



# StorNext 5 release 5.3

## Release Notes

---

<b>Product</b>	StorNext® 5 release 5.3
<b>Date</b>	December 2015

---

Made in the 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**

© 2015 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law.

### **TRADEMARK STATEMENT**

Quantum, the Quantum logo, DLT, DLTtape, the DLTtape logo, SuperLoader, Scalar, StorNext, and DXi are registered trademarks of Quantum Corporation, registered in the U.S. and other countries. Preserving the World's Most Important Data. Yours., Backup. Recovery. Archive. It's What We Do., the DLT logo, DLTSage, Dynamic Powerdown, FastSense, FlexLink, GoVault, MediaShield, Optyon, Pocket-sized. Well-armored, SDLT, SiteCare, SmartVerify, StorageCare, Super DLTtape, and Vision are trademarks of Quantum. LTO and Ultrium are trademarks of HP, IBM, and Quantum in the U.S. and other countries. All other trademarks are the property of their respective companies. Specifications are subject to change without notice.

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



---

# Contents

What's New in StorNext 5 release 5.3 . . . . .	3
StorNext Compatibility . . . . .	12
Supported StorNext Upgrade Paths and Upgrade Considerations . . . . .	13
Compatibility Between StorNext and Other Products . . . . .	15
Supported System Components . . . . .	16
Upgrading Appliances . . . . .	16
Known Issues . . . . .	18
Contacting Quantum . . . . .	34

---

## What's New in StorNext 5 release 5.3

---

### Purpose of this Release

StorNext 5 release 5.3 delivers an array of innovative new features and important enhancements. For complete details, see [New Features and Enhancements in StorNext 5 release 5.3](#) on page 3.

Security and cost are persistent barriers that often prevent users from taking advantage of new cloud based services. In this release of StorNext, the addition of compression and encryption, used independently or together, end-users wanting to utilize the integrated simplicity and cost-effectiveness Quantum Q-Cloud services can be assured data is protected and efficiently moved off-site.

In this release of StorNext, Q-Cloud has been enhanced to allow storage administrators to quickly recover StorNext backup images and help get StorNext solutions back into production.

In this release of StorNext, newly released clients like Microsoft Windows10 are fully supported, and Apple Spotlight search queries now quickly report on digital assets residing on StorNext file system volumes.

To keep pace with the explosive growth and near infinite retention of today's data, this release of StorNext includes support for the latest generation of **Linear Tape-Open (LTO)** technologies, IBM LTO-7, which doubles the capacity and performance of prior versions.

StorNext 5 release 5.3 also resolves several customer reported issues listed within the section [Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.3](#) on page 7 in this document.

---

### New Features and Enhancements in StorNext 5 release 5.3

#### Introducing Q-Cloud Vault™

StorNext 5 release 5.3 includes a new Q-Cloud offering called Q-Cloud Vault™, a storage destination which gives customers a way to store large, unchanging, and rarely-retrieved large data files off site, cost-effectively. StorNext Storage Manager has been enhanced to support the use of public cloud services as a storage tier. StorNext 5 release 5.3 extends Quantum Q-Cloud services to include a very low cost "cold storage vault" destination that allows users to copy files to a secure storage tier that is managed by Quantum.

Q-Cloud Vault™ policies are supported on all metadata controller and the new Xcellis workflow director platforms supported by StorNext 5 Storage Manager. There are no upfront fees for using Vault, and no additional software or cloud gateway hardware are required. Q-Cloud Vault™ is very efficient and provides "pay for what you use" storage.

One of the potential usage charges involve file retrieval requests sent to the service. The time between a request and first-byte delivery can be hours.

**IMPORTANT:** Retrieve charges are based on the peak retrieval request rate (Terabytes of requests per hour) for the month. You should carefully control how quickly your request files from Q-Cloud Vault™ to avoid the retrieve fee. Since Q-Cloud Vault™ is based on Amazon Glacier, it is very helpful to understand how Amazon calculates retrieval charges, explained at:

[http://aws.amazon.com/glacier/faqs#How\\_will\\_I\\_be\\_charged\\_when\\_retrieving\\_large\\_amounts\\_of\\_data\\_from\\_Amazon\\_Glacier](http://aws.amazon.com/glacier/faqs#How_will_I_be_charged_when_retrieving_large_amounts_of_data_from_Amazon_Glacier). The Q-Cloud Vault™ fee is calculated the same way, with a markup.

Pairing Q-Cloud Vault™ with an on-site tape archive is the ideal way to maintain high-speed, free retrievals and secure, low-cost, off-site file protection.

For additional information, see section **Q-Cloud** in the *StorNext 5 User's Guide* available on-line at <http://www.quantum.com/sn5docs>, or contact a Quantum sales representative.

### Support for Disaster Recovery using Q-Cloud

Beginning with StorNext 5 release 5.3, support has been added to recover the snbackup image (from Q-Cloud) by using the **Q-Cloud Access Identifier (QCAI)** and the **Q-Cloud Product Key (QCPK)**. Quantum provides both "keys" to all Q-Cloud users. With StorNext 5 release 5.3, there is not a need to contact Quantum support in order to recover the snbackup image.

### Support for Encryption and Compression with Q-Cloud Archive™ and Q-Cloud Vault™ Policies

It is vital that cloud based storage services use resources efficiently while also providing end-to-end data security. StorNext Storage Manager support for Q-Cloud has been enhanced and now includes the ability to encrypt and compress data being sent to Q-Cloud Archive™ or Q-Cloud Vault™ destinations.

---

**Note:** Client side encryption and compression are configured at the Storage Policy level. If both features are configured, data is compressed before being encrypted and transmitted to Q-Cloud.

---

StorNext 5 release 5.3 delivers LZ4 compression and AES-CBC-256 encryption. For additional information, see sections **Storage Policies** and **Client-side Encryption** in the *StorNext 5 User's Guide* available on-line at <http://www.quantum.com/sn5docs>.

### StorNext Storage Manager is Quattro "Ready"

StorNext Storage Manager is qualified and ready for Quantum "Quattro" tape libraries.

### Support for IBM LTO-7 Tape Devices

The continued rapid growth of digital data demands StorNext Storage Manager support for ever faster and higher capacity tape technologies. **Linear Tape-Open (LTO)** is a proven technology for preserving data. StorNext 5 release 5.3 adds support for IBM LTO-7 tape devices, including support for **Write Once Read Many (WORM)** use, further ensuring data cannot be altered or overwritten.

## Previously Optional Data Movement Integrity Check Option (checksum) Included

Included at no charge (beginning with StorNext 5 release 5.3), this enhancement allows storage administrators to configure policies that ensures data stored to tape or Lattus Object Storage has not changed when retrieved.

## Support for Additional Operating Systems and Platforms

StorNext 5 release 5.3 adds support for additional operating system versions of Microsoft Windows, Debian, and Ubuntu. With this release, StorNext gains compatibility with:

Operating System	Version
Microsoft Windows	Version 10
Debian	Version 7.8 (client only) Version 8.x (client only)
Ubuntu	14.04.3 Long-Term Support (LTS)

---

**Note:** Red Hat Enterprise Linux 5 (all kernels/releases) is no longer supported beginning StorNext 5 release 5.3 (or later).

---

---

**Caution:** The Debian 7 clients included in StorNext 5 release 5.3 do not launch correctly. Contact Quantum Technical Support and request the corrected Debian 7 client package. The updated Debian 7 clients will be included in the first StorNext 5.3 maintenance release.

---

For additional information, refer to the *StorNext 5 Compatibility Guide* available here: <http://www.quantum.com/sn5docs>.

## Support for Additional Web Services

Growing the StorNext solution ecosystem requires continued improvements to web services. StorNext 5 release 5.3 provides support for asynchronous request and improved progress status information about stores and retrieves.

## Support for Apple Spotlight Search

Apple Mac OS X clients (SAN attached support for Mac OS X version 10.10 and version 10.11, with LAN attached client support for Mac OS X version 10.11) can use the Spotlight search feature to find files residing on (Linux) StorNext metadata controllers.

## Comprehensive Support for Apple Xsan Configurations

StorNext 5 FX clients are qualified to fully support the following versions of Apple Xsan:

- Xsan 2.2
- Xsan 2.3
- Xsan 3.0
- Xsan 3.1
- Xsan 4.1

## Improved Lattus Object Storage Deletion Performance

StorNext Storage Manager management has been improved to better support very large Lattus Object Storage configurations. Lattus Object Storage delete requests are now multi-threaded and can be streamed to a Lattus system for quicker removal.

## Alternate Store Location (ASL)

StorNext **Alternate Store Location** is a highly configurable feature that can automatically copy files from a "local" StorNext file system to a "remote" site at the same time the files are created to tiered storage, like tape, at the main site. The **Alternate Store Location** feature copies new as well as previously created files to the remote location.

Multiple "local" sites can copy files to a single "remote" site. **Alternate Store Location** is enabled per Storage Manager Class Policy per file system. By default, all files copied use secure FTP (SCP) to the remote destination.

To configure the **Alternate Store Location** feature, contact Quantum Technical Support.

## Metadata Controller Support for Red Hat Enterprise Linux 7

StorNext 5 release 5.3 provides up-to-date support for non-appliance based solutions built using Red Hat Enterprise Linux 7.

---

**Note:** Prior to installing StorNext, the installation process verifies the `psmisc` package is installed. Consult your operating system distribution for information on how to install the `psmisc` package.

---

## Support for all M-Series, Artico and Pro Solution Appliances

All M-Series, Artico and Pro Foundation appliances can be upgraded to StorNext 5 release 5.3.

---

**Note:** StorNext 5 release 5.3 is the final feature release for M330 appliances. Support for M330 appliances is available 30 days after general availability. All StorNext 5 release 5.3 maintenance releases support M330 appliances.

---

## Supported Upgrade Paths

There are several upgrade paths to a StorNext 5 release 5.3 solution, including:

- Configurations built using StorNext 4.7 (or earlier) can purchase the upgrade to StorNext 5 release 5.3.
- Existing customers with an M-Series appliance or Pro Solution with an active support contract can upgrade to StorNext 5 release 5.3 at no charge.
- Existing customers without an M-Series appliance that have an active support service contract may purchase a license upgrade for StorNext 5 release 5.3.

## Change to Linux "df" Output

Beginning with StorNext 5 release 5.3, the **df** output for StorNext file systems has changed slightly for the Linux platform. Previously, the **Filesystem** column contained a string that combined a device name with the file system name.

For example:

```
Filesystem          Size  Used Avail Use% Mounted on
/dev/cvfsctl11_snfs1 64G  4.2G  60G   7% /stornext/snfs1
```

In the current release, the output has been simplified so the **Filesystem** column only contains the file system name.

For example:

```
Filesystem          Size  Used Avail Use% Mounted on
snfs1              64G  4.2G  60G   7% /stornext/snfs1
```

Any customer-written scripts that rely on the old output format should be updated accordingly.

## Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.3

The following table lists the customer reported issues and enhancements addressed in StorNext 5 release 5.3.

Operating System	Change Request Number	Service Request Number	Description
All	2551	156132, 1469140	Enhancement to the command <b>cvcp</b> so that it is aware of zero-fill extents
All	16037	628036	Enhancement to the command <b>cvcp</b> so that it is easier to cleanup after it is interrupted

Operating System	Change Request Number	Service Request Number	Description
All	16498	713439	Enhancement to display all SCSI sense data bytes
All	30510	1089486	The <b>FSMDBGx</b> debug messages are not identifying which routine the messages are printed from
All	31735	1156698	The command <b>msmup</b> should include a check for <b>ssi</b> and <b>t_parent</b>
All	32613	1200856, 1454150	Additional feature request for the command <b>fsmedcopy</b> to accept the <b>-i</b> option which causes media in a vault to not be reported on the fragmentation report
All	34969	3411976 1432720	FSM connections are not closing, causing <b>MaxConnections</b> limit to exceed
All	36531	n/a	The TSM <b>fs_sysparam</b> entry <b>MAX_STORE_SIZE</b> does not work as described when set to <b>0</b>
All	37336	1395540	The GUI reports an error on the library partition instead of displaying the <b>xdi_sncli</b> errors
All	38146	1603352	Central Config causes <b>fsmpm</b> to hang if there are errors in the <b>nss_ctl.xml</b> file
All	38291	1611238, 1616074, 3570382, 3568266	Configuration of standby as DDM does not work after failover
All	40579	1494608, 3518718	The command <b>snbackup</b> fails because <b>quantumdb</b> can not get to the <b>.SM_backup</b> staging directory
All	41029	1502240	<b>fsm_starter</b> thread needs to be resilient against time warp
All	41267	1517128	kernel: CVFS ASSERTION FAILED: f_path != NULL line 1659 file
All	41882	1521382	Option to <b>dbstore</b> can speed up linter2mysql dbstores
All	42512	3565330	On MSM startup if a scsi archive is issuing coming ready sense code, MSM will fail to start



Operating System	Change Request Number	Service Request Number	Description
All	44523	1622028, 1623008, 1619618, 1621444, 1606208, 1613512, 1611260, 1610606, 1615548, 1615552, 1616032, 1619446, 1620970, 1620976, 1622416, 1622270, 1626322, 1649238, 3445930, 3442938, 3527490, 3557628	Xsan clients trigger RAS events when rebooted
All	45188	1616850	Need synchronization between <b>fsmpm</b> startup and device scan
All	47022	2200556, 3304342, 2201034, 3434332	Operating system issue causing a tape file corruption (duplicate I/O block written) when <b>pse_snapshot</b> is executed
All	47219	3361648	Dismount failed due to check condition ASC=0x28/ASCQ=0x1 (Insert/Eject station opened and closed)
All	47867	3378182, 3459756	The command <b>cvdb -b</b> shows negative values
All	48334	3388550	The commands <b>fsversion</b> or <b>fsfileinfo</b> should have an option to show the store time ( <b>ADDDATE</b> ) of a file
All	48769	497240, 1601286, 3400902	The command <b>cvcp</b> should continue copying on certain errors
All	48770	497240, 1601286, 3400902	The command <b>cvcp</b> should not copy files on which it has had stat errors
All	48876	3395062	Add check of number of blocks when executing the command <b>fsrcopy</b>
All	50354	3602738	Modifications to <b>/usr/adic/MSM/config/envvar.config</b> are not preserved on upgrade

Operating System	Change Request Number	Service Request Number	Description
All	51962	3565472	With RHEL6.5 or SLES11SP2 initial mount fails after <b>cvmkfs</b> if dm-multipath is in use
All	52921	3472004	The command <b>fsmedcopy</b> hangs when proxy child death signal received before all mover statuses
All	54599	3502086, 3568640	Web Services becomes unresponsive and crashes with no warning or RAS tickets being logged
All	55179	3505084	Allow to use TCP window size auto tuning for distributed LAN client network
All	55506	3511352	Tape drives taken offline when files being stored to LTFs have non-UTF-8 compliant pathnames
All	55518	3521280	The command <b>cvfsck</b> trace file fills up the Root Partition
All	55601	3459756	The command <b>snbackup</b> does not report failure when copying <b>metadb</b> to ref directory on full file system
All	55648	3518748, 3547512	File recovery fails in the GUI but works via CLI if there are unicode characters in the path or file name
All	55964	3508654	The command <b>sncompare qt check_dirpath</b> test reports directories that have been removed
All	55965	3508654	The command <b>sncompare -c qt check_not_all_copies_made</b> test does not report errors on files with multiple active versions
All	56230	3534954	Enhance the license library to use libxml2
All	56607	n/a	Check File System runs <b>cvfsck -T /tmp</b> unless default temp dir value of <b>/usr/adic/tomcat/temp</b> is changed on page
All	56877	3546218	The command <b>sncompare</b> is killed or hangs if there are too many entries in the <b>DIRPATH</b> table
All	56969	3604020	Assert in <b>pick_slice()</b> : <b>slice != NO_SLICE</b>
All	57143	3520028	Enhancement to the GUI to remove old files from <b>/usr/adic/gui/logs/jobs</b> directory
All	57227	3560454	The command <b>snbackup</b> wrongly compares <b>bkplist cpyid</b> with <b>filecomp cpyid</b> since the entry in <b>filecomp</b> is a <b>bitmask</b>
All	57350	3579556	<b>syncha.pl</b> : <b>syncha</b> is suspended by lock file after stopping <b>cvfs</b>
All	57416	156132	Enhancement to <b>cvcp</b> to be aware of zero-fill extents
All	57624	3559792, 3563534, 3563300	Enhancement to the GUI to verify the secondary node has the <b>HA_shared fs</b> mounted
All	57654	n/a	The command <b>cvfsid</b> does not look at <b>license.dat</b> file
All	57761	3592918	RH7.1 MDC panic : kernel BUG at fs/inode.c:512!

Operating System	Change Request Number	Service Request Number	Description
All	57811	n/a	Performance improvements to the command <b>fsclean</b> and object store
All	57852	3575246	Finding columns in StorNext GUI screens that appear to sort by string value instead of numeric
All	57966	3575780,3 561090, 3595392, 3605294	Xsan clients lock up when saving file in Photoshop if MDC is running StorNext 5.x
All	58016	3562964, 3569684	Deadlock while walking directory btree using <code>Dirreaddir_old()</code> with high level of concurrency
All	58189	3580372	FSM PANIC <code>In_rele</code> : Metadump restore mishandles an inode if <code>InodeFlagSnbtCvting</code> flag is set.
All	58297	3580372	The command <b>cvfsck</b> should deal with left over conversion and restoring flags
All	58319	3580372	512-byte inode filesystems can lose inode updates sent to metadump
All	58987	1622028, 1623008, 1619618, 1621444, 1606208, 1613512, 1611260, 1610606, 1615548, 1615552, 1616032, 1619446, 1620970, 1620976, 1622416, 1622270, 1626322, 1649238, 3445930, 3442938, 3527490, 3557628	Downgrade the certain high-volume RAS tickets to Sev 2.
All	59033	3580790	5.x meta data conversion can't handle FS with <code>sb_InodeVersion = 0x208</code> - 5.0 small inode pre-31 bit directory
All	59623	3593150	Stores of large files to Lattus fail if Lattus media is write-protected
Linux	39231	1433008	Enhancement to the command <b>fspolicy -m</b> option in the emergency store

Operating System	Change Request Number	Service Request Number	Description
Linux	43246	1568166, 1631128, 3575170	Tertiary Storage Manager to Lattus: Consider less sensitive <b>offline</b> of <b>iopaths</b> and media after connectivity faults
Linux	46081	1635978	Enhancement to be more verbose when <b>force_sync_io</b> is set
Linux	56362	n/a	Added the Alternate Store Location feature
Other	55765	3560454	The command <b>snbackup</b> manifest files are not written to copy 2 on Lattus
Windows	49165	3414076, 3531344	<b>CvRename</b> should have same logic as <b>CvShRename</b> to determine, if it requires sending a DM_EVENT_RENAME to SM
Windows	56791	3547320	Windows mount failure message <b>StorNext Error (127) Mounting XXXX</b> is not meaningful
Windows	58448	3576618	Windows preallocations over-allocate causing file to use more space than needed

---

## StorNext Compatibility

For information on StorNext 5 compatibility with operating systems, kernel versions, hardware platforms, drives, libraries, StorNext Appliances, StorNext client interoperability, and other compatibility items, see the *StorNext 5 Compatibility Guide*. SNAPI, Partial File Retrieval, and Apple Xsan compatibility information is provided in separate documents.

---

### Quantum OS Upgrade Support Policy

StorNext supports any security or functional bug update that applies to the current StorNext-supported Red Hat update level or SuSE Linux Service Patch. StorNext does **not** support updating the update level or service patch beyond the currently supported levels shown in the *StorNext 5 Compatibility Guide* available here:

<http://www.quantum.com/sn5docs>

---

### StorNext and Linux Interoperability

Newer versions of the Linux **tail** command leverage the **inotify** mechanisms within Linux. The **inotify** mechanisms in Linux are not triggered by file updates coming from other StorNext nodes.

When using the **tail** command on files located in StorNext, Quantum recommends using the following option:

```
---disable-inotify
```

Recommended usage:

```
tail ---disable-inotify -f filename
```

---

## 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 5 Compatibility Guide* available here:

<http://www.quantum.com/sn5docs>

---

### Considerations for the StorNext File System Directories

---

On upgrades to StorNext 5, it may be noted 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 StorNext 4.x. This is expected behavior.

---

### Journal Size Guidelines

---

The absolute minimum Journal Size in StorNext 5 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.

---

### Buffer and Inode Cache Changes

---

In StorNext 5, the default values for `bufferCacheSize` has been increased from 32 MB to 256 MB, and the default value for `inodeCacheSize` has increased from 32768 to 131072.

- Your system requires an additional 8 GB of memory be available for use by the MDC operating system and the StorNext services.
- Quantum recommends provisioning 8 GB of memory per configured file system for the FSM `bufferCache`.
- Your system requires 8 GB of `innodb_buffer_pool_size`, but Quantum recommends up to 40 GB for larger installs, if memory is available.

---

### Distributed Data Mover (DDM) Guidelines

---

Distributed Data Movers (DDMs) must be upgraded to StorNext 5 when the Metadata Controller (MDC) is operating on StorNext 5.

---

**WARNING:** Upgrades (such as platform, service pack, etc.) are intended to be done to all of the Lattus and metadata systems present in a given deployment (for example, if M662, A10, C10, S10 are present, they all must be updated, one appliance cannot be "left behind").

---

---

## Considerations When Upgrading NFS Server Nodes to StorNext 5

---

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 5, special consideration must be made for Linux NFS servers.

In order 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.

---

## Database Schema Update During Upgrades

---

Database schema updates are applied to Storage Manager when upgrading from StorNext 4.3.x (and above) to StorNext 5. The M660 appliance can achieve approximately one hour for every hundred million entries in the filecomp tables. Smaller appliances and Software Only configurations may take considerably longer depending on CPU and memory availability.

StorNext file systems are accessible while the database schema is being updated, but Storage Manager functionality (including stores and retrieves) will be offline.

Do NOT interrupt StorNext services while the database is being updated. Interrupting the database schema update could result in an inconsistent database, and may require assistance from Quantum Support to repair or restore the database.

Use the following commands to determine the number of filecomp entries on the StorNext primary node:

- 1 List the managed filesystems configured.

```
mysql -e "select Device_key, Path from tmdb.devdb_v;"
```

- 2 For each <Device\_key> number listed display a count of the number of entries in the corresponding filecomp table:

```
mysql -e "select count(*) from tmdb.filecomp<Device_key>;"
```

---

# Compatibility Between StorNext and Other Products

This section describes various interactions between this release and StorNext components and features.

---

## Infiniband

---

StorNext 5 works with Infiniband SRP (SCSI RDMA Protocol) attached storage for Linux and Windows 2008R2.

---

## Lattus

---

Refer to the *Lattus Release Notes* for information about compatibility between Lattus and StorNext 5.3.0.

Lattus Object Storage documentation is available here:

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

---

## Partial File Retrieval

---

StorNext Partial File Retrieval (PFR) is a separately available product which enables you to quickly retrieve and utilize segments of large media files— rather than the entire file—based on timecode parameters.

---

**Note:** StorNext Partial File Retrieval is not supported for use with Lattus or Q-Cloud.

---

Refer to the *StorNext Partial File Retrieval Compatibility Guide* at this location for information about compatibility between PFR and StorNext 5.3.0:

<http://www.quantum.com/sn5docs>

---

## 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 5 Web Services Guide* available online at

<http://www.quantum.com/sn5docs>.

---

## StorNext API (SNAPI)

---

StorNext API (SNAPI) enables you to run third-party APIs with StorNext.

---

**Note:** SNAPI documentation is applicable for releases prior to StorNext 5 release 5.0. As of StorNext 5 release 5.0, the SNAPI documentation is being maintained in the *StorNext 5 Web Services Guide*.

---

Refer to the *SNAPI Compatibility Guide* at this location for information about compatibility between SNAPI and StorNext 5.3.0:

<http://www.quantum.com/sn5docs>

---

## 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.

Refer to the *Xsan Compatibility Guide* at this location for information about compatibility between Apple Xsan and StorNext 5.3.0:

<http://www.quantum.com/sn5docs>

---

## Supported System Components

---

### Supported Browsers

---

For information on browsers supported with the StorNext GUI for this release, refer to the *StorNext 5 Compatibility Guide* available here:

<http://www.quantum.com/sn5docs>

---

## General Considerations

---

### Checksum Performance Considerations

---

**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 may 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

The **Firmware Upgrade** menu option allows you to perform a firmware upgrade on StorNext Metadata Appliances. Upgrading the firmware also upgrades the StorNext software, if applicable.



---

**Note:** Use the StorNext GUI to perform all firmware upgrades.

The **Firmware Upgrade** menu option is **only** available on StorNext Metadata Appliances.

---

For instructions on upgrading your firmware, refer to the current Release Notes for your particular appliance.

For the current Release Notes for Metadata Appliances, see:

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

For the current Release Notes for the G300 Gateway Appliance, see:

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

---

## Known Issues

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

- [StorNext File System Known Issues](#) on page 19
- [StorNext Storage Manager Known Issues](#) on page 22
- [StorNext GUI Known Issues](#) on page 26
- [StorNext Installation, Replication, HA and Other Known Issues](#) on page 28

---

**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

[Table 1](#) lists known issues specific to the StorNext File System process.

Table 1 StorNext File System  
Issues

Operating System	Change Request Number	Service Request Number	Description
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 may 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>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> <hr/> <p><b>Workaround</b></p> <hr/> <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>

Operating System	Change Request Number	Service Request Number	Description
All	57304	3561252	<p>Due to a limitation in the Linux <code>rpc.mountd</code> process, mounting a StorNext file system over NFS may return an error or hang in certain cases. To encounter the issue, a directory must be exported that is below the root of the file system. For example, if the StorNext file systems is mounted locally on the NFS server as <code>/stornext/snfs1</code>, then the following export is exposed:</p> <pre>/stornext/snfs1/myshare *(rw, sync)</pre> <p>Whereas this export entry is not:</p> <pre>/stornext/snfs1 *(rw, sync)</pre> <p>Also, the NFS server must be running Linux and have StorNext 5.x installed and the version of <code>rpc.mountd</code> on the system must correspond to version of <code>nfs-utils</code> less than 1.9.0. For example, RHEL6 is exposed but RHEL7 is not.</p> <p>Finally, the directory being exported must have an inode number that is greater than or equal to 4294967296 (<math>2^{32}</math>). This can be checked by executing the following command:</p> <pre>ls -id</pre> <p>StorNext file systems having large inode numbers will usually have one or more of the following attributes:</p> <ol style="list-style-type: none"> <li>1 A file count greater than 16 million.</li> <li>2 Multiple metadata stripe groups.</li> <li>3 Stripe groups containing mixed data and metadata.</li> </ol> <hr/> <p><b>Workaround</b></p> <p>If the problem is encountered and the NFS server cannot be upgraded to a version of Linux that is not exposed to the issue, the workaround is to mount the file system on the NFS server using the <code>compat32</code> SNFS mount option.</p>
All	60814	n/a	<p><b>Description</b></p> <p>When hosted applications are run on a server node and are active inside of an SNFS file system, StorNext restarts may hang, requiring the server node to be rebooted.</p> <hr/> <p><b>Workaround</b></p> <p>To workaround this issue, any applications running locally on the active server node operating as primary must be stopped prior to any other operations that may result in StorNext being stopped or restarted. This includes rebooting the server node, exiting HA config mode, or upgrading StorNext Appliance Firmware.</p>

Operating System	Change Request Number	Service Request Number	Description
Windows	61000	n/a	<p>If a StorNext 5 release 5.3.0 Windows client has a directory mounted file system, the shares may not function properly and using the <b>Windows Explorer/File Explorer</b> to access the file system creates a <b>FS Critical Event Notification 35, 36, 37 or 38</b> and <b>SNEX</b> exceptions in the debug directory on the Windows client.</p> <hr/> <p><b>Workaround</b></p> <hr/> <p>Quantum recommends you not directory mount the StorNext file system in a StorNext 5 release 5.3.0 Windows client. Instead, mount the StorNext file system by mapping it to a drive letter (for example, Z:).</p>

### StorNext 5 Release 5.3.0 Known Issue Where a Metadump Rebuild Causes snquota Corruption

On StorNext 5 release 5.3.0 file systems, there is a known issue (**Change Request Number 60882**) with the quota database that causes a the command **snquota -L** to loop indefinitely. The conditions under which the issue occurs are as follows:

- A managed file system exists with quotas on StorNext version 4.x.
- Upgrade the file system to StorNext 5 release 5.3.0.
- Perform a metadump restore.

Follow the steps below to workaround this issue:

- 1 Execute the command below to generate a current snapshot of all the limits in the file system. The command creates a file `/usr/cvfs/data/<fsname>/quota_regen.in`, that contains instructions on how recreate all the limits and Directory Quota (DQ) namespaces.

```
snquota -F snfs1 -X
```

- 2 Execute the command below to reset the quota database.

**WARNING:** Resetting the quota database may take a long time. It is best to have a `quota_regen.in` file just as a precaution. Perform this step with extreme caution! While the command is running, the file system is unresponsive. Allocations or deallocations are not allowed while the FSM performs an IEL scan.

```
snquota -F snfs1 -Z
```

- 3 Execute the command below to reset the limits and regenerate the DQ namespaces. This step passes the `quota_regen.in` file to the shell to run. If

Directory Quotas exist, another IEL scan is performed (and the associated file system is unresponsive).

```
sh /usr/cvfs/data/snfs1/quota_regen.in
```

---

## StorNext Storage Manager Known Issues

---

[Table 2](#) lists known issues specific to StorNext Storage Manager.

---

**Note:** Due to the formatting of the table, see [Table 2](#) on page 23.

---

Table 2 StorNext Storage Manager Known Issues

Operating System	Change Request Number	Service Request Number	Description
All	43320	1581004	<p data-bbox="618 401 1469 569">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 data-bbox="618 579 1469 611"><b>Workaround</b></p> <p data-bbox="618 632 1469 789">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> <ol data-bbox="618 800 1469 1745" style="list-style-type: none"> <li data-bbox="618 800 1469 905">1 Change the <b>FS_T10K_BLOCK_FACTOR sysparm</b> from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>:  <code>FS_T10K_BLOCK_FACTOR=32;</code></li> <li data-bbox="618 947 1469 1104"><b>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 sysparm</b> should not be set to a value that exceeds 32.</li> <li data-bbox="618 1115 1469 1262">2 Restart Storage Manager to ensure the change in <b>Step 1</b> goes into effect:  <pre data-bbox="651 1188 805 1262"># tsmstop # tsmstart</pre></li> <li data-bbox="618 1272 1469 1419">3 Verify the <b>FS_T10K_BLOCK_FACTOR sysparm</b> contains the new value:  <pre data-bbox="651 1346 1187 1419"># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre></li> <li data-bbox="618 1430 1469 1493">4 Save the current copies of your <code>/etc/fstab</code> on the MDCs and the clients.</li> <li data-bbox="618 1503 1469 1671">5 Modify <code>/etc/fstab</code> on the MDCs and the clients to use the <code>auto_dma_write_length</code> and <code>auto_dma_read_length</code> mount options as follows:  <pre data-bbox="651 1598 1455 1671">snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0 0</pre></li> <li data-bbox="618 1682 1469 1713">6 Unmount and re-mount your file systems.</li> <li data-bbox="618 1724 1469 1745">7 Use new T10K media to store a copy of the file from the disk.</li> </ol> <p data-bbox="618 1755 1469 1887"><b>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.</p>

Operating System	Change Request Number	Service Request Number	Description
All	46693	n/a	<p>Executing the command <b>snbackup -s</b> while a full or partial backup is running may result in a message that <code>/usr/adic/TSM/internal/locks/backup.1f</code> is in an invalid format.</p> <p>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> <hr/> <p><b>Workaround</b></p> <p>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>
All	47833	n/a	<p><b>Description</b></p> <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>Note:</b> This behavior may cause file data segment distribution to be sub-optimal on the destination media.</p> <hr/> <p><b>Workaround</b></p> <p>Currently, a workaround does not exist for this Known Issue.</p>
All	41413	1504258, 1635952, 3471000, 3526376	<p><b>Description</b></p> <p>In certain conditions when Storage Manager receives an end of tape indication early, SCSI sense messages may be logged by the <b>fs_fmover</b> process. These the messages will have a format similar to the following:</p> <pre>Dec 28 20:38:00 MDC-Hostname sntsm fs_fmover[29837]: E1201(8)&lt;1034815852&gt;:fsScsi1311: {2}: Check condition: op=0Ah key=00h asc=00h ascq=02h END OF PARTITION/MEDIUM DETECTED</pre> <hr/> <p><b>Workaround</b></p> <p>These messages do not cause operational problems and can be safely ignored.</p>



Operating System	Change Request Number	Service Request Number	Description
All	56261	n/a	<p>The Storage Manager fsCheckTsmFilesystemsConfig health check is a process that runs as part of the regularly scheduled health checks to help make sure that all managed file systems listed in the /usr/cvfs/config/fsmlist are mounted.</p> <p>If a managed file system appears in the /usr/cvfs/config/fsmlist but is not mounted, the health check process will issue a Service Ticket indicating something like the following:</p> <pre>Quantum SERVICE REQUEST INFORMATION SR Ticket Number: 1234 SR Serial Number: SV123436789 SR Problem Summary: QUANTUM software: Storage Manager Component fsCheckTsmFilesystemsConfig: Software fault SR Problem Code: 010C: 8 SR Error Code: Software fault SR Severity: 1 SR Notes: Internal Software Error: an unhandled software error has occurred. The following TSM managed file systems do not have storageManager enabled: /stornext/FSNAME</pre> <p>There is however an issue in the processing of the mounted Storage Manager enabled file systems that contain an "_" (underscore) in the file system name which can cause the fsCheckTsmFilesystemsConfig to skip over the counting of that file system even though it is mounted. This will then generate a Service Ticket. Systems where the file system name of a Storage Manager enabled file system matches the name of another StorNext file system following the underscore character may be exposed.</p> <p>For example, if there exists a managed file system named "snfs1" that is mounted and there is another mounted file system named "another_snfs1", the health check might incorrectly report that the file system "snfs1" is not mounted.</p> <hr/> <p><b>Workaround</b></p> <ol style="list-style-type: none"> <li>1 Verify that the file system is mounted. -or-</li> <li>2 Sometimes changing the order of the file system in the /etc/fstab can reduce the likelihood of hitting this condition by changing the order that the file system appears in /proc/mounts. StorNext may need to be restart for these changes to take affect. -or-</li> <li>3 Disable the health check by commenting out the fsCheckTsmFilesystemsConfig health check from the /usr/adic/TSM/configi/filelist file. Lines beginning with "#" are ignored. For example: \$ grep fsCheckTsmFilesystemsConfig /usr/adic/TSM/config/filelist # health_check : 0 : Config : fsCheckTsmFilesystemsConfig : 0</li> </ol>

Operating System	Change Request Number	Service Request Number	Description
Linux	45718	n/a	<p>NFS clients using certain versions of Linux are exposed to a software defect in the Linux kernel, that, in a rare race condition, may lead to an application seeing an incorrect size when accessing a file immediately after it is written.</p> <p>This problem is described by RedHat Bugs 663068 and 672981.  <a href="https://bugzilla.redhat.com/show_bug.cgi?id=663068">https://bugzilla.redhat.com/show_bug.cgi?id=663068</a>  <a href="https://bugzilla.redhat.com/show_bug.cgi?id=672981">https://bugzilla.redhat.com/show_bug.cgi?id=672981</a></p> <p>However, the Bug impacts releases of other Linux distributions as well. While the problem is unlikely to occur frequently even on systems running affected releases, customers having applications that may be impacted should contact their Linux vendors for guidance on which OS versions they should run to avoid the issue.</p>
			<p><b>Workaround</b></p> <p>Contact your Linux vendors for guidance on which OS versions you should run to avoid the issue.</p>

**StorNext GUI Known Issues**

[Table 3](#) lists known issues specific to the StorNext GUI process.

Table 3 StorNext GUI Known Issues

Operating System	Change Request Number	Service Request Number	Description
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 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 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>

Operating System	Change Request Number	Service Request Number	Description
All	57856	n/a	<p>The possibility of having different NIC cards installed in the same slots across boots results in having the same Ethernet alias names being used for the network interfaces of different NICs with different speeds (1G/10G).</p> <p>However, the Ethernet alias names depicted in the StorNext Metrics GUI page do not reflect this possible change of the network device representing the alias.</p> <hr/> <p><b>Workaround</b></p> <hr/> <p>There is currently no workaround for Change Request Number 57856.</p> <p><b>Note:</b> StorNext Metrics data is only kept for 30 days; the StorNext Metrics Report auto-corrects any port-discrepancies 30 days after the configuration is changed.</p>
All	60659	n/a	<p><b>Description</b></p> <hr/> <p>There is a known GUI issue with managing ACSLS tape libraries on the <b>Configuration &gt; Storage Destinations &gt; Libraries</b> page. If you attempt to configure a new ACSLS tape library, or edit an existing ACSLS tape library, the GUI will become unresponsive until you navigate to another page.</p> <hr/> <p><b>Workaround</b></p> <hr/> <p>To workaround this issue, contact Quantum Support for a software patch that will resolve this issue if you require ACSLS tape library support in the GUI.</p>

---

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

---

[Table 4](#) lists known issues specific to StorNext installations, data replication, HA systems and other areas.

---

Table 4 StorNext Installation,  
Replication, HA and Other  
Known Issues

Operating System	Change Request Number	Service Request Number	Description
All	47041	n/a	<p>A database index named <code>classndxatimeme</code> will be automatically added to the <code>tmdb.tier000files%</code> and <code>tmdb.tier001files%</code> tables upon starting TSM for the first time after upgrading to specific StorNext 5 releases.</p> <p><b>StorNext 5 Releases affected:</b> 5.0.1, 5.1, 5.2, and 5.3</p> <p><b>Upgrading from:</b> 4.3.2, 4.3.3, 4.7.0, 4.7.0.1, 4.7.1, 4.7.2, and the initial release of StorNext 5.</p> <hr/> <p><b>Note:</b> This does not apply to direct upgrades from StorNext 5 Release 5.0.1 to StorNext 5 Release 5.1 or later. So do not use the script below for these upgrades.</p> <hr/> <p>This index improves the performance of certain operations such as truncation policies. However, the creation of this index can take multiple hours for very large databases. TSM will be unavailable after upgrading until the indexing has completed.</p> <hr/> <p><b>Workaround</b></p> <hr/> <p>To minimize TSM downtime after upgrade, the <code>classndxatimeme</code> index can be created prior to performing the upgrade using the <code>index_tierfiles.pl</code> script available in the installation media. The script can be run while TSM is running, although it may impact the performance of other operations while the index is being added to the database.</p> <hr/> <p><b>Caution:</b> The following script should only be executed against the StorNext releases indicated in the previous section.</p> <hr/> <p>To manually add the index:</p> <ol style="list-style-type: none"> <li>1 Login to the primary MDC.</li> <li>2 Source the profile: <code>./usr/adic/.profile</code></li> <li>3 Change to the directory where <code>install.stornext</code> resides on the installation media. For example: <code>cd /tmp/stornext/stornext_full/RedHat60AS_26x86_64</code></li> <li>4 Verify that the database is up by running: <code>mysql_control start</code></li> <li>5 Execute: <code>./TSM/index_tierfiles.pl</code></li> </ol>

Operating System	Change Request Number	Service Request Number	Description
All	53933	n/a	<p>The StorNext GUI has always supported SSL for https connections while scanning namespaces. Due to SSLv3 Poodle Vulnerability, the StorNext GUI now requires to support TLSv1.2 protocol.</p> <p>In order to run StorNext 5 release 5.2 HTTPS with Lattus 3.5.1 and Lattus 3.4.4 with LIN031 Poodle WAR applied see the <b>Workaround</b>.</p> <hr/> <p><b>Workaround</b></p> <p>The application of LIN031 - Poodle WAR secures the system by disabling SSLv2 and SSLv3. This WAR can be applied to Lattus 3.4.4 and 3.5.1.</p> <p><b>Note:</b> In Lattus 3.6.X this WAR is not required.</p> <p>StorNext 5 release 5.2 functions as designed with Lattus 3.6.0. If you are using HTTPS, StorNext 5 release 5.2 requires the Poodle WAR be applied to Lattus 3.5.1 and Lattus 3.4.4. However, due to the restrictions on the cipher suites available on Lattus with this WAR, the list(scan), create namespace and buckets feature in the StorNext GUI does not work in HTTPS. It does work in HTTP.</p> <p>If you want to run in HTTPS mode, you can list, create namespace and buckets in HTTP mode. When you done, you can enable HTTPS.</p>
All	57789	n/a	<p>When running StorNext replication from a deduplication-enabled StorNext filesystem to a deduplication-enabled HA StorNext filesystem, corruption is sometimes seen in files on the HA target if HA failovers occur during the replication. Attempts to read the contents of such files report EIO (5), input/output error.</p> <p>Examination of such files using snpolicy's <b>report</b> directive show no <b>TAG</b> or <b>BLK_TAG_PRESENT</b> flags in the file inode.</p> <p>It is not known if this problem is an alternate manifestation of <b>CR 58814</b>, which may also occur in these configurations.</p> <hr/> <p><b>Workaround</b></p> <p>Avoid HA failovers/reconfiguration while StorNext replication of deduplicated files is active. See also the Release Note item for <b>CR 58814</b>.</p>

Operating System	Change Request Number	Service Request Number	Description
All	60463	n/a	In some error conditions, the suspect count for Object Storage media is incremented by 1. When the suspect count reaches 3, the media is not allocatable. In this situation the stream count for the controller is not decremented correctly. This has the effect of falsely reducing the total number of available streams. The only way to clear this condition is to reset the media suspect count back to 0 and restart Storage Manager.
			<b>Workaround</b>
			Set the suspect count sysparm parameter to a higher value (for example, MEDIA_SUSPECT_THRESHOLD=0x7fff;) in the /usr/adic/TSM/config/fs_sysparm_override file. Also refer to the MEDIA_SUSPECT_THRESHOLD in the /usr/adic/TSM/config/fs_sysparm.README file.
All	58814	n/a	<b>Description</b>
			When running replication from a deduplicated file system to another deduplicated file system, corruption is sometimes observed in large files. Examination of the replicated files using snpolicy's report directive shows that there are holes in the replicated file's extended representation blobmap. These areas read as zeroes rather than the appropriate data.
			<b>Workaround</b>
			Avoid replicating deduplicated files, especially when deduplication is in progress. The problem is more likely to be seen on very large files (tens or hundreds of gigabytes or larger). It is also more likely to be seen when replication is run on a very short schedule (minutes). It is believed to occur mostly or only when deduplication and replication are running concurrently.

### StorNext 5 Release 5.3.0-Specific Pre- Upgrade Requirement for StorNext Connect

**Note:** Only apply this procedure if your M440, M660, Artico, or Pro Foundation is managed by StorNext Connect.

You should update the Connector on all StorNext Connect-managed systems that you plan to upgrade to StorNext 5 Release 5.3.0. This includes the StorNext Connect system (Node 1). This upgrade MUST be done BEFORE upgrading the appliance firmware in order to avoid a situation where the StorNext Connect Volume Storage widget on the Dashboard does not display data for StorNext Connect monitored systems.

## How to Install the StorNext Connect Connector

---

**Note:** If you did not install the new StorNext Connect Connector prior to applying the StorNext 5 Release 5.3.0 Firmware Upgrade, see [Repair a StorNext Connect System After Firmware Upgrades](#) on page 33.

---

Here are the steps you will need to download and install a new StorNext Connect Connector:

- 1 Download the StorNext Connect Connector. Do one of the following, depending on your StorNext Connect system's connection to the Internet:
  - [To Download the Connector on a StorNext Connect System With an Internet Connection](#)OR
  - [To Download the Connector on a StorNext Connect System Without an Internet Connection](#)
- 2 Install the new StorNext Connect Connector. See [Install the new StorNext Connect Connector](#) on page 33.

### To Download the Connector on a StorNext Connect System With an Internet Connection

To download the new Connector from the StorNext Connect App Store when the StorNext Connect system is connected to the Internet:

- 1 Log into StorNext Connect as an administrator.
- 2 Go to App Store.
- 3 Click **Administer** on the **Discover Components** tile.
- 4 Click the **Download** button next to the "Available Connector."
- 5 Continue to [Install the new StorNext Connect Connector](#).

### To Download the Connector on a StorNext Connect System Without an Internet Connection

To download the StorNext Connect Connector when the StorNext Connect system is not connected to the Internet:

- 1 Log into the StorNext Connect web site at:  
[www.stornextconnect.quantum.com](http://www.stornextconnect.quantum.com)
- 2 Click the **Updates** button for the system running StorNext Connect.
- 3 Under **Offline items available for download**, click **Discover content** and download the tar.gz file.
- 4 Log into the StorNext Connect system as an administrator.
- 5 Go to the App Store.
- 6 In the Connect Apps tile, click **Install File**.
- 7 In the Connect Apps screen, click **Choose file**.
- 8 Select the tar.gz file you downloaded earlier and click **Open** to unpackage the file.



- 9 Continue to [Install the new StorNext Connect Connector](#).

### Install the new StorNext Connect Connector

- 1 Click **Upload and install file**.
- 2 Install the new StorNext Connect Connector on the appliance being upgraded:
  - a Click the StorNext Connect logo in the upper left corner to go to the application tiles.
  - b Click **Discover Components**.
  - c Click the "Add or update Connectors" link.
  - d Click the "update" link under the **Connector action** column for the server node you are updating. Follow the on-screen instructions to complete the update.

### Repair a StorNext Connect System After Firmware Upgrades

If you do not install the Connector before you applied the M440, M660, Artico, or Pro Foundation StorNext 5 Release 5.3.0 firmware upgrade, Node 1 of your appliance pair will no longer be managed by StorNext Connect. If this happens, you must take the following steps on Node 1:

- 1 Open an SSH connection to the MDC node using either IP address **10.17.21.1** (Node 1) or **10.17.21.2** (Node2) on the Metadata network.

---

**Note:** Use the IP addresses assigned if different from the defaults used here.

---

- 2 Enter **stornext** for the username at the prompt.
- 3 Enter the password for the **stornext** user account. The default password is "password", but may have been changed after initial configuration.
- 4 At the command prompt enter the following to gain root user access:

```
sudo rootsh
```

- 5 Enter the password for the **stornext** user account again when prompted.
- 6 Press Enter.
- 7 Determine the IP of the StorNext Connect system and ensure it can be reached by Node 1. If you are unsure, you can log onto Node 2 and check the `/opt/quantum/connector/etc/connecthostname` file on that system.
- 8 Edit `/opt/quantum/connector/etc/connecthostname` and change the IP address (10.1.1.2) to the IP of your StorNext Connect system.
- 9 Adjust the mintd configuration by running the following command, substituting your StorNext Connect IP for W.X.Y.Z.

```
/opt/quantum/mintd/mintd_control.py set --minthost W.X.Y.Z
```

---

## Contacting Quantum

More information about StorNext is available on the Quantum Service and Support website at <http://www.quantum.com/ServiceandSupport>. The Quantum Service and Support website contains a collection of information, including answers to frequently asked questions (FAQs).

---

### StorNext Upgrades

---

To request a StorNext software upgrade for non-Quantum MDCs , visit <http://www.quantum.com/ServiceandSupport/Upgrade/Index.aspx>. To request a StorNext software upgrade for StorNext Appliances, open a support ticket at: <https://onlineservice.quantum.com/>. For further assistance, or if training is desired, contact the Quantum Technical Assistance Center.

---

### Contacts

---

Quantum company contacts are listed below.

#### Quantum Home Page

Visit the Quantum home page at:

<http://www.quantum.com>

---

### Comments

---

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

[doc-comments@quantum.com](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. Benefit today at:

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

- **eSupport** - Submit online service requests, update contact information, add attachments, and receive status updates via email. Online Service accounts are free from Quantum. That account can also be used to access Quantum's Knowledge Base, a comprehensive repository of product support information. Sign up today at:

<https://onlineservice.quantum.com/>

**Quantum.**  
**Global Services**

For further assistance, or if training is desired, contact the Quantum Customer Support Center:

<b>United States</b>	1-800-284-5101 (toll free) +1-720-249-5700
<b>EMEA</b>	+800-7826-8888 (toll free) +49-6131-3241-1164
<b>APAC</b>	+800-7826-8887 (toll free) +603-7953-3010

For worldwide support:

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

---

## **Worldwide End-User Product Warranty**

---

For more information on the Quantum Worldwide End-User Standard Limited Product Warranty:

<http://www.quantum.com/serviceandsupport/warrantyinformation/index.aspx>

