to a target directory `/stornext/backup/photos` in file system `/stornext/backup`. (For this example it doesn't matter whether the two file systems are on the same node machine or on different machines.) Since we are keeping two copies on the target, we will usually have two directories on the target:

- `/stornext/backup/photos` - *most recent replication*
- `/stornext/backup/photos.1` - *previous replication*

When the next replication occurs, the following directory changes Take place:

- The previous replication `/stornext/backup/photos.1` is removed
- The most recent replication `/stornext/backup/photos` is renamed `/stornext/backup/photos.1`
- The new replication appears in `/stornext/backup/photos`

Now consider a Linux shell process that is executing inside directory `/stornext/backup/photos`. When the next replication occurs, the directory still exists but is named `/stornext/backup/photos.1`. If the Linux shell executes the command `ls -l`, the `ls` command lists the *previous* contents of `photos` - the directory now named `photos.1`.

When the replication after that occurs, the original directory is removed. When the shell executes `ls -l`, the command displays 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.

**Note:** 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.

**Note:** Using `exportrepcopy` avoids this second issue.

What if you want to change an existing source directory into a target directory? This can be done, but without special configuration care the original source policy assignment will be lost. A directory can have only one `snpolicyd` policy assigned to it (and all of the files and sub-directories 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
Configurable via the StorNext GUI?	Yes. Select the <b>Storage Policies</b> menu's <b>Storage Manager</b> option.	Yes. Select the <b>Storage Policies</b> menu's <b>Replication / Deduplication</b> option.
Configurable via the command line?	Yes. Use <code>fs</code> commands such as <code>fsaddclass</code> and <code>fsmodclass</code>	Yes. Use the <code>snpolicy</code> command.
Where are policy internals stored?	In Storage Manager Database. One database per machine.	In the managed file system, in a private directory.
Is the policy used across file systems?	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.
Functions?	Store (to tape or SDISK), retrieve, truncate files.	Deduplicate, replicate, truncate files.



Storage Policy Option	Storage Manager Policy	snpolicyd Policy
How are truncated files retrieved?	The entire file must be retrieved.	Only portions of the file containing needed regions may be retrieved.
Schedules?	fspolicy / schedules stored in Database.	Linux crontab scheduling.
Management daemon?	multiple fs_ processes	snpolicyd
Previous file versions recoverable?	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 the following command:

```
/usr/cvfs/config/snpolicy -dumppol/stornext/snfs1/photos
```

The command's output looks like the following:

```
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 you do not have to consider. 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)

This section provides additional detail about 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 the following:

```
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 displays 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 see one or more directories similar to the following:

```
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 see the following:

```
total 4160
lrwxrwxrwx 1 root root 72 Jan 26 10:13 0x1a4062.0 -> /snfs/sn2/___CVFS_
Handle.00047DA10CB84CF8000E0000000000000000000000001A412B
-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`.

## 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, you see the following:

```
./f.2m
policy: 1720449 inode: 1720468
flags: TAG
mtime: 2010-01-26 14:20:03.590665672 CST
ingest: 2010-01-26 14:20:03.590665672 CST
size: 2097152 disk blocks: 4096
seqno: 4 blk seqno: 2
offset: 0 length: 1048576 tag: D03281B0629858844F20BB791A60BD67
offset: 1048576 length: 1048576 tag: 12665A8E440FC4EF2B0C28B5D5B28159
./f.4m
policy: 1720449 inode: 1720470
flags: TAG
mtime: 2010-01-26 14:22:56.798334104 CST
ingest: 2010-01-26 14:22:56.798334104 CST
size: 4194304 disk blocks: 8192
seqno: 4 blk seqno: 4
offset: 0 length: 1048576 tag: D03281B0629858844F20BB791A60BD67
offset: 1048576 length: 1048576 tag: 12665A8E440FC4EF2B0C28B5D5B28159
offset: 2097152 length: 1048576 tag: 7F02E08B3D8C35541E80613142552316
offset: 3145728 length: 1048576 tag: 1FEC787120BEFA7E6685DF18110DF212
./g.4m
policy: 1720449 inode: 1720471
flags: TAG
mtime: 2010-01-26 14:23:28.957445176 CST
ingest: 2010-01-26 14:23:28.957445176 CST
size: 4194304 disk blocks: 8192
seqno: 5 blk seqno: 4
offset: 0 length: 1048576 tag: D03281B0629858844F20BB791A60BD67
offset: 1048576 length: 1048576 tag: DF54D6B832121A80FCB91EC0322CD5D3
offset: 2097152 length: 1048576 tag: 7F02E08B3D8C35541E80613142552316
offset: 3145728 length: 1048576 tag: 1FEC787120BEFA7E6685DF18110DF212
```

All three files have the same contents in the first megabyte starting at offset 0. The tag for that BLOB is D03281B0629858844F20BB791A60BD67, and that BLOB is stored only once in the blockpool. The second megabyte is the same for files f.2m and f.4m (tag 12665A8E440FC4EF2B0C28B5D5B28159) but file g.4m has a different BLOB in those bytes. The final 2 megabytes of files f.4m and g.4m are the same.

Remember that the above is an artificial example. In actual practice BLOBs do not line up on 1 MByte boundaries and are not all the same length.

## Enabling Deduplication

When creating or editing a policy through the StorNext GUI, select the **Deduplication** tab and make sure deduplication is enabled (On). If you use the `snpolicy dumpo1` option, you will see `dedup=on` in the output when the policy has deduplication enabled.



## Deduplication Modification Time

Note that in the "snppolicy -dumppol" output shown earlier we also saw dedup\_age=1m. This means the file may be deduplicated after it has not changed for at least one minute. If a file is being written its file modification time (mtime) will be updated as the file is being written. Deduplication age specifies how far in the past the modification time must be before a file can be considered for deduplication.

## Deduplication and Blockpools

If replication is used, a blockpool is required even if deduplication is not used in any policy on a machine. However, in this situation the blockpool does not store any BLOBs from any file system and can therefore be small: several megabytes is all that is needed.

If you enable deduplication on any policy in the machine, StorNext stores BLOBs in the blockpool and additional space is required. Make sure you have enough space to store file system data if you enable deduplication. You also need space for BLOBs in the blockpool if the machine contains replication target directories for deduplicated replication source directories on other machines.

The current StorNext release supports only one blockpool per machine. Any file system on the machine that needs a blockpool will use that one and only blockpool.

---

## Deduplication and Truncation

Let's look again at the directory in the previous section that has the three files f.2m, f.4m, and g.4m. Using the Linux command `ls -ls` displays the following 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 displays the total number of blocks (1024 bytes per block) contained in each file. The column before the date displays 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 displays the following:

```
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" signify?

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 are displayed like the following:

```
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` displays 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.

**i 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 the following:

```
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 the following:

```

/stornext/sn1/dd1/kcm2
policy: 18 inode: 1704267
flags: TRUNC TAG
mtime: 2010-02-05 11:01:25.960012768 CST
ingest: 2010-02-05 11:01:25.960012768 CST
size: 1388936 disk blocks: 0
seqno: 16 blk seqno: 3
offset: 0 length: 1388936 tag: 0D4093057370DB4FA7EF8162C23AE416
  
```

The line beginning with "flags:" contains the keyword TRUNC. This tells us that at least part of the file is not on disk in the file system and must be retrieved to be used.

If only snpolicyd is managing a directory, snpolicyd can truncate files when the snpolicyd rules are satisfied. This means that the deduplication has happened and the file is big enough and perhaps old enough. "Large enough" and "old enough" are determined by the deduplication policy parameters.

If only Storage Manager is managing a directory, the Storage Manager truncation rules determine whether and when a file can be truncated. This usually means that all Storage Manager copies have been made and that the file is large enough and old enough. "Large enough" and "old enough" are determined by the Storage Manager policy parameters.

If *both* Storage Manager and snpolicyd are managing a directory, Storage Manager must do the truncation. Storage Manager can only truncate a file when the Storage Manager rules are satisfied and any snpolicyd data copies have been completed.

You will know that both Storage Manager and snpolicyd are managing a directory if:

- The directory is a deduplicated directory and/or a replication source directory, and the directory is a Storage Manager relation point or is within a Storage Manager relation point.
- The directory is a replication target directory within a Storage manager relation point.

The table below summarizes some of the possibilities for snpolicyd managed directories and when truncation is allowed.

Snpolicyd State of the Directory	Directory is in an SM Relation Point	Directory is <i>Not</i> in an SM Relation Point
Non-deduplication Replication Source	SM can truncate when replications are complete	No truncation
Deduplication Replication Source	SM can truncate when deduplication has happened - even before replication	snpolicyd can truncate after deduplication
Deduplication Without Replication	SM can truncate when deduplication has happened	snpolicyd can truncate after deduplication
Target of Deduplication Source	Files are replicated as truncated (0 blocks). However, SM will eventually store each replicated file, causing it to be retrieved by snpolicyd on the target. Retrieved files must be truncated by SM and can only be truncated after all SM copies are made.	Files are replicated as truncated (0 blocks)

Snpolicyd State of the Directory	Directory is in an SM Relation Point	Directory is <i>Not</i> in an SM Relation Point
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:

**i Note:** These commands assume that the Storage Manager relation point and replication policy have already been configured.

```
faddrrelation directory_path -c sm_policy_name
spolicy -assignpolicy directory_path -inherit replication_policy_name
```

Remember that the directory should be empty before using fsaddrrelation, 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"><li>• Untruncated files are deduplicated per snpolicyd policy.</li><li>• SM truncated files will not be deduplicated until SM retrieval occurs. (snpolicyd will not retrieve the files from SM.)</li><li>• Once retrieved from SM, files cannot be re-truncated by SM until deduplication is complete.</li></ul> Therefore files may not all be deduplicated.
Snpolicyd deduplication policy that is also a replication source	About the same as above. SM truncated files are not deduplicated or replicated until something causes SM retrieval of the file. Thus there may be some files not deduplicated and not replicated.
Snpolicyd policy that is a replication source with no deduplication	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. Not all files will be deduplicated unless retrieved.

## 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:

- 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
```

Below is an example of a `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-rhel5464:
(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-rhel5464: 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 co[n]figure 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 .....	549
HA Internals: HAmom Timers and the ARB Protocol .....	550
Configuration and Conversion to HA .....	559
Managing High Availability in the StorNext GUI .....	562
Configuring Multiple NICs .....	563
High Availability Operation .....	564
HA Resets .....	570
HA Tracing and Log Files .....	572
Single (Singleton) Mode .....	572
Replacing a HA System .....	573
FSM Failover In HA Environments .....	573
Moving a HA Shared File System to a New Raid .....	576
Changing the IP Address of the MDC in an HA Pair .....	577
Installing StorNext Licenses For HA Configurations From the CLI .....	577

# 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:** Setting `haFsType` to **HaUnmonitored** disables the HA monitor timers used to guarantee against split brain. When two MDCs are configured to run as an HA pair but full HA protection is disabled in this way, it is possible in rare situations for file system metadata to become corrupt if there are lengthy delays or excessive loads in the LAN and SAN networks that prevent an active FSM from maintaining its branding of the ARB in a timely manner.

---

## 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 HAmon timer interval has its effect. The FSM waits until the interval period plus a small delta expires. The period began when the FSM branded the ARB. The FSM continues to maintain the brand during the delay so that another FSM cannot usurp before activation has completed. The connection count in the ARB is set to a very high value to block a competing usurpation during the activation process.

When an FSM stops, it attempts to quiesce metadata writes. When successful, it includes an indicator in its final ARB brand that tells the next activating FSM that the file system stopped safely so the wait for the HA timer interval can be skipped.

## LAN Connectivity Interruptions

When one MDC loses LAN connectivity, clients lose access to that MDC's active FSMs, which triggers elections to find other FSMs to serve those file systems. StorNext attempts to determine which node should have control, based on connectivity, but this effort results in a tie for the HaShared file system because each node gets one vote from itself as a client. In a tie, the activated shared FSM keeps control so long as it keeps branding its ARB.

Managed FSMs are not redundant, so having clients on those file systems does not break the tie. Similarly, unmanaged FSMs can fail over without an HA reset, so clients on those file systems will not break the tie for the shared file system either.

Therefore, a third client that has the shared file system mounted is necessary to break the tie that occurs between the two nodes. The third client makes it possible to determine which of the MDCs has the best connectivity to the LAN.

---

**i 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.



**i 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 FSMPM resets the timers for those FSMs for which usurpation is not in progress. Obviously, both server computers must be up and running StorNext for this to function.

This can postpone the impending HA reset for a while, but an election could occur if this goes on too long. It is important to quickly investigate the root cause of SAN or LUN delays and then engineer them out of the system as soon as possible.

## Primary and Secondary Server Status

Databases and management data for StorNext Storage Manager or the Linux GUI must also be protected against split-brain scenario corruption. Protection is accomplished by tying the startup of processes that modify this data with the activation of the shared file system.

Activating the shared file system leads to setting a Primary status in the local FSMPM, which is read and displayed by the `snhamgr` command. Primary status and the implicit Secondary status of the peer server are distinct from the Active and Standby status of the individual FSMs on the servers.

Unmanaged file systems can be active on either server. When an HA Cluster has no managed file systems and no shared file system, neither server computer has Primary status—they are equals.

## File System Types

HA is turned on by default for all StorNext distributions, but has no effect unless FSMs request to be monitored. File system monitoring is controlled by a file-system configuration item named `HaFsType`. Each file system is one of three types: `HaUnmanaged`, `HaManaged` or `HaShared`. The `HaFsType` value is read by FSMs to direct them to set up appropriate HAmon behaviors, and it is read by the FSMPM to control how it starts FSMs.

### HaUnmanaged

Each unmanaged-file-system FSM starts an instance of the HAmon timer for itself when it first brands the ARB. Before it changes any metadata, an activating FSM waits for the timer interval plus a small amount of time to elapse. The interval for a usurping FSM begins with the last time the FSM reads new data in the ARB from a previously active FSM.

Unmanaged FSMs can be active on either server in the HA Cluster. They can be usurped and fail over without a system reset if they exit before the timer expires. The timer interval for an active FSM restarts with each update of the ARB.

### HaManaged

Managed-file-system FSMs do not start HAmon timers, and they do not wait the HAmon interval when usurping. The FSMPMs only start Managed FSMs on the Primary server, so there is no risk of split-brain

scenario. In the event that a Managed FSM exits without having been stopped by the FSMPM, it is automatically restarted after a ten-second delay and activated. The `cvadmin` tool's `FSMlist` command displays the blocked FSMs on non-Primary servers. There can be zero or more `HaManaged` file systems configured.

## HaShared

The shared file system is an unmanaged StorNext file system that plays a controlling role in protecting shared resources. It has the same HA behavior as other unmanaged FSMs, but it also sets a flag that triggers an HA reset when the `cvfsioctl` device is closed. This happens when the process exits for any reason. However, if the shared file system has been unmounted from the active server before the FSM exits, the reset-on-close flag gets turned off. This allows ordinary shutdown of CVFS and Linux without resetting the server.

When the `HaShared` FSM finishes activation, it sets the Primary status in its FSMPM process.

Protected shared data resides on the shared file system. Since only one FSM can activate at one time, the Primary status is able to limit the starting of management processes to a single server, which protects the data against split-brain scenario.

The starting of `HaManaged` FSMs is also tied to Primary status, which guarantees collocation of the managed file-system FSMs and the management processes. The GUI's data is also shared, and the GUI must be able to manipulate configuration and operational data, which requires that it be collocated with the management processes.

## The `ha_peer` and `fsnameservers` File

StorNext HA server software uses peer-to-peer communication between servers and needs to know the peer's IP address. The `fsnameservers` configuration file is not a good source for the address because some installations configure the nameservers outside of the metadata servers. Instead, the following file provides that information:

```
<cvfs root>/config/ha_peer
```

Following are the uses of the peer IP address:

- Negotiating timer resets
- Comparing the 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-configurable modes to override the default behaviors of HA. Modes persist across reboots. Following are the modes and descriptions of their purpose:

Mode	Description
<b>default</b>	HA monitoring is turned on. When the peer server is not available for communication, it is assumed to be in <b>default</b> mode.
<b>single</b>	HA monitoring is turned off. The peer server must be communicating and in locked mode, or not communicating and certified as <code>peerdown</code> (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 <b>single</b> to <b>default</b> mode without stopping StorNext.
<b>config</b>	HA monitoring is turned off. The peer server must be communicating and in locked mode (recommended), or not communicating and certified as <code>peerdown</code> . The <b>config</b> mode is meant for re-configuration and other non-production service operations. When returning to production service and the <b>default</b> mode, StorNext must be stopped. This ensures that all StorNext processes are started correctly upon returning to <b>default</b> mode.

Mode	Description
<b>locked</b>	<p>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 <b>single</b> or <b>config</b> 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 <b>default</b> mode, which is necessary for avoiding split-brain scenario.</p> <hr/> <p> <b>Caution:</b> If a secondary MDC in <b>locked</b> mode is rebooted or powered down while the primary MDC is in <b>config</b> or <b>single</b> mode, <b>snhamgr</b> 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 <b>default</b> mode. To return the HA cluster back to the <b>config</b> state, check to ensure that the secondary MDC is either powered off and <b>peerdown</b> is set on the primary MDC or that the secondary MDC is running and its <b>snhamgr</b> mode is set to <b>locked</b>. Once this is verified, restart StorNext on the primary MDC and then set its <b>snhamgr</b> mode back to <b>config</b>.</p>
<b>peerdown</b>	<p>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.</p> <p>The mode is declared by the <b>peerdown</b> 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.</p> <p>The <b>peerdown</b> mode is changed to <b>default</b> mode with the <b>peerup</b> command. The <b>peerdown</b> and <b>peerup</b> 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.</p>
<b>ha_idle_failed_startup</b>	<p>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".</p>

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-peerdn**
- **single-locked**
- **config-peerdn**
- **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 (peerdn-\*), is the case when a node that is designated as peerdn 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**
- **peerdn-\***

## HA Manager Components

The following files and processes are some of the components of the HA Manager Subsystem:


File/Process	Description
snhamgr_daemon	If the <code>cnvt2ha . sh</code> 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 <code>cnvt2ha . sh</code> script has not been run. This is the interface for StorNext control scripts to regulate component starts.
<code>/usr/cvfs/install/.ha_mgr</code>	Stored mode value, which allows the single, config, locked, and peerdn modes to persist across reboots.
<code>SNSM_HA_CONFIGURED</code>	Environment variable that points to a touch file to indicate that <code>cnvt2ha . sh</code> has been run.


File/Process	Description
/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 <b>locked</b> mode, which stops StorNext on that server, then the primary server is placed in <b>config</b> mode and stopped, and then both servers are put in <b>default</b> mode with StorNext stopped.
start	Stop each server when there is a need and transition both servers to <b>default</b> 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.  <b>Note:</b> 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 <b>locked</b> or <b>peerdown</b> mode. Then, place the local server in <b>config</b> mode. The command must be run on the primary server, or either server when CVFS is stopped on both.

Command	Description
<code>clear</code>	Remove the file referenced by the <code>HA_IDLE_FAILED_STARTUP</code> environment variable. Run this after correcting any conditions that caused the previous failure of StorNext startup scripts.
<code>force smith</code>	Trigger an immediate HA reset if the local server is in <b>default</b> mode. This command is meant for use in health-monitoring scripts. The command is two words to make accidental firing less likely.   <b>Caution:</b> It is not recommended to use the <b>force smith</b> command to administratively failover a system in a production environment. The preferred method to gracefully failover the primary system to its secondary node is to simply stop CVFS and restart it after the secondary node has become primary. For example, on the node that is primary run: <pre># service cvfs stop</pre> Wait for the secondary to become primary, then run: <pre># service cvfs start</pre>
<code>peerdown</code>	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 <b>peerdown</b> mode, and to unset the mode before powering on the server.
<code>peerup</code>	Undo the <b>peerdown</b> mode. The command will fail if the local mode is <b>config</b> or <b>single</b> . Run this command before powering on the peer server. The local server will assume the peer is in <b>default</b> mode until the peer starts <b>snhamgr_daemon</b> communications.
<code>mode=&lt;modeval&gt;</code>	Set the mode of the local server to <b>modeval</b> . The mode is stored on the local server so that it persists across reboots. <ul style="list-style-type: none"><li>• <b>config</b>: Set the mode to <b>config</b>. The peer server must be in locked or <b>peerdown</b> mode.</li><li>• <b>default</b>: Set the mode to <b>default</b>. The peer server can be in any mode.</li><li>• <b>locked</b>: Set the mode to <b>locked</b>. The peer server can be in any mode.</li><li>• <b>single</b>: Set the mode to <b>single</b>. The peer server must be in <b>locked</b> (recommended) or <b>peerdown</b> mode.</li></ul>

See the *StorNext 5 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:



Type of StorNext Configuration	Description
<b>Windows</b>	<p>The StorNext GUI has a menu item for configuring HA: <b>Tools &gt; High Availability &gt; Convert</b>. It automatically inserts the <code>haFsType HaUnmanaged</code> 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 <code>ha_peer</code> 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.</p> <p>For more information about using the <b>Convert</b> menu option, see the StorNext online help and <a href="#">Converting to HA on page 344</a>.</p>
<b>Linux SNFS without GUI support</b>	<p>Each FSM configuration file must be given the <code>haFsType HaUnmanaged</code> configuration item, and the <code>ha_peer</code> 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.</p> <p>The FSM configuration files and <code>fsnameservers</code> 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 <code>haFsType</code> item from the FSM configuration files.</p>
<b>Linux Storage Manager with only unmanaged file systems</b>	See <a href="#">Conversion to HA below</a> for more information.
<b>Linux Storage Manager with managed and unmanaged file systems</b>	See <a href="#">Conversion to HA below</a> 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.

**i Note:** 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 sub-directory of the shared file system. Finally, it creates the following touch file to indicate that conversion has completed:

```
/usr/adic/install/.snsm_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 `.snsm_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_cctl.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`

---

**i 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

See [Manage on page 353](#).

---

# 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 on page 555](#).

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.

---

**i 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

---

**i Note:** The **Windows and Linux SNFS Installations Without the HaShared File System** section is applicable to unmanaged file systems only.

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.

---

**i 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, `FSMlist`, `fsnameservers` and the optional `ha_smith_interval` files must be identical on both servers.

When stopping StorNext on one of these servers, all FSMs will stop. Standby FSMs on the redundant server will activate and resume serving their file systems. No HA reset will occur if every stopping FSM exits before their `HAmon` timer expires (after the final ARB brand update).

During operation, individual file systems could potentially fail over between servers as the result of a hardware or software failure or because an operator has directed it by the `fail` command in the `cvadmin` tool. The `fail` command can be used for load balancing after the HA cluster has completed startup.

When making file system configuration changes, one of the servers should be stopped and its FSM configurations deleted. This eliminates the possible mistake of asymmetric configurations. After making all the configuration changes on one server and updating the file systems with those new configurations, the configuration files must be copied to the redundant server. Then, the cluster can be operated with redundant servers again.

When updating StorNext software, refer to the release notes and the StorNext Upgrade Guide for current update instructions. These documents address any special considerations for HA according to the scope of the changes in the software release.

## Linux SNMS and SNFS Installations with the HaShared File System

The HaShared file system is required for SNMS and GUI-supported installations. The shared file system holds operational information for those components, which must be protected against split-brain corruption. The additional complexity this entails is simplified and automated by the HA Manager Subsystem.

### Touch Files that Control StorNext and HA

Environment variables defined in the `/usr/adic/.profile` and `/usr/adic/.cshrc` files reference touch files that are used in StorNext to track state transitions. Some of these are unique to HA configurations. The variables and their values are as follows:

1. `ACTIVE_SNFS_SERVER=/usr/adic/install/.active_snfs_server`

The file is created following activation of the HaShared file system to designate the local node as the Primary server.

2. `HA_STARTING_UP=/usr/cvfs/install/.ha_starting_up`

The file is created at the start of the 'service cvfs start' script, and removed at the end of the activation script for the HaShared file system. If it is found before it is created, that causes the creation of the `HA_IDLE_FAILED_STARTUP` file as described in the next item.

3. `HA_IDLE_FAILED_STARTUP=/usr/cvfs/install/.ha_idle_failed_startup`

When the `HA_STARTING_UP` file exists as the 'service cvfs start' starts, it is assumed that the previous attempt to start failed before completing, possibly because an HA reset occurred. The `HA_IDLE_FAILED_STARTUP` file is created to block future attempts to start, and the current script exits. This avoids an infinitely looping series of startup attempts, and allows an administrator to log in and correct problems. The HA Manager reports the existence of this file as a mode, and offers the clear command for removing the file.

4. `SNSM_HA_CONFIGURED=/usr/adic/install/.snsm_ha_configured`

The file is created by the `cnvt2ha.sh` script to indicate that the system has been converted to HA. Its existence allows the `snhamgr_daemon` to run.

5. `START_SNFS_ONLY=/usr/adic/install/.start_snfs_only`

The file is created by running one of the following commands: '`/usr/adic/bin/adic_control startonly snfs`' or '`/usr/cvfs/bin/DSM_control startonly`'.

Its existence indicates to the `snactivated` script that Storage Manager components are not to be started. The file is removed by using any of the following commands: '`DSM_control stop`', '`service cvfs stop`', or '`adic_control stop snfs`'.

## Starting a Single StorNext HA Server for Production

The command `service cvfs start` sets in motion a sequence of events that result in the starting of all the Storage Manager components.

---

**i 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 FSMPM
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 and 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*.

---

**i 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. If necessary, power down the decommissioning server.
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, verify the `/usr/cvfs/config/ha_peer` file contains the IP address of the new server; update if required.
8. On the working server, run the following two commands in order:

```
snhamgr mode=default  
snhamgr peerup
```

9. Run the conversion script on the new server as follows:

```
/usr/adic/util/cnvt2ha.sh secondary <shared fs> <peer IP address>
```

For additional information, see [Single \(Singleton\) Mode on page 572](#) and [Replacing a HA System on page 573](#).

## 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 command `snamgr force smith` by a script or manually by an administrator. The `smithlog` file is written by the `fsmpm` process, so there would not be an entry in the file when an `fsmpm` exit results in an HA Reset.

---

 **Caution:** It is not recommended to use the `force smith` command to administratively failover a system in a production environment. The preferred method to gracefully failover the primary system to its secondary node is to simply stop CVFS and restart it after the secondary node has become primary. For example, on the node that is primary run:

```
# service cvfs stop
```

Wait for the secondary to become primary, then run:

```
# service cvfs start
```

## 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  $5 \times 2 = 10$ ), an HA Reset occurs.
- Turn off tracing with `'cvdbset - all'`

## HA Resets of the Second Kind

The second method of HA Reset can occur on shutdown of CVFS if there is an unkillable process or delayed process exit under the HaShared file system mount point. This will keep the file system from being unmounted. The smithlog entry indicates when this has happened, but does not identify the process.

## HA Resets of the Third Kind

The third method of HA Reset is the most common. It occurs when the snactivated script for the HaShared FSM experiences an error during startup. The current implementation invokes the `'snhamgr force smith'` command to allow the peer MDC an opportunity to start up StorNext if it can. A similar strategy was used in previous releases. In this release, the failure to start will cause the `/usr/cvfs/install/.ha_idle_failed_startup` touch file to be created, and this will prevent startup of CVFS on this MDC until the file is erased with the `'snhamgr clear'` command.

## Using HA Manager Modes

The snhamgr rules for mode pairings are easier to understand by following a BAAB strategy for transitioning into and out of config or single mode. In this strategy, B stands for the redundant node, and A stands for the node to be placed into config or single mode. Enter the desired cluster state by transitioning B's mode first, then A's. Reverse this when exiting the cluster state by transitioning A's mode, then B's.

For the configuration-session example, place B in locked mode, then place A in config mode to start a configuration session. At the end of the session, place A in default mode, then place B in default mode.

For the single-server cluster example, shut down Linux and power off B, then designate it peerdown with the `'snhamgr peerdown'` command on A, then place A in single mode. At the end of the session, place A in default mode, then designate B as up with the `'snhamgr peerup'` command on A, then power on B.

---

# HA Tracing and Log Files

The following log files contain HA related debugging information:

Log File	Description
<code>/usr/cvfs/debug/ha_mgr.out</code>	Log messages from the <code>snhamgr_daemon</code>
<code>/usr/cvfs/debug/hamgr_cmds_trace</code>	Output from commands run by the <code>snhamgr_daemon</code> . Typically, several commands are run simultaneously. Their output becomes intertwined. This is normal.
<code>/usr/cvfs/debug/snactivated.&lt;fs name&gt;.log</code>	Output from the <code>snactivated.pl</code> command per file system.
<code>/usr/cvfs/debug/nssdbg.out</code>	Log messages from the <code>FSMPM</code> daemon. HA related messages include: the <code>HAMon</code> timer interval, anomalies in negotiations for the resetting of <code>HAMon</code> timers, setting of Primary status, activations of FSMs etc.
<code>/usr/cvfs/data/&lt;fs name&gt;/log/cvlog</code>	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 <code>snactivated</code> script etc.
<code>/usr/adic/HA/cnvt2ha.sh.log</code>	Output of the <code>cnvt2ha.sh</code> script.
<code>/var/log/messages</code>	Mounts of <code>cvfs</code> file systems.
<code>/usr/cvfs/debug/smithlog</code>	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.

---

## Replacing a HA System

To replace a HA server, refer to [Replacing an MDC in an HA environment on page 662](#).

---

## 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.

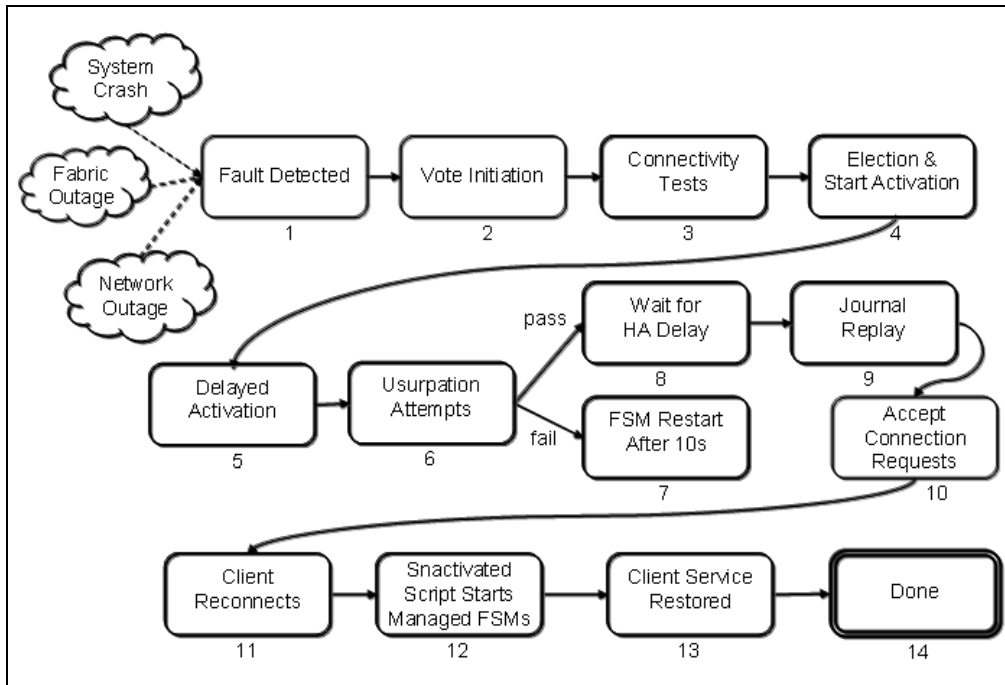
The StorNext system exclusively opens files on the HaShared file system, but assumes that only the Primary MDC does this and waives the delay for that one file system. Quantum advises against running user processes other than StorNext processes on HA MDCs for performance, reliability and availability 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 displays 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 61:** FSM Failover in an HA Cluster



Not shown in this diagram are the state transitions of the peer MDC when it incurs an HA Reset. The HA Reset is not directly tied to the failover of control to a standby FSM, but rather the detection of a loss of services triggers failovers. The HA Reset may come before or after the loss of services, or not at all. It is only important to know that by the end of state 8, no FSM on the peer MDC is controlling the arbitration block (ARB). The HA Reset mechanism guarantees that to be true.

The example failures shown here (System Crash, Fabric Outage, Network Outage) can result in a failover. Typically, the loss of heartbeat from the peer MDC's FSMPPM is the first indication that an HA Reset has occurred.

1. **Triggering Event:** The loss of heartbeat is detected and triggers an election at approximate time T3.5 seconds. Note that failover of a single unmanaged file system could also be forced with the `cvadmin` command without causing an HA Reset.
2. **Vote Initiation:** A quorum-vote election is started where the clients of the file system identify the best-connected MDC having a standby FSM for the file system.
3. **Connectivity Tests:** Each live client runs a connectivity test sequence to each server. Connections are tested in less than .5 seconds per server, when successful, and can be repeated up to four times (two seconds) when unsuccessful. At completion of the election, the time is approximately T5.5.
4. **Election and Start Activation:** The election is completed, and an activation message is sent to one server's standby FSM.
5. **Delayed Activation:** When a server has active FSMs, its FSMPPM process sends a request to the FSMPPM of its peer server to ask if the corresponding Standby FSMs are being activated. If not, the local FSMPPM can reset the HA timer of that file system's active FSM, which reduces the chance of an unnecessary HA Reset. When the peer FSMPPM gives permission, it is constrained from activating the

standby FSM for two seconds. Step 5 is for that delay of up to two seconds. The delay completes at approximately T6.5.

6. **Usurpation Attempts:** To prevent false takeovers, the ARB is polled to determine whether another FSM is active and must be "usurped". Usurpation is averted if the activating FSM detects activity in the ARB and its vote count does not exceed the active FSM's client-connection count. A typical successful poll after an HA Reset lasts two seconds. When the previously active FSM exits gracefully, the usurpation takes one second.

The activating FSM then performs a sequence of I/Os to "brand" the arbitration block to signal takeover to the peer FSM. An active FSM is required to exit when it sees that its brand has been overwritten. These operations take two seconds. The 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.

State	Failover Timing Estimates (seconds)	
	Optimal	Common
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
<b>Total</b>	<b>5.5</b>	<b>27.5</b>

---

## Moving a HA Shared File System to a New Raid

 **Caution:** Contact **Quantum Technical Support** as 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.

---

# Changing the IP Address of the MDC in an HA Pair

**⚠ Caution:** Contact **Quantum Technical Support** as 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.

---

## Installing StorNext Licenses For HA Configurations From the CLI

### Installing StorNext Licenses From the CLI for a Running HA Configuration

Sometimes it may be desired to update a new `license.dat` file from the CLI. If StorNext is running, the `/usr/cvfs/config/license.dat` file may be updated by copying the updated `license.dat` file into place on the system that is currently operating as the primary node. Once in place, StorNext services will need to be restarted for some of the features to recognize the new license file.

1. Login to the command line of the MDC node you will use to update the license.
2. To determine if the MDC you have logged into is the node currently acting as primary, execute the following:

```
# snhamgr status
```

3. Verify the output displays:

```
:default:primary:default:running:
```

4. Update `/usr/cvfs/config/license.dat` with the new `license.dat` file.
5. From the command line of the MDC node currently acting as primary, execute the following:

```
# snhamgr stop
```

6. Execute the following:

```
# snhamgr start
```

## Installing StorNext Licenses From the CLI for a Stopped HA Configuration

If StorNext is not running, the cluster will need to be placed into config mode using the `snhamgr config` command. Once the license is updated, StorNext will need to be started while in config mode to force the license to be updated. This can be done on either MDC node. Once started, StorNext will need to be placed back into default mode to resume normal operation.

---

**i Note:** If the system is not started in config mode, the updated `license.dat` will be replaced by the old version of the `license.dat` file.

1. Login to the command line of the MDC node you will use to update the license.
2. Update `/usr/cvfs/config/license.dat` with the new `license.dat` file.
3. From the command line, execute the following:

```
# snhamgr config
```

4. Execute the following:

```
# snhamgr start
```

---

**i Note:** Exiting config mode will cause StorNext services to automatically be stopped and restarted.



## Appendix D: Web Services API

---

**i Note:** For the latest Web Services documentation, see the *StorNext 5 Web Services Guide* available online at <http://www.quantum.com/sn5docs>. As of StorNext 5 release 5.0, the Web Services API documentation is being maintained in the *StorNext 5 Web Services Guide*.



# 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.

**i Note:** The StorNext Offline Notification feature may not function properly on systems using Windows 10. See [Configure StorNext Offline Notification with Anti-virus or Anti-malware Software on page 654](#).

This appendix contains the following topics:

Truncation Overview .....	580
Space Management .....	582
Disabling Truncation .....	584
Common Problems .....	585
Miscellaneous Usage Notes .....	586
Scheduling Truncation Manually .....	586

---

## 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. For more information, see [Normal Truncation below](#), [Immediate Truncation below](#), and [Daily Truncation on the next page](#).

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 relocation

## 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:

- **LOSPACE Cycle.** See [LOSPACE Cycle below](#).
- **Emergency Cycle.** See [Emergency Cycle on the next page](#).

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 relocation is:

```
fspolicy -r -y <mountpoint> -e -a <affinity>
```

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:

- a. By using truncation feature locking. See [Truncation Feature Locking below](#).
- b. By running commands that disable truncation. See [Disable Truncation Commands on the next page](#).

## Truncation Feature Locking

Truncation operations can be locked (in other words, prevented from running), by using the **fsschedlock** command. The feature name for each truncation operation is:

Parameter	Description
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 execute 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 `mintrunc` time to 1h (or 60m). Next, you must schedule via cron a truncation policy to run every half hour.

**i Note:** Even with policies running every half hour, a file may wait up until the time between policies beyond the `mintrunc` 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

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 below](#).

## 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
```

**i 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 the next page](#) 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.

---

**i 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.

---

**i 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 a deep dive of security as it pertains to StorNext, and contains the following topics:

StorNext Security (for pre-StorNext 5 release 5.4.0 Systems) .....	591
StorNext Security .....	597
Central Control .....	612
Limitations .....	613
Example .....	614
Cross-Platform Permissions .....	615
Config (.cfg) File Options .....	616

---

### StorNext Security (for pre-StorNext 5 release 5.4.0 Systems)

**i Note:** The information in this section is applicable for releases up to StorNext 5 release 5.3.x (in other words, prior to StorNext 5 release 5.4.0). Refer to [StorNext Security on page 597](#) for the new information, beginning with StorNext 5 release 5.4.0.

There are two predominate security models in legacy 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 [General Operating Guidelines and Limitations on page 505](#).

StorNext does not implement any of the Unix ACLs models or the NFSv4 ACLs model.

## ACLs on Windows (for pre-StorNext 5 release 5.4.0 Systems)

**i Note:** The information in this section is applicable for releases up to StorNext 5 release 5.3.x (in other words, prior to StorNext 5 release 5.4.0). Refer to [StorNext Security on page 597](#) for the new information, beginning with StorNext 5 release 5.4.0.

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](#)".

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.

If the file system configuration global `UnixFabricationOnWindows` is set to **true**, the Unix UID and GID values are populated based on the GUIDs of the Active Directory records.

If `UnixFabricationOnWindows` is set to **false**, the Unix UID and GID values are populated from the RFC2307 mappings in Active Directory. If the SID for the logged in user does not have a UID configured in Active Directory, a newly created file will receive a UID equal to the value of the file system configuration setting `UnixNobodyUidOnWindows`. If the SID for the logged in user does not have a GID configured in Active Directory, a newly created file will receive a GID equal to the value of the file system configuration setting `UnixNobodyGidOnWindows`.

The UNIX modes are controlled by the following file system configuration parameters:

- `UnixDirectoryCreationModeOnWindowsDefault`
- `UnixFileCreationModeOnWindowsDefault`

StorNext allows one set of values for all users of each file system.

---

**i Note:** The value of the file’s Windows read-only attribute does not affect the file’s Unix mode.

---

**i Note:** Administrators can manually change these values in the file system configuration file on the server or use the Windows or Web GUI.

For additional information, refer to the `snfs_config(5)` man-page or the following sections in the Windows Help.

- User ID Mapping Overview
- Windows Active Directory Config
- Apple/XSAN Fabricated ID’s
- Unix Permissions Background

## Enable ACLs on a StorNext MDC (for pre-StorNext 5 release 5.4.0 Systems)

**i Note:** The information in this section is applicable for releases up to StorNext 5 release 5.3.x (in other words, prior to StorNext 5 release 5.4.0). Refer to [StorNext Security on page 597](#) for the new information, beginning with StorNext 5 release 5.4.0.

ACLs were introduced on Mac OS X 10.3 (Tiger) systems. Mac and Windows systems share equivalent ACLs implementation. ACLs must be enabled on the StorNext file system. To enable ACLs on an existing file system, perform the steps below.

### Enable ACLs on a Linux MDC

1. Using the StorNext GUI, on the **Configuration** menu, click **File Systems**.
2. Select the file system from the table, and then click **Edit...**
3. Click **Advanced Parameters** to expand the section.
4. On the **Features** tab, click **Enforce ACLS**.
5. Click **Apply** to save the file system changes.

### Enable ACLs on a Windows MDC

1. Stop the file system.
2. Depending on your Windows operating system, navigate to and execute the **File System Cfg (Advanced)** server configuration utility.
3. Select an existing file and click **Open Existing Configuration** to edit an existing configuration file.
4. On the **Global Settings** tab, under **XSAN Client Behavior**, select the **Enforce ACLs** check box to enable ACLs.
5. If a check-mark already exists in the **Enforce ACLs** check box, **Enforce ACLs** is enabled.
6. Click **OK** to save the file system configuration.
7. Click **Save**, and then click **Yes** when prompted to overwrite the file system configuration.
8. Start the file system.

The `chmod(1)` and `ls(1)` commands have been modified to handle ACLs. There is also a library API for applications, `acl(3)` that allows programs to operate on ACLs.

For a detailed description of Mac OS X ACLs, see “Security Overview: Permissions” from Apples web sites and click on ACLs.

ACLs take precedence over regular UNIX permissions. If no ACE match is found for a user's requested access, UNIX permissions are checked. Therefore, a user may not match any ACE but still have access if UNIX permissions allow.

Each ACE on Mac OS X has the same 13 possible permission bits as a Windows ACE:

Directories	Files
Search Through	Execute File
List Contents	Read Data
Read Attributes	Read Attributes
Read Extended (named) Attributes	Read Extended (named) Attributes
Create Files	Write Data
Create Subdirectories	Append Data
Write Attributes	Write Attributes
Delete Subdirectories and Files	
Delete this Directory	Delete this Directory
Read Permissions (ACL)	Read Permissions (ACL)
Change Permissions (ACL)	Change Permissions (ACL)
Take Ownership	Take Ownership

Inheritance on Mac OS X is similar but does vary from Windows propagation and inheritance. Each ACE applied to a directory can be “propagated” by indicating one of 4 tags:

1. **file\_inherit**: Propagate this ACE to files created in this directory.
2. **directory\_inherit**: Propagate this ACE to subdirectories created in this directory.
3. **limit\_inherit**: After propagating this ACE to a new subdirectory, do not let its subdirectories inherit this ACE.
4. **only\_inherit**: Do not apply this ACE to this directory, just to files and/or directories created below it.

The “limit\_inherit” exists in Windows as a check box when creating an ACE on a folder that propagates. The mapping of the 3 remaining tags to the 7 Windows propagation pull down menu options are as follows:

Windows	Mac OS X
This folder only	(none)
This folder, subfolders, and files	directory_inherit, files_inherit
This folder and subfolders	directory_inherit
Subfolders and files only	files_inherit
Subfolders and files only	files_inherit, directory_inherit, only_inherit

Windows	Mac OS X
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. Windows 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.

---

## StorNext Security

There are two predominate security models in modern file systems:

- Unix permission bits
- Access Control Lists (ACLs)

StorNext supports both models and when they are used depends on the client platform and StorNext file system configuration settings. Most notably, the `securityModel` configuration variable may be set to one of the following values:

- **ACL:** With the ACL security model, permissions are enforced on Windows systems based on ACLs. On all other platforms, the security check is based on a combination of ACLs and Unix permission bits.
- **Unixpermbits:** With the Unixpermbits security model, permissions are enforced on all platforms (including Windows) based on Unix permission bits.
- **Legacy:** With the Legacy security model, permissions are enforced on Linux and Unix systems based on Unix permission bits.
  - Mac systems use either ACLs or Unix permission bits or a combination of the two, depending on other StorNext file system configuration settings.
  - Windows systems use ACLs when the **Windows Security** setting is enabled; otherwise, Windows systems do not perform security checking.

The behavior of the various models is summarized in [Table 8 on the next page](#).

**Table 8:** StorNext Permission Method Summary

	Windows Native Clients and Windows SMB Clients to a Windows Server	Xsan	Linux/Unix	SMB clients to a Samba server	NFS
securityModel = acl	ACLs	Combination of ACLs and Unix Permission Bits	Combination of ACLs and Unix Permission Bits	ACLs <sup>1</sup> but fall back to Unix Permission Bits if no ACL	Combination of ACLs and Unix Permission Bits
securityModel = unixPermBits	Unix Permission Bits	Unix Permission Bits	Unix Permission Bits	Unix Permission Bits	Unix Permission Bits
securityModel = legacy, windowsSecurity = false	No Access Check	Unix Permission Bits	Unix Permission Bits	Unix Permission Bits	Unix Permission Bits
securityModel = legacy, windowsSecurity = true, enforceACLs = false	ACLs	Unix Permission Bits	Unix Permission Bits	ACLs <sup>1</sup> but fall back to Unix Permission Bits if no ACL	Unix Permission Bits
securityModel = legacy, windowsSecurity = true, enforceACLs = true	ACLs	Combination of ACLs and Unix Permission Bits	Unix Permission Bits	ACLs <sup>1</sup> but fall back to Unix Permission Bits if no ACL	Unix Permission Bits

<sup>1</sup> ACLs only supported for Samba when using the StorNext NAS stack and enabling ACLs for a share

The selection of security model is made on a per file system basis. After a file system has been created, in many cases it is possible to change the security model without reinitializing the file system. However, there are a few exceptions:

- It is **not** possible to transition from the ACL model to the Unixperm bits model.
- It is **not** possible to transition to the Unixperm bits model from the Legacy model if **Windows Security** is enabled.

Which model is appropriate will depend on security requirements. Unix permission bits are simpler but not as flexible as ACLs. Both the ACL and Unixperm bits security models allow a single type of file permissions (ACLs or Unix permission bits) to be enforced uniformly across all platforms in a heterogeneous environment providing a consistent experience. However, in some cases, due to identity mapping constraints, the legacy model must be used. In such environments, both Unix permission bits and ACLs must be separately maintained.

## Permission Enforcement Details

Permission enforcement varies based on the platform and StorNext configuration parameters.

### File Permission Enforcement on Linux, Unix, and Mac OS Environments

When using the `unixpermbits` or legacy Security Model, Linux and Unix systems use the standard Posix permission bit model. When ACLs are enabled for all systems by using the `acl` Security Model or on Mac OS systems using the `enforceAc1s` file system configuration variable, these systems will first check whether the given file or directory contains an Access Control Entry (ACE) that applies to the user and operation being performed. If so, the ACE will be used to determine access. Otherwise, the Unix permission bits are used.

**i Note:** While Unix permissions have only 3 bits (READ/WRITE/EXECUTE), the ACEs contained in an ACL may contain up to 14 bits. Not all of the ACE permission flags have meaning outside of Windows, but many do, depending on the file operation being performed. [Table 9 below](#) describes how these flags are used for common file system operations on Linux, Unix, and Mac OS platforms.

**Table 9:** Effect of ACE Permission Flags on Non-Windows Platforms when ACLs are Enabled

Permission Flag	DENY on Directory	ALLOW on Directory	Deny on FILE	ALLOW on FILE
Read Data/List Folders	Disallows <code>readdir(3)</code>	Allows <code>readdir(3)</code>	Disallows <code>open(2)</code> for READ	Allows <code>open(2)</code> for READ
Write Data/Create Files	Disallows <code>creat(2) / open(2)</code> that results in a new file	Disallows <code>creat(2) / open(2)</code> that results in a new file	Disallows <code>open(2)</code> for WRITE	Allows <code>open(2)</code> for WRITE
Append Data/Create Subfolders	Disallows <code>mkdir(2)</code>	Allows <code>mkdir(2)</code>	Disallows <code>open(2)</code> for APPEND	Allows <code>open(2)</code> for APPEND
Read Extended Attributes	Disallows <code>getxattr(2)</code>	Allows <code>getxattr(2)</code>	Disallows <code>getxattr(2)</code>	Allows <code>getxattr(2)</code>
Write Extended Attributes	Disallows <code>setxattr(2)</code>	Allows <code>setxattr(2)</code>	Disallows <code>setxattr(2)</code>	Allows <code>setxattr(2)</code>



**Table 9:** Effect of ACE Permission Flags on Non-Windows Platforms when ACLs are Enabled

Permission Flag	DENY on Directory	ALLOW on Directory	Deny on FILE	ALLOW on FILE
Execute/Traverse Folders	Disallows path traversal through a directory	Allows path traversal through a directory	Denies binary/script execution	Allows binary/script execution
Delete Subfolders and Files	Disallows deletion of immediate descendant, unless target object ALLOWS delete	Allows deletion of immediate descendant	N/A	N/A
Read Attributes	Disallows stat(2)	Allows stat(2)	Disallows stat(2)	Allows stat (2)
Write Attributes	Disallows utimes(2)	Allows utimes(2)	Disallows utimes(2)	Allows utimes(2)
Read Permissions	Disallows stat(2) and displaying ACL (e.g. snacl -l)	Allows stat(2) and displaying ACL	Disallows stat(2) and displaying ACL	Allows stat (2) and displaying ACL
Write Permissions	Disallows chmod(2) and ACL modifications	Allows chmod(2) and ACL modifications	Disallows chmod(2) and ACL modifications	Allows chmod(2) and ACL modifications
Delete	Disallows rmdir(2) unless delete_child is permitted	Allows rmdir (2)	Disallows unlink(2) unless delete_child is permitted by parent	Allows unlink (2)
Change Owner	Disallows ACL modifications	no effect	Disallows ACL modifications	no effect
Synchronize	no effect	no effect	no effect	no effect

## File Permission Enforcement on Native StorNext Windows Clients and Windows SMB Clients Attached to a Windows SMB Server

When using the unixpermbits Security Model, Windows systems map Unix permission bits to Windows rights as shown in [Table 10 on the next page](#).

**Table 10:** Windows Rights when using Unixpermbits Security Model

Unix Permission Bit	Granted Windows Rights - Regular Files	Granted Windows Rights - Folders
READ	Read Data	List Folders
	Read Extended Attributes	Read Extended Attributes
	Read Attributes	Read Attributes
	Read Permissions	Read Permissions
WRITE	Write Data	Create Files
	Append Data	Create Subfolders
	Write Extended Attributes	Write Extended Attributes
	Write Attributes	Write Attributes
	Delete	Delete
		Delete Subfolders and Files
EXECUTE	Execute	Traverse Folders

The ability to update permissions (which for ACLs would correspond to the **Change Permissions** right) is granted to the owner of the file or to the Domain Administrator. The ability to change file ownership (which for ACLs would correspond to the **Take Ownership**) is granted only to the Domain Administrator.

When ACLs are enabled, Windows systems use the ACL on a file or folder for determining access and the Unix permission bits are ignored.

ACL enforcement on Windows systems works as follows:

The ACL on each file/folder contains zero or more ACEs that specify how a user or group can access or not access the file or folder. Each ACE specifies a principal (in other words, a user or group), the type (ALLOW or DENY), permission flags and inheritance flags. [Table 11 below](#) provides the possible permission flags in each ACE.

**Table 11:** Possible Permission Flags in each ACE

Folders	Files
Traverse Folder	Execute File
List Folder	Read Data
Read Attributes	Read Attributes
Read Extended Attributes	Read Extended Attributes
Create Files	Write Data

**Table 11:** Possible Permission Flags in each ACE

Folders	Files
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

All file and folder access is determined by matching the accessing user and associated groups to the ACL of the object being accessed. The ACEs in the ACL are compared in order with the accessing user and groups 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.

If a file or folder does not have an ACL, because, for example, it was created on a system that didn’t support them, then on the Windows client, StorNext will automatically synthesize one based on the parent folder. If the parent does not have an ACL, StorNext searches backward all the way to the root for an ACL to use. If the Security Model is set to legacy and Windows Security is disabled, then Windows clients will not perform a security check.

## File Permission Enforcement on SMB Clients Using a Samba Server Running StorNext NAS

When using the unixpermbits or legacy Security Model, SMB clients use the standard Posix permission bit model. When ACLs are enabled, these systems check to see whether given file or directory contains an ACE that applies to the user and operation being performed. If so, the ACE will be used to determine access. Otherwise, the Unix permission bits are used.

---

**i Note:** While this sounds similar to how file permissions are enforced for Linux, Unix, and Mac OS environment, there is a subtle difference.

# Display and Modification of File Permissions

## Display and Modification of Unix Permission Bits

On non-Windows systems, Unix permission bits and file ownership are displayed using the **ls** command and modified using the **chmod**, **chown**, and **chgrp** commands.

On native Windows StorNext clients, Unix permission bits can be displayed and modified using the StorNext **cvstat** command from a DOS shell. When using the unixpermbits security model with either **mdc** or **ldap** identity mapping, Unix permission bits may also be displayed and modified using the **Windows Security** tab.

---


**i Note:** The **Windows Security** tab is accessed in Windows Explorer by right-clicking on a file or folder, selecting **Properties** and navigating to the **Security** tab

In this case, StorNext synthesizes an ACL based on the Unix permission bits. The following example illustrates Unix permissions as displayed from a Linux system.

```
# ls -lt testfile
-rw-r--r-- 1 jdoe domain users 0 Aug  7 16:23 testfile
```

[Figure 62 on the next page](#) illustrates the same Unix permissions displayed as a synthesized ACL in Windows Explorer.

**Figure 62:** Unix permissions Displayed as a Synthesized ACL in Windows Explorer



When using the unixpermbits security model, the **Windows** security tab displays three Access Control Entries (ACEs).

- One for the owner of the file
- One for the group
- One for Everyone

The displayed ACE permissions correspond other Unix permissions on the file. In the figure, the Unix mode on the file is 0644 which means that the owner of the file (jdoe) has READ and WRITE permission for the file, the group (domain users) has READ permission, and finally Everyone has READ permission which corresponds to the unix permission bits for “others”.

**i Note:** Permissions can be modified by clicking the **Edit** button. Also, it is possible to change the ownership of the file by clicking the **Advanced** button and navigating to the **Owner** tab. However, this operation may only be performed by the Domain Administrator.

**i Note:** If the UID cannot be uniquely mapped to a Windows user, the **Security** tab displays the owner as **CREATOR OWNER**. In this case, modifying permissions from Explorer is not allowed. Similarly, if the GID of the file cannot be mapped to a Windows group, the **Security** tab displays the group as **CREATOR GROUP**.

## Display and Modification of ACLs

[Table 12 below](#) provides tools (which vary by platform and protocol) used for displaying and modifying ACLs.

**Table 12:** Tools Used for Displaying and Modifying ACLs

	Display ACLs	Modify ACLs
Windows	Explorer security tab icacls cacls	Explorer security tab icacls cacls
Mac OS	ls -le	chmod +a chmod -a chmod =a etc.
Linux and Unix Native StorNext Clients	snacI -l	snacI +a snacI -a snacI =a etc.
Linux NFSv4 clients <sup>1</sup>	nfs4_getfacl	nfs4_editfacl nfs4_setfacl
Solaris NFSv4 clients <sup>1</sup>	ls -lV	chmod A+spec chmod A-spec chmod A=spec etc.
AIX NFSv4 clients <sup>1</sup>	aclget	acledit  aclput
HP-UX NFSv4 clients <sup>2</sup>	N/A	N/A

<sup>1</sup>Display and modification of ACLs over NFS is only supported when the NFS server is running on a StorNext appliance.

<sup>2</sup>While HP-UX NFS clients cannot display or modify ACLs, they are still enforced and inheritance is still in effect.

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 the Apple `chmod(1)` man-page or the StorNext `snac1(1)` man-page for additional details.

## ACL Inheritance

As mentioned in [Display and Modification of File Permissions on page 603](#), ACLs can be assigned to files after they are created. On Windows, it is also possible for an application to explicitly assign a specific ACL to a file at creation time. For example, an ACL can be passed in as part of a security descriptor in the `IpSecurityAttributes` parameter in the **CreateFile** function. However, the most common way that ACLs are assigned to files and folders is through inheritance.


### How Inheritance Works

Each ACE in the ACL on a folder contains a set of flags that determine inheritance behavior. When creating new objects, the inheritance flags on an ACE may dictate that only files, only subfolders, both files and subfolders, or neither types of objects inherit the ACE. Finally, another flag determines whether an ACE should be inherited by children of a folder, but not be used in the permission check for the folder itself.

**Note:** These flags are assigned to each ACE individually. This means that a newly created file or folder may inherit zero, all, or a subset of the ACEs from the parent folder.

On Windows, the inheritance flags can be adjusted implicitly in the Explorer security tab by clicking on **Advanced > Change Permissions > Edit**. On other platforms, these flags can be adjusted explicitly using a CLI. Refer to the [Display and Modification of ACLs on the previous page](#) for the specific command to run on a given platform. [Table 13 below](#) provides the list of ACE inheritance flags and descriptions.

**Table 13:** ACE Inheritance Flags

Flag	Description
<b>File inherit</b>	ACE should inherit to (non-directory) files.
<b>Directory inherit</b>	ACE should inherit to directories.
<b>Limit inherit</b>	Prevents newly created sub-directories inheriting the ACE from further inheriting the ACE to its children.
<b>Only inherit</b>	Causes the ACE to be inherited to files and sub-directories but not used for permission checking on the directory.  <b>Note:</b> The <code>only_inherit</code> flag is never inherited.

On Windows, the `limit_inherit` is mapped to a check box labeled **Apply these permissions to objects and/or containers within this container only**. The mapping of the three remaining tags to the seven Windows propagation pull down menu options are as follows:

Windows	Non-Windows (Mac OS, Linux, Unix)
This folder only	(none)
This folder, subfolders, and files	directory_inherit, files_inherit
This folder and subfolders	directory_inherit
Subfolders and files only	files_inherit
Subfolders and files only	files_inherit, directory_inherit, only_inherit
Subfolders only	directory_inherit, only_inherit
Files only	files_inherit, only_inherit

On non-Windows, 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.

- On Mac OS X, the `chmod` command with the `+a` option can be used to apply an ACE. This can be done for large sub-trees with the `chmod -R` option.
- On Linux and Unix systems, the StorNext `snac1 -R +a` command can be used to apply an ACE to a sub-tree.

## Identity Mapping

The credentials used by Unix-based systems are based on User Identifiers (UIDs) and Group Identifiers (GIDs). These integers that uniquely identify an account holder and the groups they belong to. Username/UID and Groupname/GID mappings exist either in flat files such as `/etc/passwd` and `/etc/group` or a directory service such as NIS or LDAP.

On the other hand, credentials used by Windows systems are based on Security Identifiers (SIDs). The following is an example SID:

```
S-1-5-21-372942913-2183038205-7263820342-8-1077
```

Where the first 1 is the SID revision level, the 5 component is the authority value, the 21-372942913-2183038205-7263820342 component identifies the local computer or domain, and the 1077 is the Relative ID that is unique to a User or Group within the domain.

**Note:** The number of values in this component may vary but typically there are four.

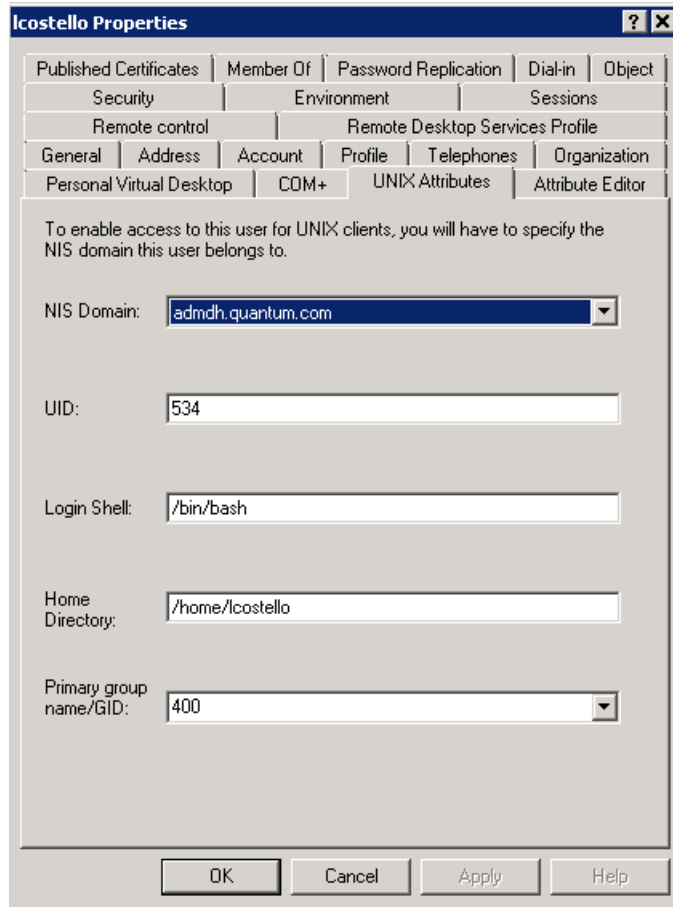
Since UIDs/GIDs and SIDs have different structures, in a heterogeneous environment, a mapping is required so that non-Windows systems can operate on UIDs and GIDs and Windows systems can operate on ACLs; this is referred to as **Identity Mapping**.

A common way to establish this mapping is to use an Active Directory server with RFC2037 extensions. In this setup, the AD database contains both a SID and a UID for a given user. When a Windows system



needs access to the UID for a user, it performs a simple Active Directory query to retrieve it based on the SID. Similarly, if a Unix system needs to retrieve the SID, it can do so by performing an Active Directory query to retrieve it based on the UID. A similar process is used for mapping GID to/from SIDs. RFC2307 ID mapping is supported by all platforms.

**Figure 63:** Updating RFC2307 Mappings in Active Directory



When using the ACL security model, the use of RFC2307 requires that all systems bind to Active Directory. For Linux and Unix systems this requires the use of **Winbind**. This also requires setting the StorNext file system configuration global `UnixIdMapping=winbind`.

Linux StorNext appliances licensed for the NAS feature can be bound to Active Directory by running the following command in the NAS API:

```
x86_64:myhost> auth config ads administrator domaincontroller.mycompany.com  
mycompany.com rfc2307
```

If the SID for the logged in user does not have a UID configured in Active Directory, a newly created file will receive a UID equal to the value of the file system configuration setting `UnixNobodyUidOnWindows`. If the SID for the logged in user does not have a GID configured in Active Directory, a newly created file will receive a GID equal to the value of the file system configuration setting `UnixNobodyGidOnWindows`.

The UNIX modes are controlled by the following file system configuration parameters:

- `UnixDirectoryCreationModeOnWindowsDefault`
- `UnixFileCreationModeOnWindowsDefault`

StorNext allows one set of values for all users of each file system.

---

**i Note:** The value of the file's Windows read-only attribute does not affect the file's Unix mode.

For Active Directory environments that do not require ACLs on Linux and Unix systems, but do require them for Mac and Windows clients, another form of Identity Mapping can be used called **Unix ID Fabrication**. In this configuration, the UID for an account is derived by extracting the first four bytes of the `objectGUID` in the Active Directory account record. To enable this behavior, Mac systems must bind to Active Directory using this form of ID mapping, the StorNext Security Model must be set to "legacy", and the StorNext file system configuration global `UnixIdFabricationOnWindows` must be set to `true`.

For environments using Open Directory that do not have Windows systems but where the use of ACLs on Macs and Linux systems is required, a different form of ID Mapping is required. In this case, StorNext must be configured so the SID for a user can be derived using an algorithmic approach. The **Security Model** for this configuration should be set to `acl` and `unixIdMapping` should be set to `algorithmic`. In addition, a file named `/usr/cvfs/config/domainsid` must be populated on the MDC. This file needs to contain the domain SID being used by the Open Directory domain. This value can be determined by running a command on a Mac. For more information, refer to the `domainsid(5)` man-page.

When ACLs are used in pure Linux/Unix environments, SIDs can be synthesized using a simple algorithmic approach. In this case, the Security Model should be set to `acl` and the `unixIdMapping` should be set to `algorithmic`. The generated SIDs will then use the StorNext built-in domain SID, `S-1-5-21-3274805877-1740924817-4269325941`.

Finally, when using the `unixpermbits` security model, the `windowsIdMapping` variable controls the type of identity mapping performed by Windows clients.

There are three options:

1. **Ldap:** Windows clients map SIDs to UIDs and GIDs using LDAP queries to the Active Directory server. In this case, the Active Directory server must be set up with RFC2307 extensions with the Unix UIDs and GIDs populated.
2. **MDC:** Windows clients map user names to UIDs and GIDs by making remote procedure calls to the StorNext MDC. For this to work, the MDCs must have user account names that exactly match the ones used by Windows. For example, if a Windows user name is "lynn", there must be a Unix account name on the MDC also named "lynn" whose name can only vary by case.
3. **None:** All users map to the user "NOBODY"

## Configuration Examples

The following are examples of high-level steps to configure StorNext for various security configurations.

### Example: Cross-platform ACLs using Active Directory

This configuration uses Active Directory for ID mapping and enforces ACLs on all platforms.

1. From the StorNext GUI, for each file system where ACLs are to be used, set `securityModel=ac1` and `unixIdMapping=winbind`.
2. Join StorNext Windows and Xsan clients to Active Directory. Xsan clients should be configured to use RFC2307 mapping.
3. Join all Linux systems to Active Directory using Winbind using rfc2307. Refer to operating system documentation for instructions. Or, when using StorNext appliances licensed for the NAS stack, the NAS CLI can be used to join to Active Directory. For example:

```
x86_64:mymdc1> auth config ads administrator ad.mycompany.com mycompany.com  
rfc2307
```

- 
- i Note:** Before running these commands, ensure that the `license.dat` files on the MDCs contain proper NAS licenses and that the system clocks on the MDCs and the active directory server are relatively in sync.

## Example: Cross-platform unixpermbits using Active Directory

This configuration uses Active Directory for ID mapping and enforces Unix permission bits on all platforms.

1. From the StorNext GUI, for each file system where Unix permissions bits are to be used, set `securityModel=unixpermbits` and `windowsIdMapping=ldap`.

- 
- i Note:** Complete this step when the file systems are created. Typically, it is not possible to switch to the `unixpermbits` security model after a file system has been created.

2. Join StorNext Windows clients to Active Directory

## Example: Cross-platform unixpermbits using MDC ID Mapping

This configuration uses the password service on the MDC for ID mapping and enforces Unix permission bits on all platforms.

1. Update the password file/database on the MDC so that account names exactly match the local accounts on Windows clients; however, the names are allowed to vary by case.
2. From the StorNext GUI, for each file system where Unix permissions bits are to be used, set `securityModel=unixpermbits` and `windowsIdMapping=mdc`.

- 
- i Note:** Complete this step when the file systems are created. Typically, it is not possible to switch to the `unixpermbits` security model after a file system has been created.

## Example: Cross-platform ACLs using Mac OS Open Directory without Windows Clients

In this configuration, Mac OS clients use Open Directory for ID mapping. Linux clients using algorithmic mapping and it is assumed that no Windows clients are in the environment.

1. Join Mac OS and Linux clients to Open Directory.

**Note:** For Linux clients, use `sssd` instead of `winbind`.

2. From the StorNext GUI, for each file system Open Directory is to be used, set `securityModel=ac1` and `unixIdMapping=algorithmic`.
3. On a Mac OS system run the following command, where `username` is the name of any regular user account in Open Directory.

```
$ dsmemberutil getsid -U username
```

This will return a string such as the following:

```
S-1-5-21-2553502104-2799725507-638401443-3106
```

The Domain SID is the string without the trailing RID so in this example, it has the value `S-1-5-21-2553502104-2799725507-638401443`. The following command may be run on primary and standby MDCs to set this domain SID:

```
mdc# echo S-1-5-21-2553502104-2799725507-638401443 >  
/usr/cvfs/config/domainsid
```

After configuring the `domainsid`, file systems must be restarted on the FSM to have it take effect.

## Example: ACLs on Linux/Unix without Mac OS or Windows Clients

In this example, it is assumed that the environment contains no Xsan or Windows clients but ACL enforcement is still desired.

1. Ensure that all Linux/Unix clients are using the same database for UIDs and GIDs. For example, they are all bound to the same NIS domain, are using identical password files, or are bound to the same Active Directory server, etc.
2. From the StorNext GUI, for each file system where ACLs are to be used, set `securityModel=ac1` and `unixIdMapping=algorithmic`.

For additional information, refer to the `snfs_config(5)` man-page or the following sections in the Windows Help.

- User ID Mapping Overview
- Windows Active Directory Config
- Apple/XSAN Fabricated ID's
- Unix Permissions Background

# 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 keyword `#SNFS_ALL#` is used in place of a file system name 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.

## Controls

Currently seven controls are supported. Each control has the following format:

```
<control value="true|false"/>
```

The `value` can be either `true` or `false`. The `control` is one of the following controls:

Parameter	Description
<code>mountReadOnly</code>	Controls whether the client should mount the given file system as read only. Value “ <code>true</code> ” means the file system is mounted as read only. Value “ <code>false</code> ” means the file system is mounted as read/write. If this control is not specified, the default is read/write.

Parameter	Description
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”. <b>i Note:</b> 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. <b>i Note:</b> 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 ( <code>securityControl</code> ), then the client has no access privilege to the specified file system.

---

## Limitations

Currently only the Linux platform is supported to be a nss coordinator server capable of parsing this xml file. If you have a non-Linux machine as the fsm server, in order to enforce this cluster-wide central control, you must use a Linux machine as your nss coordinator with this central control file in place. The nss coordinator typically can be a very low-end machine as it is not stressed heavily.

---

## Example

Following is an example of a `nss_cctl.xml` file (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 defaults for other file systems using the special keyword `#SNFS_ALL#`.

```
<snfsControl xmlns="http://www.quantum.com/snfs/cctl/v1.0">
  <securityControl fileSystem="snfs1">
    <controlEntry>
      <client type="host">
        <hostName value="192.168.230.132"/>
      </client>
      <controls>
        <mountReadOnly value="true"/>
        <mountDlanClient value="true"/>
        <takeOwnership value="false"/>
        <exec value="true"/>
        <suid value="false"/>
      </controls>
    </controlEntry>
    <controlEntry>
      <client type="netgrp">
        <network value="192.168.1.0"/>
        <maskbits value="24"/>
      </client>
      <controls>
        <takeOwnership value="true"/>
        <mountReadOnly value="true"/>
      </controls>
    </controlEntry>
  </securityControl>
  <securityControl fileSystem="#SNFS_ALL#">
    <controlEntry>
      <client type="host">
        <hostName value="linux_ludev"/>
      </client>
      <controls>
        <snfsAdmin value="true"/>
        <snfsAdminConnect value="true"/>
      </controls>
    </controlEntry>
  </securityControl>
</snfsControl>
```

## Cross-Platform Permissions

When file permissions are properly configured, the ACL and unixpermbits security models can provide fairly consistent behavior, even when the environment contains multiple client platforms. However, some differences are bound to be encountered, especially when ACLs are used. As previously described, this is due to the fact that Windows systems only use ACLs for file permissions whereas other platforms use a combination of ACLs and Unix permission bits. Some of these differences are obvious and others are more subtle. One subtle behavior occurs when an environment initially does not contain Windows clients but later adds them.

When ACLs are enabled, the first time a Windows client mounts the file system, it checks whether there is an ACL on the root of the file system. If not, it applies an inheritable, permissive, “Everyone” ACE on the root directory. From that point on, this ACE will propagate down to every file and subfolder unless the ACE is removed from a subdirectory or its inheritance is disabled. The presence of this ACE allows users to perform actions on files.

If the environment is initially set up without Windows clients and files are created, they typically will not have an ACL. If no ACLs are explicitly applied to files and subfolders and a Windows client is later added, there is no impact to the Windows client because if a file or folder doesn’t contain any ACL, the Windows client will use the ACL on its parent and if its parent doesn’t have an ACL, it will search backward all the way to root and use the first ACL it finds which can be the root “Everyone” ACE.

However, if ACLs are placed on files before any Windows clients ever mount the file system, some unexpected, behavior may occur after Windows clients are added. For example, a Mac user named “jane” may put an ACE to a file that allows read access by “joe”:

```
jane@mac$ chmod +a "joe allow read" myfile
jane@mac$ ls -le myfile
-rw-r--r--+ 1 jane managers 0 Sep 30 11:02 myfile
0: user:joe allow read
```

This results in a file with one ACE since it did not inherit any ACEs from the parent directory. Suppose that a Windows system is then added to the environment. If “jane” attempts to access the file from Windows, she is denied access whether she tries to read the file, write the file, or simply display the security information. The reason for this is that while the file contains an ACL, it does not specify any ACEs that allow her to access the file, and Windows does not fall back to using the Unix permission bits.

The workaround for this scenario is to apply ACLs from the Mac on all files before attaching Windows clients. For example, “jane” could apply ACEs that allow full access to herself, and read access to the group “managers”:

```
jane@mac$ chmod -R +a "jane allow list,add_file,search,delete,add_
subdirectory,delete_
child,readattr,writeattr,readextattr,writeextattr,readsecurity,writesecurity,chow
n,file_inherit,directory_inherit" /Volumes/snfs1/homes/jane
jane@mac$ chmod -R +a "managers allow
list,search,readattr,readextattr,readsecurity,file_inherit,directory_hinerit"
/Volumes/snfs1/homes/jane"
```



---

## Config (.cfg) File Options

The StorNext config file contains the following options that relate directly or indirectly to security or permissions:

Parameter	Description
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.

---

Parameter	Description
Quotas	<p>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.</p> <p>Files created on Windows with <code>WindowsSecurity ON</code> 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.</p> <p>To fix this problem, the UID/GID “space” and SD “space” must be consolidated into one “space”.</p> <p><b>i Note:</b> Quotas can only be enabled or disabled by modifying the <code>Quotas</code> parameter of the file system configuration file. The CLI <code>snquota -L -F file-system</code> command informs you whether the file system has quotas enabled.</p>
UnixDirectoryCreationModeOnWindows	<p>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.</p>
UnixFileCreationModeOnWindows	<p>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.</p>

Parameter	Description
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. 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 online at <http://qsupport.quantum.com/kb/>.

This appendix contains the following topics:

Troubleshooting StorNext File System .....	619
Troubleshooting OS Issues .....	626
Troubleshooting Replication .....	628
Troubleshooting HA .....	630
Troubleshooting StorNext Installation and Upgrade Issues .....	635
Troubleshooting Other Issues .....	636
Debugging StorNext for Object Storage Systems and Cloud Providers .....	638

---

## Troubleshooting StorNext File System

This section contains troubleshooting suggestions for issues which pertain to StorNext File System.

## Question: Why does my StorNext client or MDC not see my newly labeled LUN(s)?

Sometimes the local FSMPM does not automatically see newly labeled LUNs. This results in a system log message like the following:

```
Sep 16 10:32:07 stornext01 fsm[3584]: StorNext FSS 'snfs1[0]': Server could not find any Meta-Data devices!
```

or

```
Sep 16 10:32:07 stornext01 fsm[3584]: StorNext FSS 'snfs1[0]': Node [1] has 1 missing or DOWN disks
```

**Answer:** To resolve this issue, you should force a rescan of the disks by executing the command `cvadmin -e 'disks refresh'` on each system that is unable to see the LUN(s).

## 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)?

- 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)?

- 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 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 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. The following is an example of the `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 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` displays 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 cannot 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 3 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, and then gets stored again by the policy.
- **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 -&gt; RELAUNCH,
next event in 60s
[0327 14:41:59] 0x40305960 (debug) FSS 'stornext-fs1' RELAUNCH -&gt; 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 -&gt; RELAUNCH,
next event in 60s
[0327 14:42:59] 0x40305960 (debug) FSS 'stornext-fs1' RELAUNCH -&gt; 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 following command on the StorNext client:

```
ls -l /usr/cvfs/config/fsmlist
```

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:~ # cvversionsServer 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 fsnameservers (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'.

**i 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 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.000474F892EBB65E000E0000000000000000000000000292BF4". Error opening file
"/stornext/source/__CVFS_
Handle.000474F892EBB65E000E0000000000000000000000000292BF4". 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. For an in-depth look at HA systems and operation, see [High Availability Systems on page 548](#).

## 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 570](#) in the HA topic.

## 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 'fsmList'` command, which displays the local FSMPM's internal FSM and process table. A remote FSMPM's table can also be viewed with `'fsmList on <MDC name or address>'`.

Finally, the message: 'Admin Tap Connection to FSM failed', is an intermittent response that occurred because the timing of the `cvadmin select` command was during the period after the FSM failed the fifth usurpation attempt and before the FSM was restarted (a ten-second delay). At other times, the response will show an activating FSM. Note that the `cvadmin`-displayed asterisk simply indicates that the FSM has been told to activate, not that it has been successful at usurpation and activation.

**Question:** Using the same configuration above (Shasta and Tahoe), an HA Reset occurs if I pull the fibre connection from Shasta when it is the Primary MDC, but it takes 30-40 seconds. Why does it take so long?

**Answer:** When the fibre connection is lost, Shasta's FSMs cannot maintain their brands on their ARB blocks, so the HA timers do not get restarted in the *read, write, restart-timer ARB branding* loop. After five seconds the timers would expire and reset the MDC. However, there is a second method for resetting the timers that uses the LAN.

Every two seconds, the 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 file 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/cvfsct11_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.

**i 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.

## Does a StorNext installation only support a single 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.

## How can I find the Product Serial Number?

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).

---

**i Note:** The serial number is not available through the StorNext application.

## A system panic caused connectivity loss to metadata. Is there anything I can do to prevent this from happening?

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.

## What should I do after a database (MySQL) crash?

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.

## 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

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.

## How do I resolve the "No new media found." issue in systems with archives that have multiple mailboxes available?

Importing media can fail with the message "**No new media found.**" This occurs after selecting **Storage Destinations > Library > Add Media Mailbox** from StorNext's **Setup** menu. To fix this problem, try putting the media in one of the other mailboxes and then re-run the import. If the operation still fails, you can run the import manually by performing these steps:

1. Open up a UNIX root shell on the MDC server.
2. Source the profile by running `./usr/adic/.profile`
3. Obtain a list of available mailboxes for an archive by running `/usr/adic/MSM/bin/mmportinfo <archivename>`
4. Import media into an archive from a specific mailbox by running `/usr/adic/gui/scripts/library.pl add_media --archive=<archivename>`

`--importmethod=mailbox --mailbox=<mailbox>`

Example:

```
/usr/adic/gui/scripts/library.pl add_media --archive=archive01
```

```
--importmethod=mailbox --mailbox=16:LTO:0,0,15,16
```

---

# Debugging StorNext for Object Storage Systems and Cloud Providers

When an error is detected, StorNext generates as much error information as possible to describe the action and the error encountered. Errors may originate internally within StorNext or be generated by one of the StorNext components, including libcurl, SSL, or the HTTP server.

For example, an error in the log file may look like:

```
... rc=412, error: HTTP status 404: 404 Not Found ..
```

where, 412 is the unique error formed by StorNext and 404 is an HTTP error.

Or a specific libcurl error may look like:

```
.. CURLE_OUT_OF_MEMORY ..
```

For more information, consult the following links:

- <https://curl.haxx.se/libcurl/c/libcurl-errors.html>
- <https://www.w3.org/Protocols/rfc2616/rfc2616-sec10.html>

In general, all providers provide the same error status code.

## Debugging Steps to Determine the Source of the Error

1. What operation are you doing when the error occurred?
2. Find the error message in:
  - a. `/usr/adic/TSM/tac/tac_00`
  - b. `/usr/adic/TSM/trace/trace_04` or `trace_05`
  - c. Look at the extended error message. Not all messages have extended verbose error messages from the server, but most do. These extended messages are very useful. An extended error message looks like:

```
.. <Error><Code>SignatureDoesNotMatch</Code> ..
```
3. Which component failed?
  - a. StorNext? Storage Manager or CVFS, etc.
  - b. Is the failure in libcurl?
  - c. Is the failure an HTTP error - Server error?

## Example of an Invalid Address Error

```
Oct 03 10:37:06.486626 cngam-mdc-1 sntsm fsobjcfg[22238]: E1201
(8)<05247>:mdtWASStorage1226: WASStorageRestIF::validateUrlConnectivity: can't
access url https://google.com/cngam-c2s1, rc=412, error: HTTP status 404: 404 Not
Found^M
for url: https://google.com/cngam-c2s1 Oct 03 10:37:06.487035 cngam-mdc-1 sntsm
fsobjcfg[22238]: E1200(7)<05247>:mui5objcfg3606:
mui5WASStorageValidateUrlConnectivity: No connectivity to url
https://google.com/cngam-c2s1
```

## Example of Invalid Credentials (Password)

```
Oct 03 10:29:30.245375 cngam-mdc-1 sntsm fsobjcfg[16600]: E1201
(8)<05194>:mdtWASStorage1226: WASStorageRestIF::validateUrlConnectivity: can't
access url https://s3.amazonaws.com/cngam-c2s1, rc=427, error: HTTP status 403:
403 Forbidden^M
for url: https://s3.amazonaws.com/cngam-c2s1
Additional response info: <?xml version="1.0" encoding="UTF-8"?>
<Error><Code>SignatureDoesNotMatch</Code><Message>The request signature we
calculated does not match the signature you provided. Check your key and signing
method.</Message><AWSAccessKeyId>AKIAJ4WZ7RFOD5Q7NEDQ</AWSAccessKeyId><StringToSi
gn>AWS4-HMAC-SHA256 20161003T152930Z 20161003/us-east-1/s3/aws4_request
5cc6438b6ddf7f91058ec6a9daab21062b49b9ff4f32ba7f675dc50fc41fcf8d</StringToSign><S
ignatureProvided>2d13c2c3a314a207b4907ca71d24e64c8cf92b9cd6fb51de2e113028d5f9c78a
</SignatureProvided><StringToSignBytes>41 57 53 34 2d 48 4d 41 43 2d 53 48 41 32
35 36 0a 32 30 31 36 31 30 30 33 54 31 35 32 39 33 30 5a 0a 32 30 31 36 31 30 30
33 2f 75 73 2d 65 61 73 74 2d 31 2f 73 33 2f 61 77 73 34 5f 72 65 71 75 65 73 74
0a 35 63 63 36 34 33 38 62 36 64 64 66 37 66 39 31 30 Oct 03 10:29:30.245784
cngam-mdc-1 sntsm fsobjcfg[16600]: E1200(7)<05194>:mui5objcfg3606:
mui5WASStorageValidateUrlConnectivity: No connectivity to url
https://s3.amazonaws.com/cngam-c2s1
```







# 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.

**i Note:** The StorNext Offline Notification feature may not function properly on systems using Windows 10. See [Configure StorNext Offline Notification with Anti-virus or Anti-malware Software on page 654](#).

This appendix contains the following topics:

StorNext Offline Notification Overview .....	641
Installing the Notification Application .....	643
Starting the Notification Application .....	650
Configuring the Notification Application .....	650
Uninstalling the Notification Application .....	654

---

## StorNext Offline Notification Overview

The StorNext Offline Notification feature is a Windows only application which helps users prevent 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 topic 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.

---

**i 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.

---

**i 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 64 below](#). If you select this option, you must log in with the credentials for the Administrator account as shown in [Figure 65 on the next page](#).

Figure 64: Run as Administrator

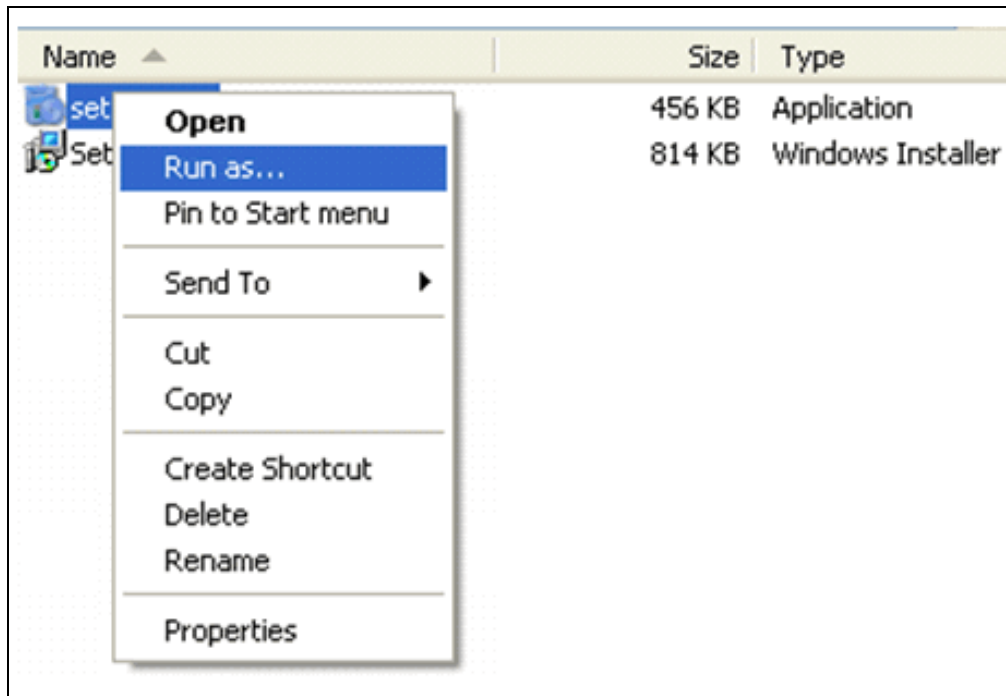
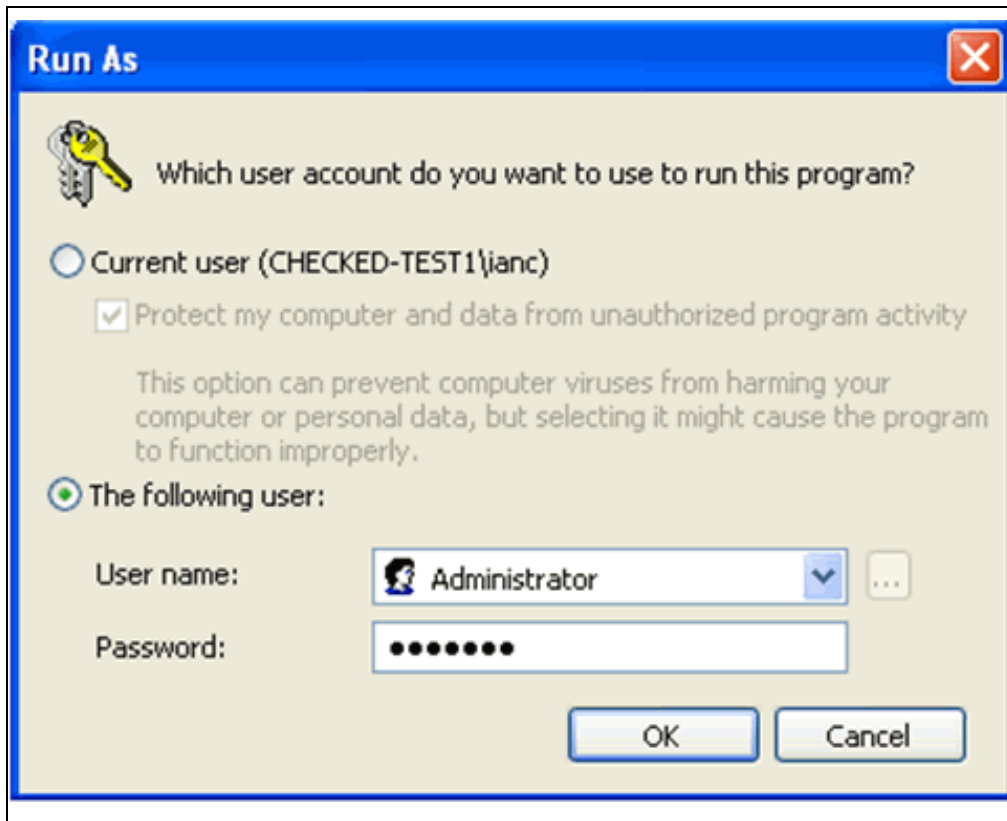
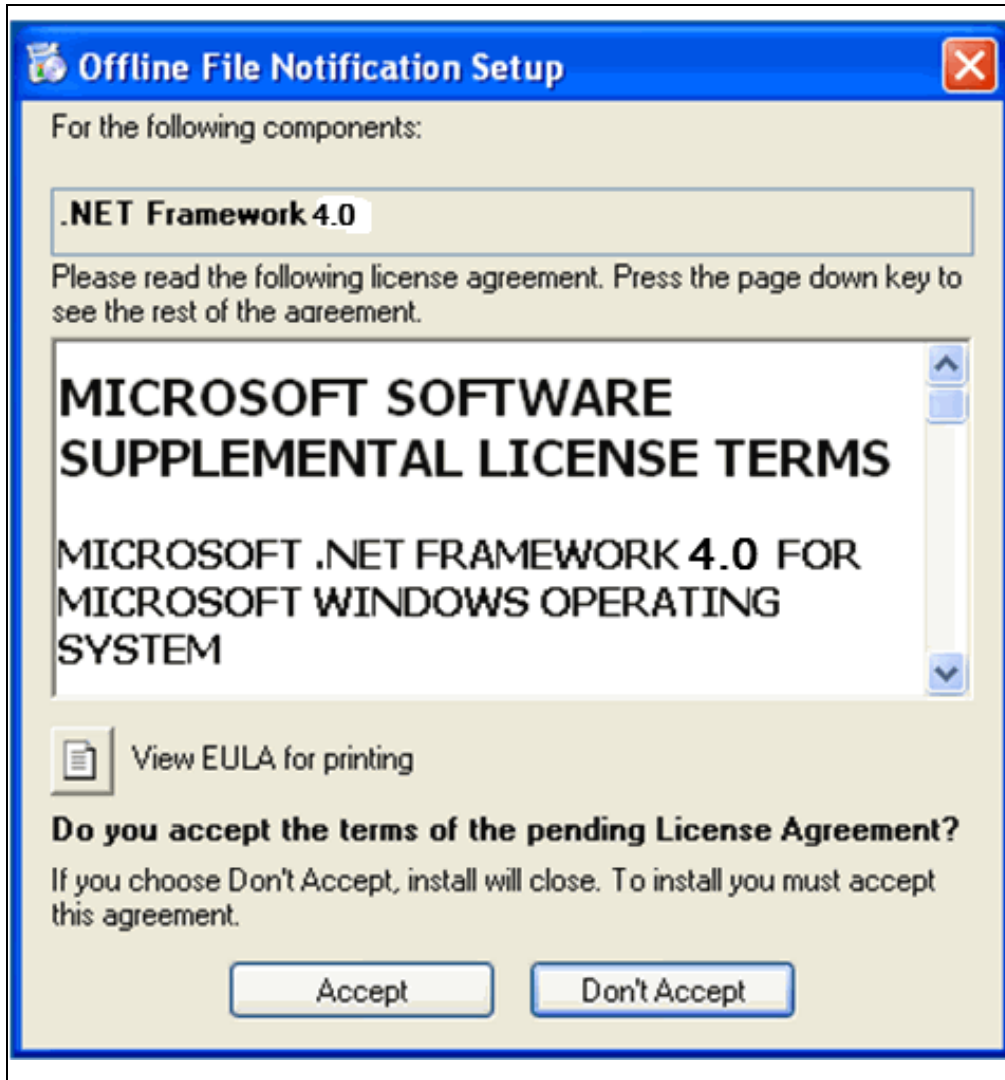


Figure 65: Logging in to the Administrator Account



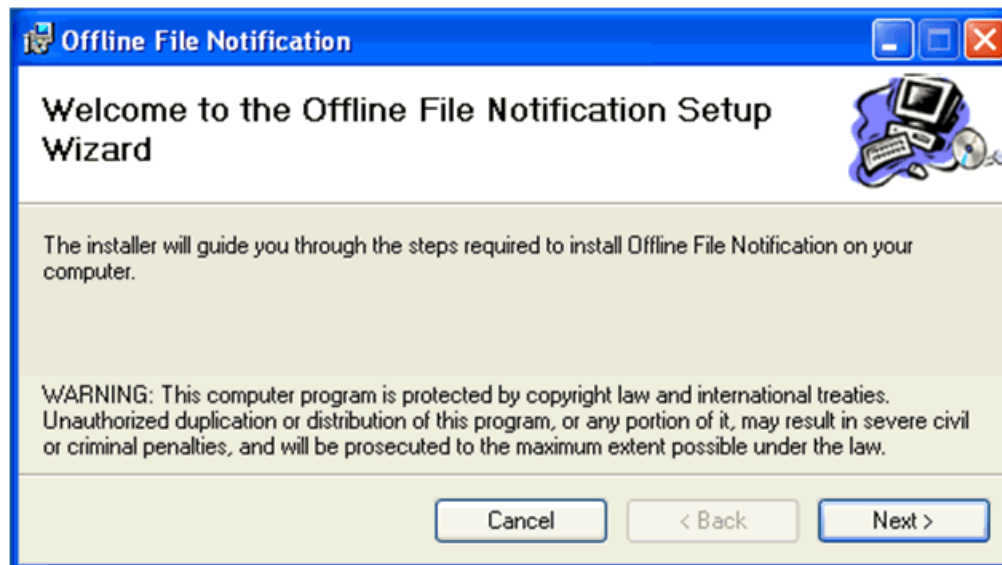
2. The application installer requires .NET updates to function correctly. If you are installing onto a machine without any of the .NET updates, you will be prompted to update to .NET Framework 4.0 as shown in [Figure 66 on the next page](#).

Figure 66: 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 67: Offline Notification Setup Wizard



5. Click **Next** to continue. The Quantum End User License Agreement window appears.

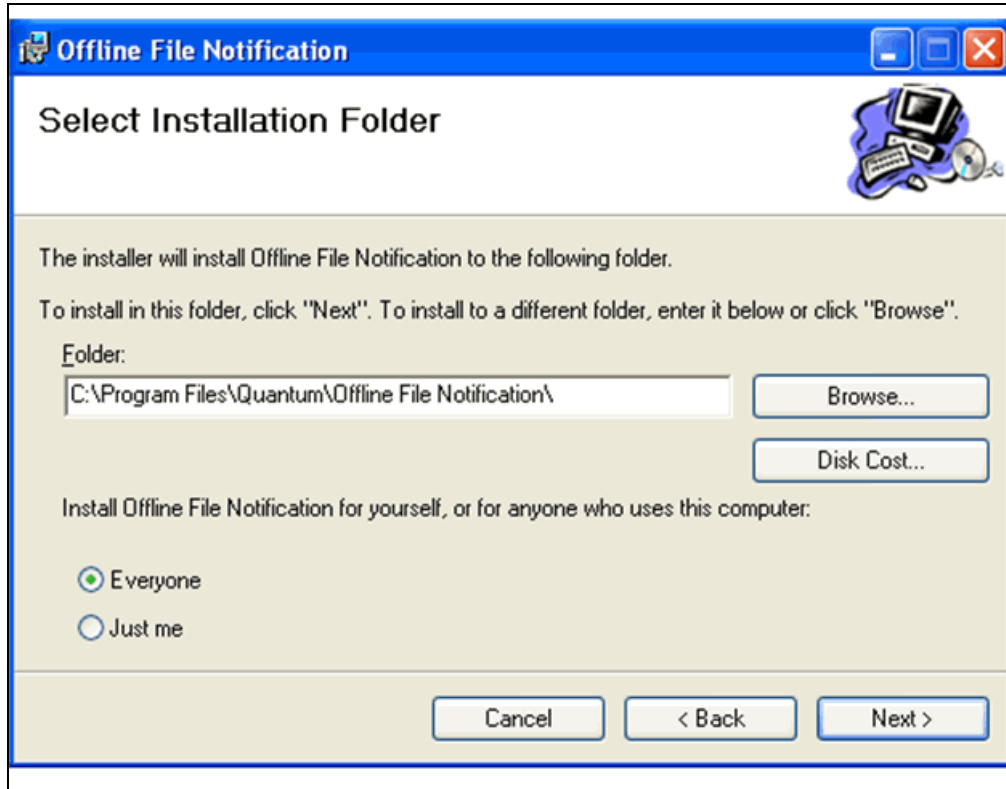
Figure 68: Quantum End User License Agreement





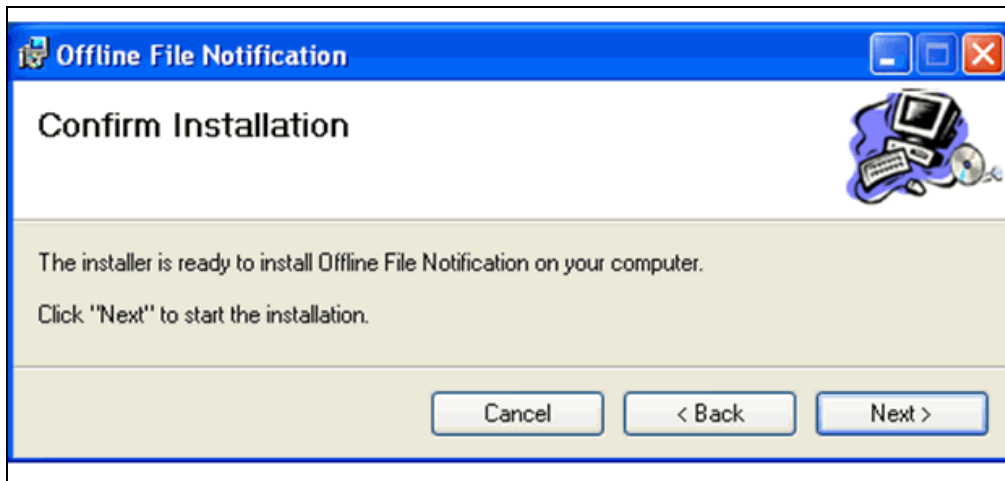
6. Read the end user license agreement. If you accept the terms of the agreement, select **I Agree** and click **Next** to continue. The **Select Installation Folder** window appears. If you do not accept the terms of the agreement, click **Cancel** to stop the installation.

Figure 69: 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 70: Confirm Installation

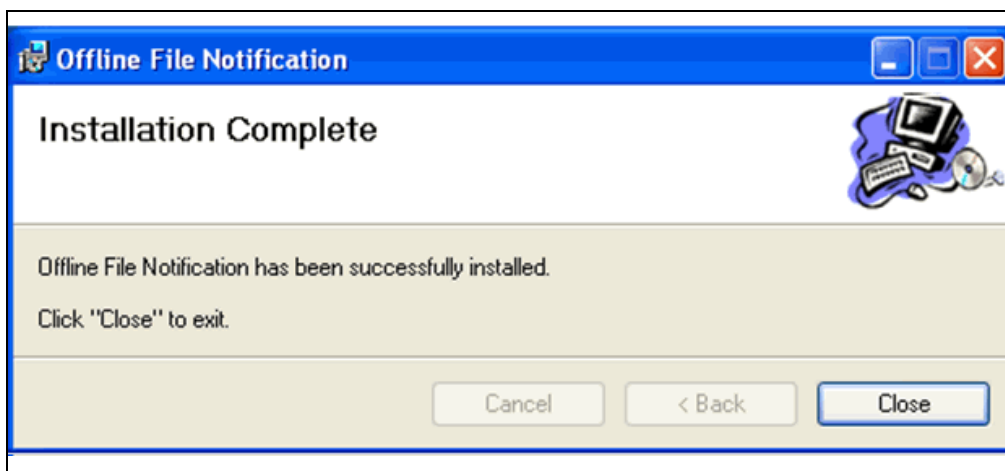


9. Before you proceed, make sure you have enough disk space to install the binaries. A minimum of 10MB is required.
10. Click **Next** to begin the installation.

**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 71: Installation Complete



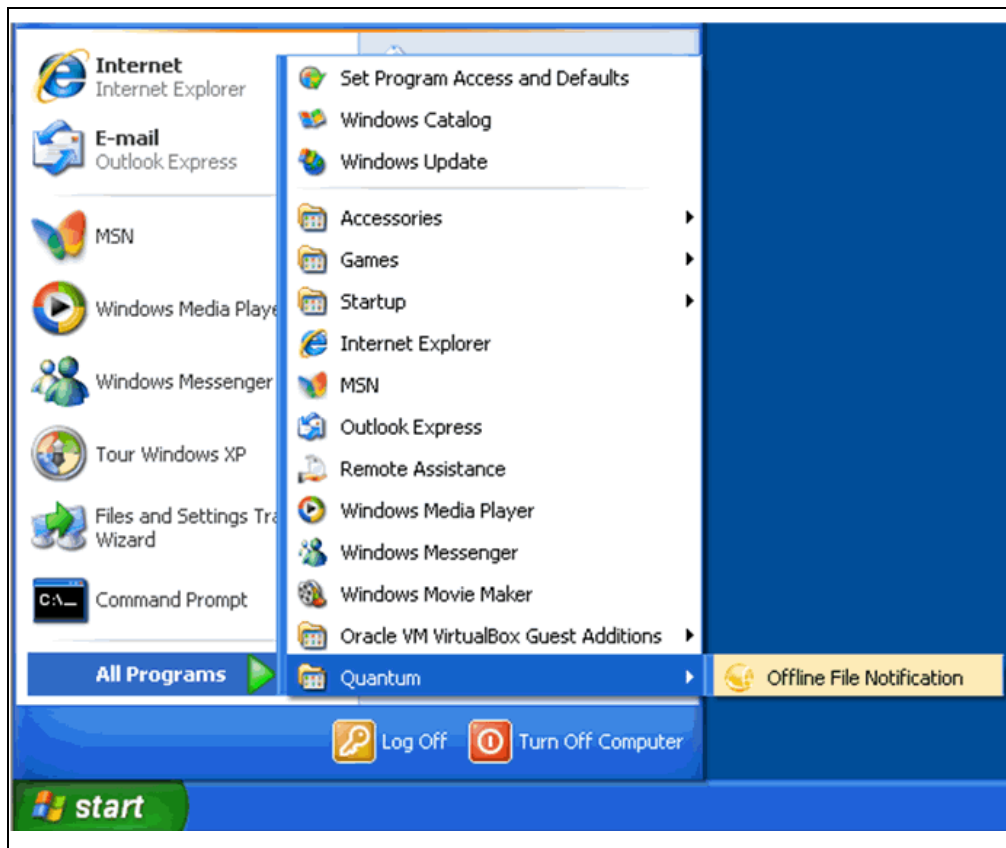
11. Click **Close** to exit the installation .

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, click **All Programs > Quantum > Offline File Notification**. The application starts.

Figure 72: 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:

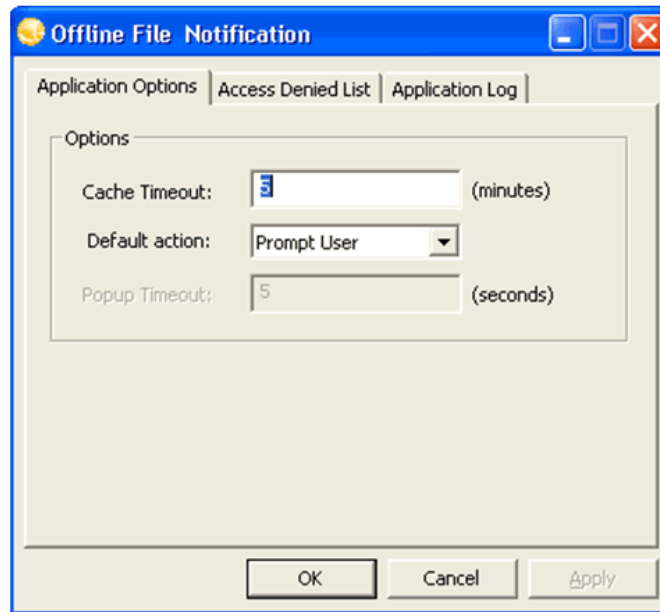
Parameter	Description
<b>Preferences</b>	View and enter application options, view a blocked file list, and view the application log.
<b>About</b>	View the StorNext Offline Notification application's current version number and copyright date.
<b>Exit</b>	Terminate the application.

## 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 73: Application Options



2. View or adjust the following fields on the **Application Options** tab:

Parameter	Description
<b>Cache Timeout</b>	The value in the <b>Cache Timeout</b> 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.

Parameter	Description
<b>Default Action</b>	This drop down list provides three options
<b>Prompt User</b>	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.
<b>Always Deny Access</b>	When this option is selected, access to offline files is always denied, preventing files from being retrieved from offline storage.
<b>Always Allow Access</b>	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 74 on the next page](#)) 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 74:** 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 75:** 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.

## Configure StorNext Offline Notification with Anti-virus or Anti-malware Software

The **StorNext Offline Notification** feature may not function properly when anti-virus or anti-malware protection is enabled on the StorNext file system. The Offline Notification feature cannot block anti-virus or anti-malware protection, such as Microsoft's **Windows Defender**. **Windows Defender** has the undesirable behavior of re-populating truncated files.

---

**i Note:** On systems using Windows 10, **Windows Defender** is enabled by default. To enable StorNext Offline Notification on systems using Windows 10, configure **Windows Defender** to exclude the StorNext directory. See [Configure StorNext Offline Notification on Windows 10 Systems below](#).

### Configure StorNext Offline Notification on Windows 10 Systems

Perform the procedure below to configure **Windows Defender** to exclude the StorNext directory and enable StorNext Offline Notification on a system using Windows 10.

1. On a system using Windows 10, search for **Windows Defender**.
2. Click **Settings**. The **Update & Security** page appears.
3. In the **Exclusions** menu, click **Add an exclusion**, and then click **Exclude a folder**.
4. Select the StorNext mount point/policy directory; that is, the directory or directories that contain the files which will be truncated by the Storage Manager. Once this directory is excluded, the StorNext Offline Notification feature will block access to truncated files in the excluded directory.

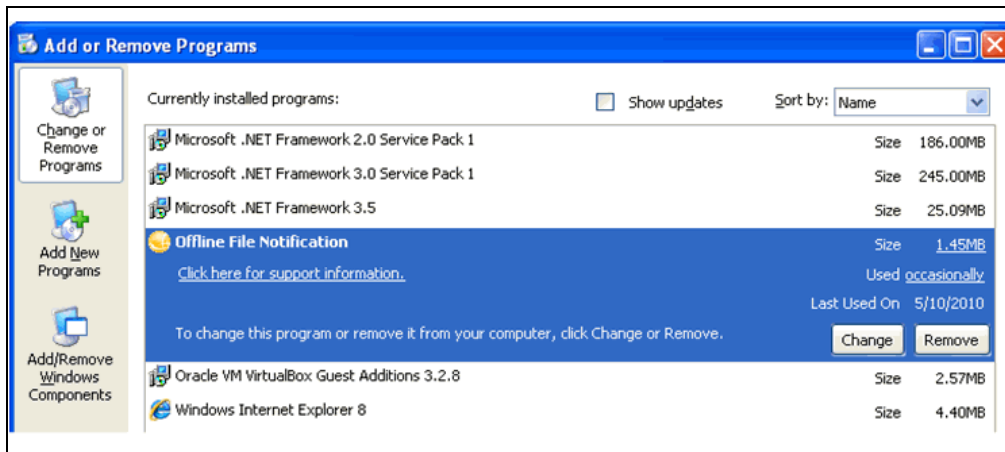
---

## 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 76: 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 5 release 5.3, 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.

**i 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. Contact Quantum Support (see [Preface on page xxi](#)) for an updated `license.dat` file.

This appendix contains the following topics:

Replacing an MDC in a non-HA environment(non-backup/restore method) .....	658
Replacing an MDC in non-HA environment(backup/restore method) .....	660
Replacing an MDC in an HA environment .....	662

# 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.

**i 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
```

1. 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
```

2. Create a tar archive of database metadumps directory:

```
# tar -zcvhf $DESTDIR/meta_source.tar.gz /usr/adic/database/metadumps
```

3. Copy /etc/fstab to the external file system (for later reference):

```
# cp /etc/fstab $DESTDIR
```

4. Remove StorNext with the -remove option to preserve log files:

```
# install.stornext -remove
```

5. Create tar archive of preserved files for later reference:

```
# tar -zcvhf $DESTDIR/preserved_logs.tar.gz /usr/adic
```

6. (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.

7. Install StorNext on new destination using the identical version that the original MDC was running.

```
# install.stornext
```

8. 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.

---

**i 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
```

9. Make sure that StorNext has stopped:

```
# service cvfs stop
```

10. Restore database from tar backup created in Step 5:

```
# tar -zxvf $DESTDIR/db_source.tar.gz -C /
```

11. Restore metadump from tar archive created in Step 6:

```
# tar -zxvf $DESTDIR/meta_source.tar.gz -C /
```

12. 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
```

13. The mount points and fstab entries will need to be recreated. Refer to the **/etc/fstab** file copied to **\$DESTDIR** in Step 7.

14. Start StorNext:

```
# service cvfs start
# service stornext_web start
```

15. Verify that backups are working correctly by running a full backup.

```
# snbackup
```

16. Verify StorNext is functioning by storing and retrieving files.

17. 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, 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.

---

**i 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 mount points and **fstab** entries may need to be recreated. Refer to the **/etc/fstab** file copied to **\$DESTDIR** in Step 4.

12. Restart StorNext with the following commands:

```
# service cvfs stop
# service cvfs start
```

```
# service stornext_web start
```

13. 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
```

14. Verify that backups are working correctly by running a full backup:

```
# snbackup
```

15. Verify StorNext is functioning by storing and retrieving files.

16. 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.

- i 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

- i 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 or to Object Storage 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/fsnameservers` 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 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.