

# StorNext 6.2 Release Notes

## Contents

What's New in StorNext 6.2 .....	2
Supported StorNext Upgrade Paths and Upgrade Considerations .....	26
Compatibility Between StorNext and Other Products .....	27
General Considerations .....	28
Upgrading Appliances .....	29
Appliance Release Notes .....	29
Known Issues .....	29
Contacting Quantum .....	43

---

© 2019 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. Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, DXiV-Series, FlexSync, FlexTier, Lattus, 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.

---

# What's New in StorNext 6.2

## Purpose of this Release

The StorNext 6.2 release introduces several new features.

### Most Comprehensive SAN and NAS Access - Now with Protocol Locking

As more organizations depend on broad collaboration among their staff to meet project deadlines and workflows extend from on-premise into the cloud, being able to effectively manage and access data across complex, hybrid environments has become a requirement for modern file systems. Now with the ability to provide cross protocol locking for shared access across SAN, NFS, and SMB, Xcellis is an optimal solution for organizations that need to share content across both Fibre Channel and Ethernet.

### Integration with S3 Cloud-enabled Workflows

With StorNext 6.2 Quantum now offers an S3 interface to Xcellis appliances, enabling them to serve as targets for applications designed to write to RESTful interfaces. Organizations gain greater flexibility to use Xcellis as either a gateway to the cloud or as an S3 target for web-based applications.

### New Cloud-Based Monitoring Tool Delivers Insights, Maximizes Uptime

With storage environments becoming increasingly complex, users need more powerful tools to proactively manage their storage resources. Xcellis environments can now be managed with a new cloud monitoring tool that enables Quantum's support team to monitor critical customer environmental factors, speed time to resolution, and ultimately increase uptime. When combined with Xcellis Web Services – a suite of services that enables users to set policies, adjust system configuration, and more – overall system management is simplified and streamlined.

Quantum StorNext 6.2 includes new cloud-based monitoring tools.

### Enhanced Multi-Site Data Synchronization

Available with StorNext 6.2, enhanced FlexSync® replication capabilities enable users to create local or remote replicas of multi-tier file system content and metadata. With the ability to protect data for both high-performance systems as well as massive archives, users now have more flexibility to protect a single directory or an entire file system.

## Bolstered Functionality for Media Archives

Organizations that provide storage as a service to departmental users require additional management tools to support their business processes. StorNext 6.2 enables administrators to provide defined and enforceable quotas, implement quality of service levels for specific users, and allows for simplified reporting of used storage capacity. These new features make it easier for administrators to efficiently manage large-scale media archives.

For more information, see [New Features and Enhancements in StorNext 6.2 below](#).

The StorNext 6.2 release also provides important software fixes. See [Fixed Issues and Enhancements Addressed in StorNext 6.2 on page 14](#).

---

 **Caution:** If your system is running StorNext 6.1.x or earlier, see [Product Bulletin 103](#).

## New Features and Enhancements in StorNext 6.2


### Important Information Regarding Storage Manager Performance

As part of the StorNext 6.2 release, an upgraded version of the back-end MySQL database engine has been included to help resolve several security issues. The provider of this engine has changed certain settings to improve the product from corruption in the event of a crash.

Quantum agrees with the changes being included in the StorNext 6.2 release due to the added protection that they provide. As a result, there is a potential that retrieves of small files from a Storage Disk Media/Tier might have a performance impact. For 4K files, Quantum measured approximately a 20% to 25% reduction in retrieve performance from Storage Disk Media/Tier. As the file size increases, the reduction is amortized and has less of an impact on the resultant performance.

Quantum does not recommended you change the settings; however, if this is not an acceptable option and you have to change the settings and are willing to take the risks involved with changing this settings, contact the StorNext Engineering Team and reference Technical Support Bulletin Part Number 6-01194-80.

### Important Information Regarding Upgrades to StorNext 6.2

 **WARNING:** To implement some functionality and performance improvements to the Metadata Archive (**mdarchive**) in the StorNext 6.2 release, the schema for the database used by **mdarchive** has changed. As a result, upgrading to StorNext 6.2 causes FSMs to build new metadata archives, removing all metadata history generated prior to the upgrade. This means that tools **snhistory** and **snaudit** will not report any events that occurred prior to the upgrade and **snrecover** (not to be confused with **fsrecover** for managed file systems) will not recover any files deleted prior to the upgrade.

## Important Information Regarding StorNext Licenses / Licenses for Features

Effective with StorNext 6.2, the following licenses are no longer supported:

- **Deduplication:** A **Deduplication** license was required to use the StorNext Data Deduplication (blockpool) feature. This license is no longer available for either File System only or Storage Manager environments. Deduplication was licensed per terabyte (TB) of the blockpool.
- **Encryption:** An **Encryption** license was required to create encryption keys and use them in storage manager policies for client-side encryption. This license is no longer available.

Also, the StorNext 6.2 release is the final supported version for the following licenses:

- **Solaris LAN Client:** The **Solaris LAN Clients** used IP protocols to read and write data to the StorNext File System through Quantum appliances that can operate as a StorNext LAN Gateway. You must have a LAN Client license for each LAN Client you intend to use with StorNext (in addition to any SAN Client licenses).
- **Replication:** A **Replication** license is required if you want to use the StorNext Data Replication feature. Replication is licensed on a per-MDC (or MDC pair) basis. If replication is used between multiple MDC sets, each MDC set must have a replication license. If deduplication is used in conjunction with replication, a separate deduplication license is also required.

---

**i Note:** There are two types of replication licenses - the File System only Replication license, and the Storage Manager replication license. If Storage Manager is in use on the MDCs, the Storage Manager replication license must be purchased even if the file system being replicated is not managed.

- **UNIX (AIX, HP-UX, or Solaris) SAN Client:** The **UNIX (AIX, HP-UX, or Solaris) SAN Clients** enabled a host computer to mount a StorNext file system with direct block-level access to the disk arrays using Fibre Channel or iSCSI connections. The StorNext file system is licensed on a per-client basis. Any machine that directly mounts the file system is considered a client. If you are an Xsan customer, then you can use an unlimited number of Xsan clients when connected to a StorNext MDC.

## Changes to the Maximum Inactive Versions Option for a Storage Manager Policy

Beginning with StorNext 6.2, the default number of versions of a file that Storage Manager keeps is 2 versions. For prior releases, the default was 10 versions.

If you want the default to be 10, you can perform one of the following tasks:

- Using the StorNext GUI, you can configure the option when you create a policy class. See the **Maximum Inactive Versions** section within the **General Tab** heading in [Add a Storage Manager Policy](#).
- Using the command line interface (CLI), you can run the **fsaddclass** command to specify the default when you create or modify a class.
- Using the CLI, you can run the an editor prior to creating a policy class to set the system parameter to **CLASS\_MAX\_VERSIONS=10**. The system parameter is located within the **fs\_sysparm\_override** file in the **/usr/adic/TSM/config** directory.

---

**i Note:** If you use the CLI to configure the system parameter, you must reboot Storage Manager for the configuration changes to take effect.

- If you are upgrading from a release prior to StoNext 6.2, then the system parameter **CLASS\_MAX\_VERSIONS** is set to 10 in the **fs\_sysparm\_override** file, allowing the default to remain at 10 for all new classes created. If the system parameter **CLASS\_MAX\_VERSIONS** was previously set in the **fs\_sysparm\_override** file, then it will not be reset. If you want the default to be 2, you can perform one of the following tasks:
  - Using the StorNext GUI, you can configure the option when you create a policy class. See the **Maximum Inactive Versions** section within the **General Tab** heading in [Add a Storage Manager Policy](#).
  - Using the command line interface (CLI), you can run the **fsaddclass** command to specify the default when you create or modify a class.
  - Using the CLI, you can run an editor prior to creating a policy class to set the system parameter to **CLASS\_MAX\_VERSIONS=2**. The system parameter is located within the **fs\_sysparm\_override** file in the **/usr/adic/TSM/config** directory.

---

**i Note:** If you use the CLI to configure the system parameter, you must reboot Storage Manager for the configuration changes to take effect.

- i Note:** Whether you are performing a fresh installation or an upgrade, the backup class is automatically set to 10 versions. To change this, you can use the StorNext GUI, or the CLI to edit this value.
- The options to the **fsaddclass** and **fsmodclass** CLI commands for **account number** and **security code** are being removed in a future release of StorNext. The respective options are listed below:

```
-o acctnum
```

```
-l securitycode
```

Consequently, the JSON and XML outputs are no longer being reported in a future release of StorNext. The respective outputs are listed below:

```
"acctNumber": 12345,
```

```
"securityCode": "NONE",
```

The output **NO LONGER** displays the information below as part of the **fsclassinfo** command:

Acct Number: 23456

Security Code: 233

## Enhancements to Storage Manager Processes

In StorNext 6.2, enhancements were made to the two Storage Manager processes below that perform directory scans as part of their normal operations:

- **fsaddrelation**
- **fspolicy** (in rebuild mode)

For both of these processes, the existing scan technology (mapping) was replaced with a new scanning mechanism.

### Changes to the **fspolicy** Command

A new **-C** option was added to the **fspolicy** command that allows you to rebuild all candidate lists as part of the rebuild process.

Additionally, a new **-A** option was added to the command in rebuild mode. When this option is provided, a check for failed **fsaddrelation** processing is done and if any are found to be incomplete, then the processing is completed by the rebuild policy. This option is used by the regularly scheduled rebuild processing to catch any **fsaddrelation** attempts that may have been interrupted and you are not aware of.

For additional information, refer to the **fsaddrelation** and **fspolicy** man pages in the [StorNext 6 Man Pages Reference Guide](#) available on the [PDFs Downloads](#) page in the [StorNext 6 Documentation Center](#).

## Enhancements to the **fsfiletapeloc** Web Service Command

Prior to StorNext 6.2, the **fsfiletapeloc** web service occasionally produced duplicate **fileinfo** keys in its JSON output, causing the JSON data to be invalid. This occurred only when reporting on multi-segmented files. An example is shown below.

---

**i Note:** The **header**, **footer**, and **statuses** objects/arrays have been collapsed for brevity.

```
{
  "header": {},
  "fileInfo": {
    "segmentNumber": 1,
    "mediaId": "ACD683",
    "libraryId": "Library1",
    "format": "ANTF",
    "startBlock": 56580,
    "offset": 128,
```

```

    "segmentSize": 99900000000,
    "blockSize": 524288
  },
  "fileInfo": {
    "segmentNumber": 2,
    "mediaId": "ACD682",
    "libraryId": "Library1",
    "format": "ANTF",
    "startBlock": 190550,
    "offset": 128,
    "segmentSize": 99900000000,
    "blockSize": 524288
  },
  "statuses": [],
  "footer": {}
}

```

Beginning with StorNext 6.2, the **fileinfo** has been replaced by a **segments** array within a **files** array. An example is shown below.

**i Note:** The **header**, **footer**, and **statuses** objects/arrays have been collapsed for brevity.

```

{
  "header": {},
  "files": [
    {
      "fileName": "/stornext/snfs1/poll/file1.txt",
      "segments": [
        {
          "index": 1,
          "mediaId": "ACD683",
          "libraryId": "Library1",
          "format": "ANTF",
          "startBlock": 56580,
          "offset": 128,
          "segmentSize": 99900000000,
          "blockSize": 524288
        },
        {
          "index": 2,
          "mediaId": "ACD682",
          "libraryId": "Library1",
          "format": "ANTF",
          "startBlock": 190550,
          "offset": 128,

```

```

        "segmentSize": 99900000000,
        "blockSize": 524288
    }
  ]
},
"statuses": [],
"footer": {}
}

```

The **fsfiletapeloc** web service can also produce XML output. The XML output has also been corrected in a manner similar to the JSON fix. Below is an example.

**i Note:** The **header**, **footer**, and **statuses** elements have been collapsed for brevity.

```

<?xml version="1.0" encoding="UTF-8"?>
<fsfiletapeloc xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="fsfiletapeloc.xsd">
  <header></header>
  <files>
    <file>
      <fileName>/stornext/snfs1/pol2/file1.txt</fileName>
      <segments>
        <segment>
          <index>1</index>
          <mediaId>ACD683</mediaId>
          <libraryId>Library1</libraryId>
          <format>ANTF</format>
          <startBlock>56580</startBlock>
          <offset>128</offset>
          <segmentSize>99900000000</segmentSize>
          <blockSize>524288</blockSize>
        </segment>
        <segment>
          <index>2</index>
          <mediaId>ACD682</mediaId>
          <libraryId>Library1</libraryId>
          <format>ANTF</format>
          <startBlock>190550</startBlock>
          <offset>128</offset>
          <segmentSize>99900000000</segmentSize>
          <blockSize>524288</blockSize>
        </segment>
      </segments>
    </file>
  </files>
</fsfiletapeloc>

```



```

</files>
<statuses></statuses>
<footer></footer>
</fsfiletapeloc>

```

For additional information, refer to the [File / Tape Location](#) section of the [File Commands](#) topic in the [StorNext 6 Documentation Center](#).

## Multiple Options Added to the `fsmedcopy` Command

### -R Option

Beginning with StorNext 6.2, you can recreate the contents of an unusable medium on one or more other media.

---

**i Note:** You can only use feature if there are multiple copies stored for each file on the bad medium because the data necessary to recreate the bad tape is pulled from the other copies.

You can use the CLI command, `fsmedcopy`, to provide the desired capability. A new `-R`, option was added which indicates that the files on a medium should be recreated on other media from other copies. An argument to the `-R` option allows you to specify which copy to use to recreate the bad tape. You can also specify `0` and allow Storage Manager to select the copy(s) to use to recreate the bad tape.

There are limitations with the new option:

- The Media Status of the bad medium should be **Unavailable**.
- If any alternate source media are in an unusable state or vaulted, then they are not used.
- The size of the file segments on the alternate source media to be used and the media to be recreated must be the same, or those segments are not copied. Typically, this only occurs when the media types are different for the alternate source media and the media being recreated.

---

**i Note:** A notification is displayed if there are any files that can not be copied due to mismatched segment sizes.

### -L Option

A new `-L` option was added so that you can throttle and limit the number of concurrent streams used for the request, when doing migrations. The default value is ten (**10**).

---

**i Note:** The `-L` option is especially useful with object storage configurations.

### -I Option

A new `-I` option was added to the `fsmedcopy` command and the `fsfilecopy` command, that you can use to ignore the cleanup of the Object Storage when the source is an object media. If the source is not an object media, then this option is ignored. If you provide the option, then the cleanup of the Storage Manager database occurs, but there is no attempt to delete the objects from the Object Storage. This option is useful when migrating data off of an Object Storage appliance that is no longer being used.

For additional information, refer to the **fsmedcopy(1)** command and the **fsfilecopy(1)** command in the [StorNext 6 Man Pages Reference Guide](#), and also the **Media /Copy Media** API command in the [StorNext 6 Web Services Guide](#).

## Enhancements to the fsretrieve Command

Beginning with StorNext 6.2, a new **-A** option was added which allocates disk blocks to files in alphabetic order in an attempt to provide sequential block allocation to all files in the specified directory.

---

**i Note:** The **-A** option is only valid with the **-D** option.

Also, the **fsretrieve** command now supports the **-t affinity** option to retrieve files to a specific disk affinity on the file system. The affinity must be defined for the file system in which the files reside. An affinity denotes a specific disk or set of disks (a stripe group) on which the file system resides.

For additional information, refer to the **fsretrieve(1)** command in the [StorNext 6 Man Pages Reference Guide](#) and the **File / Retrieve Files** API command in the [StorNext 6 Web Services Guide](#).

## Changes to Commands that Provide Recursive Processing on Directories

Beginning with StorNext 6.2, the following commands are enhanced to include a new CLI option, **-D**, that only processes the directory and its contents.

- **fschdiat**
- **fsexpcopy**
- **fsfilecopy**
- **fsfileinfo**
- **fschfiat**
- **fsretrieve**
- **fsrcmcopy**
- **fsrcmdiskcopy**
- **fsstore**

In addition, the CLI option, **-D** for the two commands below is changed to **-C** in order for the **-D** option to work the same way for all commands processing a directory.

- **fschdiat**
- **fschfiat**

Changes were made to the commands below that affect the handling of symbolic links (symlinks) when they are encountered by the commands.

- **fschdiat**
- **fschfiat**
- **fsexpcopy**

- **fsfilecopy**
- **fsfileinfo**
- **fsrcopy**
- **fsrcmdiskcopy**
- **fsstore**

With the updates, there is now more flexibility provided as to how symlinks are handled. Prior to this update, these commands mirrored the behavior of commands like **ls** or **cd**, symlinks were followed. Now, with command line options, behavior like the default of the **find** command can be requested where symlinks are not followed. While the default behavior of the **find** command is not to follow symlinks, the **-H** and **-L** options are provided by that command to allow for specifying symlink behavior.

The new options that were chosen for the SM commands are those same options provided by **find**. Users of symlinks are likely to be familiar with the operations of these options. As with the **find** command the new options to the SM commands behave as follows:

- **-H**, when provided do not follow symlinks that are encountered.
- **-L**, when provided follow symlinks that are encountered.

**i Note:** One last change was made with respect to these new options. The **fsfilecopy** command already had a **-L** option which was used to specify a limit to the number of concurrent I/O operations. The **-L** option now specifies symlink behavior as with the other commands. The **-O** option is now used with **fsfilecopy** for specifying the operations limit.

For additional information, refer to the [StorNext 6 Man Pages Reference Guide](#) and the [StorNext 6 Web Services Guide](#).

## Enhancements to the Export and Import Functionality

The ability to export and import LTFS media was introduced in the StorNext 6.0 release. Beginning with StorNext 6.2, the export and import functionality is enhanced to support ANTF media. The ANTF export/import feature closely mirrors the support already in place for LTFS media, but there are two major differences.

- StorNext LTFS media does not support segmented files, while ANTF media does support segmented files. The ANTF export/import feature fully supports ANTF media that contain segmented files.
- The LTFS export/import feature works in either a **Media Ingest** mode or a **File Ingest** mode when importing data from media. The ANTF export/import feature supports the **Media Ingest** operation for ANTF media, but not the **File Ingest** operation. However, there is a manual process documented at [Media Manifests](#) that you can follow to achieve a **File Ingest** operation for ANTF media.

The overall export and import functionality is also enhanced in the following ways.

- Tape manifests for both ANTF and LTFS media are now part of the export and import processing. A manifest file is generated for each tape that is exported. When media content is imported on a destination system, these manifest files provide more robust and efficient import processing.
- Import processing now supports importing multiple media with one command.

## Enhancements to the `snquota` Command

In releases prior to StorNext 6.2, the grace period argument to `snquota (-t)` would accept modifiers for **hours**, **days**, **weeks**, **months** and **years** (**h**, **d**, **w**, **m/M**, **y**).

---

**i Note:** Both lowercase and uppercase **M** specified months.

In the man page, lowercase **m** is supposed to be **minutes** and uppercase **M** is **months**. Beginning with StorNext 6.2.0, the `snquota` command has been changed to do just that. However, if you were relying on lowercase **m** to mean **months** you must change to using uppercase **M**.

Below is an example of the new behavior to set a grace period of 60 minutes:

```
snquota -F data -S -u lisa -h 10g -s 9g -t 60m
```

Below is an example of the new behavior to set a grace period of 2 months:

```
snquota -F data -S -u lisa -h 10g -s 9g -t 2M
```

For additional information, refer to the `snquota(1)` command in the [StorNext 6 Man Pages Reference Guide](#), and also the **Quota / Manage Quotas** API command in the [StorNext 6 Web Services Guide](#).

## Important Information About the RTIO and RVIO Functions

StorNext 6 introduced a new quality of service feature called Quality of Service Bandwidth Management (QBM). All quality of service enhancements are being made to QBM only. No further enhancements are planned for the RTIO and RVIO functions, which are being deprecated in a future release of StorNext software. For additional information, see [Quality of Service Bandwidth Management \(QBM\)](#).

## Enhancements to LTFS Compression Functionality

Prior to StorNext 6.2, all data written to LTFS media was uncompressed and there was no way to modify this. The drive configuration setting did not affect this.

LTFS media require that, at format time, the setting for compression be set, otherwise the media is uncompressed.

Beginning with StorNext 6.2, any new media that is not formatted has compression enabled by default, if the LTFS format type is used. Media that already have data written to them remain unchanged and future data written to those media remain uncompressed.

---

**i Note:** If you want to disable compression on your media, you can set the Storage Manager system parameter `LTFS_COMPRESSION=false`; in the `/usr/adic/TSM/config/fs_sysparm_override` file. After setting the system parameter, you must restart Storage Manager.

## Features No Longer Supported Effective with StorNext 6.2

- Data deduplication option
- End-of-Life (EOL) DAS libraries
- End-of-Life (EOL) AML libraries
- File systems using small inodes
- MDC support for SUSE Linux
- Obsolete 3rd party tape drives
- Partial File Retrieval (PFR)
- Solaris LAN client
- StorNext 4.x Web Services (WSAPI)
- StorNext API (SNAPI)

If you are running EDLM or ActiveVault in the Quantum i6000 or i500 tape library, your tape library must be at the following minimum firmware version:

- Scalar i6000, i13 or later
- Scalar i500, 710G or later

Also, with EDLM and ActiveVault, you must enable and configure Web Services on the library in order to communicate with StorNext (see the Tape Library documentation for details on this configuration). You must perform this configuration, since SNAPI is no longer supported. In prior versions of the Tape library, communication with StorNext relied on SNAPI.

- StorNext Web Services V1
- UNIX SAN clients

## Compatibility and Support

The [StorNext 6 Compatibility Guide](#) provides the basic compatibility for StorNext 6.2, includes the StorNext components supported, operating systems and service packs, libraries and drives, browsers, virtual machines, and appliance support. Listed below are just a few of the types of information available to you in the [StorNext 6 Compatibility Guide](#).

- **Upgrade Paths:** Provides information on what upgrades to this release are supported.
- **Appliance Support:** Provides information on what StorNext and Lattus appliances are supported with this release or are compatible with it.
- **Operating Systems and Platforms:** Provides information on what StorNext components run on various operating systems and service packs. Also includes which operating systems have been newly added or removed.
- **Client Interoperability:** Provides information on what StorNext clients running other versions of StorNext are compatible with metadata-controllers (MDCs) running this release.
- **Virtual Machine Support:** Provides information on what StorNext components running on selected operating systems and service packs are supported in virtual machines.

- **Compatibility with Other Products:** Provides information on references to additional StorNext sold-separately products that are supported with this release.
- **Browser Support:** Provides information on what versions of browsers are supported with the GUI in this release.
- **Drives and Libraries:** Provides information on what Quantum and 3rd party drives and libraries are supported with this release.

## Fixed Issues and Enhancements Addressed in StorNext 6.2

Operating System	Change Request Number	Service Request Number	Description
All	13880	1104418, 1320360, 1310204, 1598810, 3359478, 1603352, 3352398, 3369488, 3423152, 3442602, 3745562, 402821, 399556, 375593, 414277, 425032	Need way to copy good files from a tape that's bad
All	31660	1159754	Enhancement for snfsdefrag messaging and exit status
All	31707	1154098, 1289298, 1330012, 1358812, 1360904, 1500604, 3447454, 3471990, 365899, 372013	Enhancement: email the "failed to retrieve" admin alert
All	35857	1322120, 1390768, 3655030, 296253	fsaddrelation against a populated directory can leave things incomplete
All	43848	3394824, 3552900, 3644524	snfsdefrag should be able to move all files off of a stripe group, including sparse (holey) and future mtime
All	47077	322038, 342913	Metadata Error : Quota update would have caused current size to be negative
All	55069	3499882, 311199	Enhancement to allow file systems with Named Streams enabled to be managed
All	55914	n/a	rebuild policy does not correct/add missing attributes

Operating System	Change Request Number	Service Request Number	Description
All	59146	3658938, 3546014, 411825	fspostrestore failed to complete on 1.6B file system after 2 weeks runtime
All	62969	3646340, 437668	fs_mtranx & fs_fmover may consume huge amounts of RAM
All	63858	3695950	Xcellis pair does not have same uid for quantumdb
All	64564	3524236, 3685640, 3644044, 3658054, 3675606, 3682962, 3696610, 347921	Identify and fix the cause of FCNT becoming incorrect on tapes
All	65324	203093, 451962	fsm ASSERT failed pclient->cl_rsvd_counted == 0
All	65613	292237	Time for cvgather in pse_snapshot can be too short for it to complete
All	65712	411187	Improve qrdb space utilization and debugging
All	66241	304804	GUI Web Services debug logging enables do not work, i.e. does not log anything
All	66335	343125	fsclean hangs if objects being deleted fails
All	67571	468099	fs_moverd: not mounted file systems is detected but does no corrective actions.
All	67851	401898	recreate can bring a file back into existence that was deleted intentionally. PHASE 2
All	68544	342232, 354576	WebServices are not logging the host/user that send a request via webapi
All	68655	424274	Request for linux 4.8 kernel support ... needed someday and used already by Ubuntu HWE

Operating System	Change Request Number	Service Request Number	Description
All	68768	346928	StorNext should provide a way to garbage collect unused security descriptors
All	69147	348367	extra space in some TSM log entry "invalidated due towrite event"
All	69348	355297, 450619	GUI does not support umlauts for File and Directory Actions eg. fsrecover
All	69477	360713	When a dpserver entry is present, whilst a MDC is licensed for 'gateway' /var/log/message gives proxy license error when FSM is restarted
All	69496	360349	Add retry attempts to locate correct device path into fs_drv
All	69607	0340017	cvgather multiple paths devices are not displayed uniquely
All	69671	364890	debug statement to trace s_biggest prints values inconsistent
All	69839	370627	unable to add media to library in SN GUI, when media extended barcode support (XDI_USEBARCODEMEDIATYPEID) is enabled within MSM
All	69970	1322120, 1390768, 3655030, 296253, 3527552, 3505060	Update fsaddrrelation/rebuild policy to improve processing
All	70006	n/a	When using the Unixpermbits security model, quotas should be allowed
All	70007	n/a	Unixpermbits should support more than 32 supplemental groups
All	70023	425973, 442545, 445787	Licenseflex_private_cloudhas reached max capacity of 0



Operating System	Change Request Number	Service Request Number	Description
All	70065	374873	import of LTOW WORM Media fails with "Attempted enter of invalid Media" when extended barcode reporting is disabled on library
All	70191	377480	sncompare should check FILECOMP tables for active versions entries of removed files and fix endtime for them so that fsclean will clean them up
All	70202	377480	filecomp* endtime was not set [hence stored copies not invalidated] for file removed
All	70259	382872	OS issue causing a tape file corruption (duplicate I/O block written) when /usr/adic/TSM/util/tsm-stacks is run should not allow foolish run
All	70359	373989	sncompare should escape single quotes and other special chars in filenames and paths, in order to ensure repair*.sql won't fail
All	70407	n/a	The conversion/upgrade processing in fs_policyd is not checking the exclusions file
All	70462	389650	Enhancement: Health Check creates RAS ticket for SDISKS when policy is not configured
All	70464	407536, 395041, 423901, 454583, 472848	MSM failure to start (archive not getting ready) caused SMITHs
All	70475	396650,396981, 407598	GUI should complain for / partition filled up
All	70602	396890	snquota option -twith qualifier 'm' for minutes instead sets months (M) when setting the GracePeriod

Operating System	Change Request Number	Service Request Number	Description
All	70625	398258	Enhancement request - add functionality to NSS_CCTL to allow/disallow globalSuperUser per client machine
All	70632	394615	fs_altstore should check that filesystem is mounted before attempting to stat files associated to ASL requests
All	70640	398079	MDC smith during upgrade because TSM did not start properly.
All	70879	403645	hm_connect() TestLink fail message does not provide ip and port information
All	70888	403392	PSE snapshot should include "chkconfig --list" command output
All	70956	403555	ASL ARL related: admin alert issued for failed retrieve, although alternate retrieval was successful
All	70999	404184	HAMgr: hm_common_processing for cmd 1 needs to be converted to a human-readable string in the message
All	71050	402207	checkDiskSpaceTsm can't handle a customised location for mdarchives i.e. non-default metadataArchiveDir
All	71058	388584	New API calls hang if there are more than 150 calls being processed by the MDC
All	71114	n/a	PFR fails when MDC is disabled as MoverHost for DDM - No usable DDM hosts (E_DDM_NO_HOST_MS)
All	71127	408927	Batch retrieve doesn't skip over unavailable media like recursive retrieve

Operating System	Change Request Number	Service Request Number	Description
All	71147	451310	Clients using NSS2 cannot see file systems if they are not on same subnet as coordinators
All	71205	410232	cvgather_multipath provides insufficient information when verbosity in /etc/multipath.conf is set to 1 or less
All	71210	415249	RDAR: 31152810: Xsan client panic in _BumpPathReference()
All	71233	416280	Admin alert to convert MySQL should not be issued on MDCs with no Storage Manager licensed
All	71257	429099	NSS2 - StorNext clients are unable to mount file systems using non-preferred network
All	71260	417420	Enhancement request - permit globalSuperUser=no on managed file systems
All	71295	444615	Adjust to renaming inode_change_ok() to setattr_prepare() and changing signature in kernel 4.13+
All	71299	456036, 459515, 460901	nexpose run results in fs_copymand/fs_moverd core
All	71304	411188	specifying OpHangLimitSecs to 0 did not disable it in mdarchive.c
All	71345	n/a	mdarchive rebuild prints too many messages overflowing /var/log/messages
All	71348	419415	cvfsck failed to fix icb when not using -Y option.
All	71392	417808	snapshots should include the status of the mdarchive

Operating System	Change Request Number	Service Request Number	Description
All	71448	403677,427190,425216, 452189, 460912	snbackup fails because fsstore command failed to store a multi-segment file
All	71498	423130	snac1 +a# should not allow allow ACL before deny
All	71572	420913	Better handle RID overflow in algorithmic ACL
All	71577	459572	Update percona-server to 5.7.23-23.1
All	71602	426019	cvfsck is not reporting correct file size stats
All	71605	n/a	Utility for determine how many retrieves were done with recently truncated files
All	71612	422667,426264	Stornext GUI incorrectly reports library slot count as vault slot + library slots
All	71616	426400	Linux client incorrectly applied inherited ACL when security model was legacy
All	71619	426372	Minor metadata restore performance improvements
All	71683	414900	fs_resourced not able to allocate drive it is looping until policy timeout
All	71704	430418	CvApi_SetNtAttributes() does not set the immutable flag when setting the Readonly attribute
All	71733	429099	fsmpm can not find cluster causingfilesystem not mounted on gateway
All	71735	430228	When trying to move an unmanaged directory to a managed directory the command hangs

Operating System	Change Request Number	Service Request Number	Description
All	71777	432068	How many sgmanage/sgoffload commands can be run in parallel?
All	71789	425088	fsimport killed due to OOM during import of LTFS media with "--ingesttype files" option
All	71814	n/a	SN 6.2 - Add support in fsretrieve Web SVC call for new -g glacier-types
All	71839	355440	Unable to move the RAS "queue" directory off the root file system on the fsnameserver box
All	71884	432914	fsclean -a -R -D does not remove files if directory path contains spaces
All	71901	423376, 413967, 422579, 407282, 451471, 457294, 460363	snbackup failure due to files with multiple current versions
All	71944	424025	Security Scanner cause FSM on secondary MDC to die Weekly
All	71976	422355, 442979, 449687	wsar_agent : segmentation fault causes webservice to be down
All	71979	435103	GUI Information for SSL Certificate should be more meaningful and less confusing.
All	71996	438298	nss_cctl file being synced but not re-read by fsmpr on secondary node causing smith after failover due to blocked access
All	72025	435562	sn_log_update check for primary MDC fails after changes to fsstate -f
All	72031	436505	/usr/adic/TSM/util/snnas_usage returns without error when no relation point exists

Operating System	Change Request Number	Service Request Number	Description
All	72038	1104418, 1320360, 1310204, 1598810, 3359478, 1603352, 3352398, 3369488, 3423152, 3442602, 3745562, 402821, 399556, 375593, 414277, 425032	Need way to copy good files from a tape that's bad from gui
All	72047	426964	request for Windows Offline File Notification to give option to right-click and truncate or retrieve like the MAC version
All	72061	440161	frequent soft lockup warnings, threads in PurgeCvNode()
All	72157	441240	SQL code injection possible with ws-api call.
All	72219	462208	cvfs build fails on Debian 8.11
All	72221	443945, 422536, 457091	Data partition on LTFS tape is full but TSM keeps trying to write to it ignoring errors
All	72230	446081	cvmkfs -r incorrectly issues an error when metadata is moved into its own stripe group
All	72242	440828	file corruption when running vidiomap -r on files larger than 64MiB
All	72243	443169	snhistory triggers a segmentation fault while generating the full pathname for a new hard link
All	72250	445916	Multiple "fsrmdiskcopy -R" commands running on same directory at the same time can leave files in an incorrect state
All	72274	432124	snprobe reports a host as down when it actually is up
All	72288	444657	REST URL advertised by FSM is not always usable

Operating System	Change Request Number	Service Request Number	Description
All	72289	444657	SNFS should not allow bad timestamps to be set or exposed
All	72523	449338	Kernel panic kernel BUG at fs/dcache.c:1383!
All	72526	448021	mdarchive backup failed with cannot run: bash -c umask 177cat
All	72565	449257 3593004	TSM creates LTFS ill-formed dummy path names in case of problems getting file path
All	72569	449055, 406908, 424025, 469386	fsmpm stuck in new_input() handling RAS event can cause filesystem failover
All	72590	449738	Special characters in web services have incorrect encoding
All	72658	452368	StorNext Linux clients incorrectly disallow setting xattrs on directories containing the sticky bit for the owner and superuser
All	72704	451847	new coherency model w/ global share mode enabled can hit fsm "ASSERT failed "(optr->open_flags & OPEN_HAVE_BOTH) == OPEN_HAVE_BOTH" upon InodeOpen()
All	72726	3734274, 336772, 409963, 430039	fsmedscan can cause the root filesystem to fill if run on StorNext appliances
All	72791	n/a	Create a new option for sgmanage to list the files that have an extent on a stripe group
All	72839	455314	cvgather will leave files in /tmp when run. Should clean up these files when cvgather completes to avoid filling up root filesystem

Operating System	Change Request Number	Service Request Number	Description
All	72890	406908	The mdarchive query used by sgdefrag/sgoffload is not efficient, causing performance issues
All	72964	406908	Improve various situations which can cause mdarchive query delays
All	72976	455446	Restarting only MSM removes the "ipc:msm_arcdisp" socket
All	72977	72977	GUI Better handling of manually configured VLANs
All	73016	456603	fsexport fails with "File does not exist" if a file or directory name contains a space character
All	73154	461529	usr/cvfs/bin/fsm ASSERT fail 'p_ip->i_idinode idi_nchildren >= 2' file dir.cline 1935
All	73246	461691	GUI needs option to transcribe inactive versions from a medium
All	73307	459454	Client file open request to the FSM hangs forever
All	73404	463603	fs_resourced segfault occurred upon exception in updateDelayedReleaseList leading to TSM abnormal termination
All	73410	464639	snac1 certain INHERIT flags naming schemes not consistent
All	73411	n/a	GUI should allow to defrag 100% of user blocks, similar to sgdefrag --blocks -1, ideally as checkbox option
All	73422	n/a	GUI StripeGroup defrag option doesn't accept values for "Blocks to move" > 9999



Operating System	Change Request Number	Service Request Number	Description
All	73430	456604	snprobe displays the wrong FSM address in client entries for a file system
All	73561	451310	Windows SNFS 6.x cannot connect with auth_secret if using fsforeignserver
Linux	45856	398545	fsm [14412]: StorNext FSS '2_man_fs[0]': dqns_load: GetIdentityNameBySID failed for key 0xb0e31fb0b0e31fb error 37
Linux	49439	3420306	sncompare should log a message related to running rebuild policy after updating ALL_COPIES flags
Linux	53486	425247	Appliance: Srvclg RAS tickets lost if HaShared file system not mounted
Linux	72136	406908	sgdefrag command fails due to a timeout
Linux	72537	449523, 453432	Linux kernel dump from f_cvp->cv_opencnt == 0 line 1478 file vnops.c
Mac OS	48194	3394288, 3384400, 3437750, 3499376, 3507270, 3635546, 3437750, 3499376, 3507270, 318356, 320308, 355814	Restriction of moving from a managed fs relation point to unmanaged problems (including quotas, trash can and recycle bin)
Windows	72481	433666	Microsoft MPIO in Round-Robin mode is causing continuous disk rescan
Windows	72873	454707	Real-Time QOS Token Leak

# Supported StorNext Upgrade Paths and Upgrade Considerations

## StorNext Software Upgrade Matrix

For information on which StorNext versions allow you to upgrade directly to this release, refer to the **StorNext Software Upgrade Matrix** section in the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

## Considerations for the StorNext File System Directories


On upgrades to StorNext 6.2, note that the attributes of many directories in the StorNext file system show much smaller sizes, even zero sizes, where these same directories showed non-zero sizes in previous releases of StorNext. This is expected behavior.

## Journal Size Guidelines

The absolute minimum Journal Size in StorNext 6.2 is 4 MB. If a file system is configured with a Journal Size smaller than 4 MB, the Journal Size must be increased prior to upgrading. The recommended Journal Size is 64 MB. New file systems must have a Journal Size of 64 MB or larger.

## Distributed Data Mover (DDM) Guidelines

Distributed Data Movers (DDMs) must be upgraded to the same version of StorNext that the Metadata Controller (MDC) is running.

 **WARNING:** Upgrades (such as platform, service pack, etc.) are intended to be done to all systems present in a given deployment. For example, if Xcellis, M660, M440, Pro Foundation, Artico, and G300 are present, they all must be upgraded. One appliance cannot be "left behind".

## Considerations When Upgrading NFS Server Nodes to StorNext 6.2

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 6.2, special consideration must be made for Linux NFS servers.

To prevent issues with mounted NFS clients, NFS clients must be unmounted prior to upgrading StorNext on the NFS server. If unmounting all NFS clients is not an option during the upgrade, Quantum suggests using the "compat32" mount option on NFS servers.

# Compatibility Between StorNext and Other Products

The following sections provide information regarding compatibility between this release and StorNext components and features.

## Appliance Controller

To view supported Appliance Controller software configurations, see *Appliance Controller Compatibility* available online at <http://qsupport.quantum.com/kb/Flare/Content/appliances/ACC/DocSite/Compat.htm>.

## Infiniband

StorNext 6.2 works with Infiniband SRP (SCSI RDMA Protocol) attached storage for Linux and Windows 2008R2.

## Lattus

See the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#) for information about compatibility between Lattus and StorNext 6.2.

---

**i Note:** Object Storage documentation is available online at <http://www.quantum.com/lattusdocs>.

## StorNext Web Services

StorNext Web Services enables you to run third-party application program interfaces (APIs) with StorNext. To view the latest commands supported by the StorNext Web Services, refer to the [StorNext 6 Web Services Guide](#) in the [StorNext 6 Documentation Center](#).

## Apple Xsan

Xsan is software that enables multiple Mac computers to concurrently access hundreds of terabytes of content on Xserve RAID or Promise RAID storage over high-speed Fibre Channel which allows you to share data faster and consolidate projects. Quantum supplements this solution with StorNext data management software, enabling Apple Xsan customers to use applications running on Windows, Linux, and UNIX with their Xsan and share content across more systems.

For information about compatibility between Apple Xsan and StorNext 6.2, refer to the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

## Supported Browsers

For information on browsers supported with the StorNext GUI for this release, refer to the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

For all other components and features, see the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

---

# General Considerations

This section provides information about items to consider for StorNext 6.2.

## Checksum Performance Considerations

**i Note:** Generating MD5 checksums is a CPU-intensive operation.

Current StorNext metadata controller and Mover hardware is able to calculate MD5 checksums at around 300 MB/s to 500 MB/s. For newer generation tape technology, the maximum throughput might exceed the rate at which the system can generate checksums. In this case, the MD5 checksum calculation will define the throughput of a single data movement operation. With multiple movement streams, MD5 calculations will be done in parallel across the streams and aggregation of performance will be seen.

---

# Upgrading Appliances

**⚠ Caution:** If you have a Replication, Deduplication, or Object Storage license, see Change Request [73945 on page 42](#) in the [StorNext Installation, Replication, HA, and Other Known Issues on page 41](#) section before you 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](#)

---

## Known Issues

The following sections list known issues in this release of StorNext, as well as associated workarounds, where applicable:

**i Note:** If you encounter one or more of the issues listed in this section, please contact Quantum Customer Support and report the issue(s) you encountered. Also inform the support representative whether you were able to successfully work around the issue(s) by using the provided workaround. Doing these things will help Quantum prioritize the order in which known issues are addressed in future StorNext releases.

## StorNext File System Known Issues

The table below lists known issues specific to the StorNext File System.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	54834	3505208, 3516356	<p>If a file is being copied to the StorNext file system using Windows Explorer and Windows Explorer crashes before it finishes copying all the data, the file might contain data blocks from old, deleted files. This problem occurs because Windows Explorer sets EOF to the size of the file before it writes the data to the file. This leaves a gap of uninitialized data in the file.</p> <p><b>i Note:</b> This problem can also occur with other programs that set EOF beyond the end of data.</p> <p>This problem does not occur if Windows Explorer encounters an error while writing the file; Windows Explorer will delete the partially written file.</p> <p><b>Workaround:</b></p> <p>To prevent this problem from occurring on StorNext, you can use the StorNext "client configuration" application's advanced mount option "Restrict Pre-allocation API" on Window systems and the "protect_alloc=yes" mount option on Linux systems. This option will set the unwritten parts of the file to zero. When this option is set, non-root users are unable to use the preallocation ioctl. This option also implies sparse=yes.</p> <p>For more information on this option, see the man page <code>mount_cvfs(8)</code>. The sparse option will introduce some overhead when using Windows Explorer. Before setting the <code>protect_alloc</code> option, see the sparse option in <code>mount_cvfs(8)</code> for a description of how it changes StorNext behavior.</p>
All	67363	n/a	<p>StorNext 5.4.0.x incorrectly allowed the <b>Unix ID Mapping</b> type to be set to <b>none</b> when the <b>Security Model</b> is set to <b>acl</b>. As a result, file systems fail to start when the <b>Unix ID Mapping</b> type is set to <b>none</b> when the <b>Security Model</b> is set to <b>acl</b>.</p> <p>Beginning with StorNext 6, the FSM does not start when this invalid combination of settings is used.</p> <p><b>Workaround:</b></p> <p>To prevent this issue, set the <b>Unix ID Mapping</b> to either <b>winbind</b> or <b>algorithmic</b> for any file system where the <b>Security Model</b> is set to <b>acl</b>. You can make the adjustment before or after upgrading.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	74202	n/a	<p>There are two basic scenarios, as follows:</p> <p><b>Scenario # 1</b></p> <p>When StorNext stops, it signals NAS to vacate the StorNext file systems that NAS is using. However, the <b>snnas_controller</b> attempts to automatically restart NAS services if these file systems remain available for too long.</p> <ul style="list-style-type: none"> <li>The restarting of the NAS services can resume use of StorNext resources and interfere with the operation of stopping StorNext.</li> <li>The interference can result in mounted StorNext file systems that do not have a running fsm process, thereby blocking access to the StorNext file system.</li> <li>In the event that StorNext fails to stop due to NAS resource usage, you might need to perform the following steps:             <ol style="list-style-type: none"> <li>Check to see if StorNext file systems are mounted:                 <pre>grep cvfs /proc/mounts</pre> </li> <li>Mount the HaSahred file system:                 <pre>mount /usr/adic/HAM/shared</pre> </li> <li>Start the fsm for each StorNext file system mounted on the MDC:                 <pre>cvadmin -e "start FSNAME on localhost"</pre> </li> </ol> <p><b>Note:</b> You might need to reboot the MDC if StorNext still fails to stop after restarting the required fsm processes.</p> </li> </ul> <p><b>Scenario # 2</b></p> <p>Maintenance Mode. An extension of <b>Scenario #1</b>, where NAS must vacate the file systems for an extended period of time.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>By resuming NAS services, those service processes may interfere with maintenance operations.</p> <p>To affect both scenarios, timers exist that you can adjust to extend the amount of time required between the <b>stornext stop</b> operation and before NAS services resume. The timers are controlled by the following <b>Controller Registry</b> values:</p> <ul style="list-style-type: none"> <li>• <b>nas.heartbeat.check_state_secs</b></li> <li>• <b>stornext_service.stop_period</b></li> </ul> <p>You can modify the timers by using the Controller shell interface <b>reg</b> commands.</p> <p>For example, to observe existing values, execute the following commands:</p> <pre data-bbox="626 867 1456 936">su sysadmin -c 'reg show nas.heartbeat.check_state_secs'</pre> <p>or</p> <pre data-bbox="626 1062 1456 1131">su sysadmin -c 'reg show stornext_service.stop_period'</pre> <p>Both timers default to 120 (seconds).</p> <p>For example, to set new values, where X is an integer, execute the following commands:</p> <pre data-bbox="626 1341 1456 1411">su sysadmin -c 'reg set nas.heartbeat.check_state_secs X'</pre> <p>or</p> <pre data-bbox="626 1537 1456 1606">su sysadmin -c 'reg set stornext_service.stop_period X'</pre> <p><b>Workaround:</b></p> <p>For <b>Scenario #1</b>, Quantum recommends that you extend the time period</p>



Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>to 10 minutes, a value of 600 for those two variables. The goal is to set a value that exceeds the amount of time necessary for typical StorNext shutdown.</p> <p>For <b>Scenario #2</b>, choosing an X value requires knowing how long maintenance will take and staying under that time. It is difficult to extend the maintenance period once its started. To avoid the complexity of controlling the maintenance period, Quantum recommends you perform the following steps:</p> <ol style="list-style-type: none"> <li>1. Stop NAS by executing the following command: <pre data-bbox="683 709 1458 779" style="background-color: #f0f0f0; padding: 5px;">/usr/cvfs/lib/snnas_control stop</pre> </li> <li>2. Stop the Controller by executing the following command: <pre data-bbox="683 905 1458 1010" style="background-color: #f0f0f0; padding: 5px;">/usr/local/quantum/bin/sml_service_tool stop snnas_controller</pre> </li> </ol>
All	74299	n/a	<p>Beginning with StorNext 6.2, when the <b>snfsdefrag</b> command skips any file for any reason and the verbose option is not supplied, the follow message is displayed:</p> <pre data-bbox="626 1178 1458 1283" style="background-color: #f0f0f0; padding: 5px;">Processing one or more files failed, run Verbose option for details</pre> <p>When this occurs, the command also returns a non-zero exit status, suggesting that the command encountered an error. However, there are normal conditions where <b>snfsdefrag</b> can skip a file. For example, a file can be skipped if it does not need to be defragmented because its extent count is already low enough. Also, <b>snfsdefrag</b> skips device files and symbolic links. So the fact that <b>snfsdefrag</b> is reporting "failures" in this case is unexpected.</p> <p><b>Workaround</b></p> <p>You can ignore the error and non-zero exit status.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Mac OS	66948	322824, 336945	<p>If you access StorNext file systems from Apple Xsan clients, then you might encounter I/O error messages in the system log that do not contain details about real I/O errors detected on the Xsan client.</p> <p><b>Workaround</b></p> <p>If you encounter the errors on an Xsan client, contact Apple.</p>

---

## StorNext Storage Manager Known Issues

The table below lists known issues specific to StorNext Storage Manager.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	43320	1581004	<p>File retrieves from media to disk can be suboptimal for fast tape drives like the Oracle STK T10K drives. This scenario can occur when the retrieve event is initiated on a host that is different from the host running the mover process, which requires the use of synchronous direct I/O.</p> <p><b>Workaround:</b></p> <p>To work around this issue and achieve optimal performance for both file stores and retrieves with the T10K drives, increase the default I/O size used by the mover process and make the mover process use asynchronous buffered I/O when the use of synchronous direct I/O is not required, using the following steps:</p> <ul style="list-style-type: none"> <li><b>i Note:</b> This workaround might also help improve the performance of the faster LTO drives by updating the <b>FS_LTO_BLOCK_FACTOR</b> system parameter.</li> <li><b>i Note:</b> Changes to <b>FS_xxx_BLOCK_FACTOR</b> only affects tapes formatted after the change.</li> </ul> <ol style="list-style-type: none"> <li>1. Change the <b>FS_T10K_BLOCK_FACTOR</b> system parameter from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre>FS_T10K_BLOCK_FACTOR=32;</pre> </div> <ul style="list-style-type: none"> <li><b>i Note:</b> The T10K default I/O block size is 512 KB or 8 * 64 KB. With the block factor changed to 32, the new T10K I/O block size will be 2 MB or 32 * 64 KB. Presently, the <b>FS_T10K_BLOCK_FACTOR</b> system parameter must not be set to a value that exceeds 32.</li> </ul> </li> <li>2. Restart Storage Manager to ensure the change in <b>Step 1</b> goes into effect: <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre># tsmstop # tsmstart</pre> </div> </li> <li>3. Verify the <b>FS_T10K_BLOCK_FACTORsysparm</b> contains the new value:</li> </ol>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<pre data-bbox="695 352 1243 422"># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> <p data-bbox="634 485 1446 667">4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients.</p> <p data-bbox="634 569 1438 667">5. Modify <code>/etc/fstab</code> on the MDCs and the DDM clients to use the <code>auto_dma_write_length</code> and <code>auto_dma_read_length</code> mount options as follows:</p> <pre data-bbox="695 730 1386 831">snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0 0</pre> <p data-bbox="634 898 1146 930">6. Unmount and re-mount your file systems.</p> <p data-bbox="634 951 1365 982">7. Use new T10K media to store a copy of the file from the disk.</p> <p data-bbox="683 993 1458 1192"><b>i Note:</b> <b>Step 7</b> is very important; when the new copy is made to the new tapes, the new tapes are labeled with a 2 MB block size, which is used for subsequent writes or reads to and from the media. Tapes on which <b>fsformat</b> was run before the change will use the block factor in use at that time. This change will not impact those tapes.</p>
All	46693	n/a	<p data-bbox="602 1230 1458 1360">Executing the command <b>snbackup -s</b> while a full or partial backup is running might result in a message that <code>/usr/adic/TSM/internal/locks/backup.1f</code> is in an invalid format.</p> <p data-bbox="602 1371 1458 1440">This is due to the <b>snbackup -s</b> process reading the <b>backup.1f</b> status file while the backup process is updating it.</p> <p data-bbox="602 1451 781 1482"><b>Workaround:</b></p> <p data-bbox="602 1493 1458 1591">Ignore the message; to clear-up the process, re-execute the command <b>snbackup -s</b> (provided that the backup is not writing to the <b>backup.1f</b> status file while <b>snbackup -s</b> is trying to read it again).</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	47833	n/a	<p>When copying files between media using the CLI command <b>fsmedcopy</b>, the file is not re-segmented to match the segment size of the destination media. Rather, the original segments are copied to the target media type and the distribution of segments across destination media will, therefore, be the same as the distribution on the source media.</p> <p><b>i Note:</b> This behavior might cause file data segment distribution to be sub-optimal on the destination media.</p> <p><b>Workaround:</b></p> <p>Currently, a workaround does not exist for this known issue.</p>
All	69265	n/a	<p>Your DDMs might experience a timeout if you try to connect to the database. The issue is identified by an error log in <b>/usr/adic/TSM/logs/tac</b> which contains the text:</p> <pre>Process fs_moverd on &lt;host&gt; timed out trying to connect to the database. This usually indicates network connectivity trouble. Try increasing the timeout value by setting the connect_timeout value in /usr/adic/mysql/my.cnf. The default setting is 10 seconds so the new value should be larger.</pre> <p><b>Workaround:</b></p> <ol style="list-style-type: none"> <li>1. Increase the database connection timeout value by adding the following line to <b>/usr/adic/mysql/my.cnf</b> under the section labeled <b>[mysqld] connect-timeout=240</b>.</li> <li>2. Cycle the Storage Manager in order to pick up the updated timeout value.</li> </ol>
All	69341	n/a	<p>If you have the IBM APFO driver installed and configured, then when you perform an <b>fsmedread</b> operation of a partial tape block from a full tape block, the operation can fail with <b>errno=12</b>.</p> <p><b>i Note:</b> This issue affects all IBM APFO versions 3.0.19 and earlier, and has an impact primarily on disaster recovery procedures.</p> <p><b>Workaround:</b></p> <p>To correct this, perform an <b>fsmedread</b> operation without the IBM APFO driver.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	72993	452722	<p>When DDM is enabled for non-primary systems, kernel error messages are logged for <b>reservation conflict</b> because the Primary MDC mounts the tape which sets the reservation to itself and the DDM sets the reservation to itself once the tape is ready.</p> <p>This issue causes a benign message in the <code>/var/log/messages</code> file for <b>reservation conflict</b>. Since the unmounting of the tape also resets the reservation back to the primary MDC, this message is generated the next time a DDM (non-primary MDC) accesses a tape.</p> <p><b>i Note:</b> This issue can result in a large amount of log messages on any machine running DDMs.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, you can filter and drop the messages so they are no longer logged, as follows.</p> <p><b>i Note:</b> Reservation conflicts that are a problem also produce RAS alerts, so you can also drop these errors.</p> <p>Create the following rules on all DDM clients, based on the OS, to drop the messages from <b>rsyslog</b>:</p> <p><b>For RedHat 6</b></p> <pre># echo ':msg, contains, "reservation conflict" ~' &gt; /etc/rsyslog.d/ignore-reservation-conflict.conf  # service rsyslog restart</pre> <p><b>For RedHat 7</b></p> <pre># echo 'if \$programname == "kernel" and \$msg contains "reservation conflict" then stop' &gt; /etc/rsyslog.d/ignore-reservation-conflict.conf  # systemctl restart rsyslog.service</pre>

## StorNext GUI Known Issues

The table below lists known issues specific to the StorNext GUI.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69360	n/a	<p>Using <b>autofs</b> to mount a StorNext file system on an MDC is not supported when the same file system also has a native mount point.</p> <p>For example, if the StorNext file system <b>snfs1</b> is mounted as <b>/stornext/snfs1</b>, then the MDC should not also have an <b>autofs</b> configuration that mounts it on the MDC in another location such as <b>/space/snfs1</b>. Doing so, causes the <b>fsCheckAffinities</b> and <b>fsCheckTsmFilesystemConfig</b> health checks to fail and generate RAS tickets.</p> <p>Additionally, this might cause the StorNext GUI to fail unexpectedly for certain operations.</p> <p><b>Workaround:</b></p> <p>There is currently no workaround for this issue. If you experience this issue, contact <a href="#">Quantum Technical Support</a>.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69958	373823	<p>There is a known issue where stripe group expansion using the StorNext GUI can fail and requires manual intervention to restore file system operation.</p> <p>Stripe group expansion allows an additional disk (LUN) to be added to an existing stripe group, growing the file system user data space.</p> <p>The other way to add space to a file system is to create a new stripe group and add this to the file system.</p> <p>When the GUI executes stripe group expansion, it stops the file system, modifies the configuration file and runs <b>cvupdatefs</b>. If <b>cvupdatefs</b> takes more than 5 minutes to complete, the GUI kills the <b>cvupdatefs</b> utility and reports an error. At this point the file system does not start because the configuration file does not match the current state of the metadata.</p> <p><b>Workaround:</b></p> <p>To address this issue, you can do one of two things:</p> <ul style="list-style-type: none"> <li>• Optimally, you would not attempt the stripe group expansion at all. Instead add a new stripe group to the file system.</li> <li>• If stripe group expansion is deemed necessary, use the <b>cvupdatefs</b> CLI directly instead of through the GUI. This method does not have any time limitations.</li> </ul> <p>In the case that the StorNext GUI attempt was made and hit the 5 minute timeout, file system operations can be resumed by performing the following:</p> <ol style="list-style-type: none"> <li>1. Restore the previous version of the configuration file. This can be found in the following directory: <div data-bbox="695 1360 1242 1394" data-label="Text"> <pre>/usr/cvfs/data/&lt;fs&gt;/config_history</pre> </div> </li> <li>2. Run the <b>cvfsck</b> utility to verify and potentially correct the metadata.</li> <li>3. Start the file system.</li> </ol>



Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	47954	n/a	<p>The Safari browser becomes unresponsive when you attempt to configure an Email server using the StorNext GUI.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, perform the following procedure:</p> <ol style="list-style-type: none"> <li>1. Shut down the Safari browser window(s).</li> <li>2. Restart the Safari browser, and then retry the operation.</li> <li>3. Uncheck the <b>Verify SMTP Server Connectivity</b> box, and then retry the operation.</li> <li>4. Set <b>Authentication</b> to <b>NONE</b>, and then retry the operation.</li> <li>5. Disable the Safari <b>User names and passwords</b> AutoFill under <b>Safari &gt; Preferences &gt; AutoFill</b>, and then retry operation.</li> </ol>

## StorNext Installation, Replication, HA, and Other Known Issues

The table below lists known issues specific to StorNext installations, data replication, HA systems, and other areas.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	68849	n/a	<p>After an appliance firmware upgrade, you might be unable to use previously functioning tape devices because the <code>lin_tape</code> device driver was automatically unloaded during the upgrade.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, rebuild the <code>lin_tape</code> device driver as shown in the following example:</p> <pre>rpm -e lin_taped rpm -e lin_tape rpmbuild --rebuild /root/lin_tape-1.76.06-1.src.rpm rpm -ivh /root/rpmbuild/RPMS/x86_64/lin_tape-1.76.06-1.x86_64.rpm rpm -ivh /root/lin_taped-1.76.0-rhel6.x86_64.rpm</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	73557	n/a	<p>Beginning with StorNext 6.2.0, if you change the system clock backwards while MySQL is running and then try to shutdown StorNext, the MySQL database may block when trying to stop until the current time matches the time just prior to changing the time on the system.</p> <p><b>i Note:</b> If you use a Quantum appliance and use the StorNext GUI to change the time, there is no issue.</p> <p><b>Workaround:</b></p> <p>This workaround applies to the following systems:</p> <ul style="list-style-type: none"> <li>• If you use a customer supplied system.</li> <li>• If you use a Quantum appliance and do not use the StorNext GUI to change the system clock.</li> </ul> <p>To workaround this issue, perform the following procedure.</p> <ol style="list-style-type: none"> <li>1. Shut down StorNext.</li> <li>2. Change the time.</li> <li>3. Restart StorNext.</li> </ol> <p><b>i Note:</b> If shutting down StorNext is not an option, then wait for the system time to match the time prior to the time change, and then shut down to avoid this issue. For example, if you adjust the system time back one hour, then you must wait one hour before you halt or reboot your system.</p>
All	73945	n/a	<p>If you have a Replication, Deduplication, or Object Storage license installed in <code>/usr/cvfs/config/license.dat</code>, then during HA upgrade time, the first time one of the nodes is activating and becoming primary on the new StorNext version, <b>snpolicyd</b> could be restarted. The restart might take up to 5 minutes to stop.</p> <p><b>Workaround:</b></p> <p>Execute the following commands on both nodes in order to shorten the <b>snpolicyd</b> activation time before upgrading to a new version of StorNext:</p> <ol style="list-style-type: none"> <li>1. <code>mkdir -p /usr/cvfs/update/config</code></li> <li>2. <code>touch /usr/cvfs/update/config/objs.conf</code></li> <li>3. <code>touch /usr/cvfs/config/objs.conf</code></li> </ol>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	70282	n/a	<p>The <a href="#">Stornext Connect Utilization App Version 1</a> does not recognize the HGST ActiveScale™ P100 (Quantum Lattus P100) integrated object storage system, and does not incorporate capacity or data movement associated with the HGST ActiveScale™ P100 (Quantum Lattus P100) system as a target.</p> <p><b>i Note:</b> This does not otherwise affect the functionality of the HGST ActiveScale™ P100 (Quantum Lattus P100) integrated object storage system.</p> <p><b>Workaround:</b> There is currently no workaround for this issue. If you experience this issue, contact <a href="#">Quantum Technical Support</a>.</p>

# Contacting Quantum

## Contacts

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

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

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>

## 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)