StorNext 5 release 5.2.2
Release Notes

Product	StorNext® 5 release 5.2.2
Date	August 2015

Contents

What's New in StorNext 5 release 5.2.2

Purpose of this Release

StorNext 5 release 5.2.2 is a maintenance release of StorNext that provides expanded support for new clients, including support for the enhanced distributed LAN client that is included with Apple OS X version 10.11, and also addresses several customer reported issues listed in the section Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.2.2 on page 3 in this document.

New Features and Enhancements in StorNext 5 release 5.2.2

Support for Additional Mac Operating System Versions

Apple has enhanced Xsan 4.1 that is included with the El Capitan (OS X 10.11) release, anticipated to be generally available during the third quarter of the calendar year 2015, to include Distributed LAN Client (DLC) support. Quantum has tested the OSX 10.11 Xsan 4.1 client to ensure that it works correctly when attached to a StorNext MDC platforms in both SAN and DLC mode.

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

Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.2.2

The following table lists the customer reported issues and enhancements addressed in StorNext 5 release 5.2.2.

Operating System	Change Request Number	Service Request Number	Description
All	32811	1213826	**fsrecover** hits infinite loop recovering one file into a non-existent relpt.
All	41038	1505742, 1570378	Scheduled name space realization stopped happening.
All	46051	1638584	**fsmedinfo -l** do not show the segment number or offset for segment.
All	48708	3393794	Unclear rename message.
All	53923	3452842	Unsigned integers display as negative values in logs.
All	54629	3508876	StorNext GUI allows for use of preceding spaces when creating an SDisk, but will not list new SDisk afterwards.
All	54751	3509432	The WARNING message "ras file already exists" is not warning us of anything - no action required.

Operating System	Change Request Number	Service Request Number	Description
All	55083	3499380	The RAS log message for I/O error is sent when error is VOP_ESTALE but no slog.
All	55479	3519756, 3502040	Stripe Group mark offline on reconnect cannot recover until remount.
All	55907	3531752	GUI **Recover Files** does not list FS if name and mount point do not match.
All	56059	3534174	When using Linux NFS servers, write performance is abysmal if file is in shared write mode.
All	56188	3521282	Timeout in leaving config mode put HA in a bad state.
All	56220	3540552	`cvfsck`: metadata preload violates repair ordering/layering.
All	56228	3462410	`fsrmdiskcopy`: BUG: unable to handle kernel NULL pointer dereference.
All	56229	3540552	Errors in conversion can result in undetected incomplete metadump.
Linux	54620	3507308	Should enable **windowsSecurity** if **enforceAcls** is set to true.
Linux	56279	3539894	G300 NAS server samba threads deadlocked in truncate requests for a file.
Linux	56662	3546966	`FixStaleAttrs()` update FSM with stale data causing file corruption.

What's New in StorNext 5 release 5.2.1

Purpose of this Release

StorNext 5 release 5.2.1 is a maintenance release of StorNext that provides additional support for Red Hat Enterprise Linux, SUSE Linux Enterprise Server and Ubuntu operating systems, adds support for StorNext NAS, and also addresses several customer reported issues listed in the section Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.2.1 on page 5 in this document.

New Features and Enhancements in StorNext 5 release 5.2.1

StorNext NAS Support

StorNext 5 release 5.2.1 adds support for SMB and NFS NAS share access through network connections and the optional NAS feature when installed and configured on a StorNext appliance.

For additional information on how to install and configure the NAS feature, contact Quantum Customer Support (see Contacting Quantum on page 56).

Support for Additional Operating System Versions

StorNext 5 release 5.2.1 adds support for additional operating system versions of Red Hat Enterprise Linux, SUSE Linux Enterprise Server, and Ubuntu. With this release, StorNext gains compatibility with:

Operating System	Version
Red Hat Enterprise Linux	Version 6 Update 6
	Version 7 (client only) Version 7 Update 1 (client only)
SUSE Linux Enterprise Server	Version 12 (client only)
Ubuntu	Version 14.04.2

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

Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.2.1

The following table lists the customer reported issues and enhancements addressed in StorNext 5 release 5.2.1.

Operating System	Change Request Number	Service Request Number	Description
All	40301	1486782, 3477454	`cvlabel` output "The disk is too big to label. You will lose some blocks"
All	42464	1554048	`fsschedule` command fails when modifying an `activevault` schedule
All	45118	1628792	StorNext 5 to set default of MySQL `innodb_buffer_pool_size` to 40 GB
All	46817	3300044	Need documentation that `fsretrieve -n` will only run on local MDC
All	52223	3468068	`Cvupdatefs` failure doesn't provide sufficient error reporting in trace log
All	52678	3460814	`fs_feature` core dumps due to uninitialized g_display_fn when running `activevault`
All	52922	3479708	NFSD hung waiting for cvfs IO buffer.
All	53035	n/a	Add additional device inquiry strings to the TGV license validation code

Operating System	Change Request Number	Service Request Number	Description
All	53776	3518924, 3526170	GUI: retrieve of file to a new filename retrieves only 1-byte unless user knows filesize
All	53810	3492202	A bogus candidate entry can cause fspolicy to put files and directories into a bad state.
All	53890	3494282	`cvfsck` should stop if I/O errors happen - StorNext 5
All	54153	3443236, 3495936	Need recommended/required OS packages prior to StorNext install
All	54168	3511046, 3512842	sar logs need to be changed to a 24 hour time format
All	54599	3502086	webservices becomes unresponsive/crashes with no warning or ras tickets being logged
All	54600	3502738, 3512842, 3516022	When DDM is configured with aliases in /etc/hosts, the DDM hostnames cannot be resolved
All	55293	3516508 3518794	VopGetattr() should not panic FSM when called with bogus inode
All	55598	3522076	Metadump restore doesn't support named stream inodes correctly
All	55716	3525532	Need smb.conf "printcap name" default be /dev/null to disable CUPS warnings (dxi:42551)
All	55769	3526302	`fs_feature` may pass the wrong date to the feature it is running
Linux	49268	3409512, 3515864	ipython activation may cause kernel stack overrun on Linux
Linux	52536	n/a	SUSE Linux Enterprise Server 12 SNFS client support
Linux	54219	3488100	Enhancement Request: add screen rpm utility to enable multiple login sessions
Linux	54568	3508950	We do not resinstall local_repo on install retry after failure
Linux	54835	3510896	Changes required for supporting SNFS clients on Ubuntu 14.04.02 LTS
Other	55104	3515176	Client Software for UBUNTU not avail from the drop-down menu on the MDC440 in 5.2.x
Windows	55290	3518186	5.2.0.1 SCM builds fail to install on windows
Windows	55783	3525140	Need new nss-pam-ldapd rpm to address 32-bit uidNumber problems

What's New in StorNext 5 release 5.2.0.1

Purpose of this Release   StorNext 5 release 5.2.0.1 is a minor update to StorNext 5 which addresses Change Request Number 54642 (see the details in the following table).

Operating System	Change Request Number	Service Request Number	Description
All	54642	n/a	Corrected an issue that could lead to an I/O hang during ingest if a second file system is mounted.

Support for Ubuntu Operating System

StorNext 5 release 5.2.0.1 supports the following Ubuntu releases:

- Version 14.04.0
- Version 14.04.1

For additional information, see the *StorNext 5 Compatibility Guide*.

What's New in StorNext 5 release 5.2

Purpose of this Release   StorNext 5 release 5.2 for StorNext Metadata Appliance customers adds support for Q-Cloud Archive™, a new cloud-based tiering option for StorNext. This release expands support for Mac OS X by interoperating with Xsan 4 (part of OSX 10.10 "Yosemite") as well as Xsan 2.2 and Xsan 2.3. Release 5.2 also adds StorNext Certified Disks and the necessary licensing. Finally, release 5.2 addresses several customer reported issues listed in the section Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.2 on page 9 in this document.

New Features and Enhancements in StorNext 5 release 5.2   ### Q-Cloud Archive™

The Storage Manager has been enhanced to support the use of public cloud as storage tier though a service called Q-Cloud Archive™. This service allows StorNext users to create policies to copy or move files to a private account that is managed and serviced by Quantum. Q-Cloud Archive™ has the following benefits:

- Very easy to provision and use
- No upfront fees

- No additional software, cloud gateway hardware, code development, or manual transfers required

- Built on globally-recognized public cloud services for reliability and data integrity

- True OPEX storage; pay only for usage

- Competitively priced and a simplified billing model

For additional details, see the *StorNext 5 User's Guide*, section **Q-Cloud** or talk to your Quantum sales contact.

Support for Additional Mac Operating System Versions

StorNext 5 release 5.2 includes support for additional versions of Xsan and Mac OS. With this release, StorNext gains compatibility with:

- Xsan 4 / Mac OS X 10.10 ("Yosemite")

- Xsan 2.2 / Mac OS X 10.5.8+ / 10.6

- Xsan 2.3 / Mac OS X 10.7 ("Lion")

This additional interoperability will allow users to get the benefits of StorNext 5 release 5.2 without having to upgrade their Macintosh environment if they choose not to. For details, refer to the section **Mounting the StorNext File System on Xsan 4 and Later Clients** in the *StorNext 5 Installation Guide*.

Also refer to the *Xsan Compatibility Guide* for additional information about compatibility between Apple Xsan and StorNext.

Support for StorNext Certified Disk

Starting with StorNext 5 release 5.2, Quantum will begin using disk licensing to support a new model where 3rd-party disk vendors can market and sell their systems as "StorNext Certified." Please note that all storage systems branded as Quantum or sold by Quantum as a resold product is by definition, certified.

Customers seeking to add storage to existing StorNext systems or purchase storage for new systems will benefit from Quantum support professionals that understand their storage environment and can better address issues that may arise from 3rd-party storage systems.

Newly deployed systems that are based on StorNext Certified storage systems include an additional benefit of reduced service costs for the entire environment.

To enable this, StorNext 5 release 5.2 users will be asked to procure a StorNext disk license for $0 (zero dollars). This license will aid in the maintaining a profile of the disks installed in each StorNext deployment. For additional information on disk licensing, see **StorNext Disk Licensing** and **StorNext Disk Certification Catalog (DCC)** in the *StorNext 5 and StorNext FX 5 Licensing Guide*.

StorNext Lattus Changes

This release addresses the SSL "Poodle" security vulnerability by restricting SSL connections to the TLS 1.0 or higher cryptographic protocol. Lattus customers should apply Lattus patch LIN031c and update to SSL 1.0.1c. The default

Stornext SSL certificate and path have changed for Lattus customers. See the *StorNext 5 User's Guide* section **HTTPS Default CA ROOT Certificate File or Path**.

Starting with StorNext 5 release 5.2, StorNext provides the capability to convert an AXR namepsace to an S3 bucket. To convert one or more AXR namespaces to S3 buckets, see **Converting an AXR Namespace to an S3 Bucket** in the *StorNext 5 User's Guide*.

Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.2

The following table lists the customer reported issues and enhancements addressed in StorNext 5 release 5.2.

Operating System	Change Request Number	Service Request Number	Description
All	50361	1292936, 1203510, 1244184, 1252666, 1256238, 1264130, 1283186, 1284344, 1284436, 1329830, 1478788, 3436562, 3436750, 3424342	Save disk labels for possible recovery if the need to restore the labels occurs
All	50362	3357910, 3432000	UNKNOWN VPINST: UNKNOWN EVENT: 85 TEXT: Failed to roll event file /usr/adic/TSM/internal/event_dir/ 0x0004ea0ad4bd3bd9.create_d, errno: 2 fs: /stornext/ cvfs1
All	50364	3400798, 3414562	StorNext 5.X qustat hourly reports are no longer reset each hour
All	50377	n/a	GUI Hardware status Node 2 serial number is blank
All	50451	1245606	Customer did a delete media from GUI and then canceled the eject from GUI.

Operating System	Change Request Number	Service Request Number	Description
All	50452	1515914, 1525624, 1557926, 1629758, 1635034, 1644314, 1271440, 1312338, 1327440, 1330014, 1329864, 1338932, 1371132, 1372660, 1418712, 1426826	GUI LOI: canceling media move after a media has been ejected causes problems
All	50453	1231270, 1309642, 1305202	Need method to detect and associate orphaned media with an archive
All	50454	1176108	Cannot use **chgmedstate** to put a intransit/unknown media into an archive if no other media of that type already exists in the archive
All	50455	1215318, 1526752	The **chgmedstate** utility should not allow a media to be associated with an archive (library or vault) if there is a pending enter/eject for the media
All	50536	3442896	**fsobjcfg -r -f** invalid usage in StorNext 5 release 5.1
All	50547	1204448	Improve syslog documentation in `log_params` file
All	50548	1435330	TSM uses maximum logging if `log_params` file is missing and does not issue a warning
All	50851	3389468	Add warning on ssh login when an upgrade is still in progress on the nodes
All	50852	3389416, 3380000	**hwclock** does not get set when system get smith
All	50853	3363752	Blacklisting dm-multipath does not work in initramfs
All	50981	n/a	Request to have 10GbE Ethernet interfaces available in a M441D StorNext appliance
All	50982	3423152	Support HP Advanced Path Failover for StorNext AEL Archives & Scalar libraries, ETA calendar Q4 2014

Operating System	Change Request Number	Service Request Number	Description
All	51172	3371316, 3380000, 3415176, 3428360, 3436244	**fsmpm** should be less verbose when setting `deviceparams`
All	51487	3438412	**cvfsck** dies with a segfault/core
All	51905	3458612, 3392684, 3467430	FixStaleAttrs() return VOP_EINTR causing inode to be mark stale
All	52035	3464048	Appliances need symbolic link from /opt/graphite/storage/whisper to /scratch/whisper
All	52156	n/a	**fsmedinfo** XML/JSON output can be malformed when DB errors occur
All	52238	n/a	Replication failed due to failing to open the dump file on the src
All	52242	3432566	Name space realization failed due to cleanup_tree failure on target
All	52250	n/a	Deadlock in sharing connection for metadata sending and response receiving
All	52260	3442480	G300: system hanged while trying to reboot at the end of the upgrade (SN 5.1 => SN 5.2)
All	52341	3432566	Add more log info for replication failure due to the failure of file open by handle
All	52492	3457398	FSM PANIC In_rele: ip(0xf67af0 : 0x13800000824180) i_ref_count(-1) has gone negative!
All	52753	3453464	ntp.conf too restrictive on Appliance causing server to be ignored.
All	52878	3476712	The allocator can pick a slice that is beyond the # of slices of a SG resulting in fsm panic
All	52882	3477190	M440 with 10GbE NIC on slot3: fresh install network config failed
All	52927	3467430, 3477406	Journal replay failed at Keeper after a SMITH
Linux	50371	3429870	VopLookup() deadlock in the LOCK_ACTIVE_CLIENT_WR() call.
Linux	50534	3399452, 3440750	5.0.1: ASSERT(IP_IS_SPACE_TREE(ip) ||this.idiext_frblock + extsize <= next.idiext_frblock)

Operating System	Change Request Number	Service Request Number	Description
Linux	50620	3319200, 3429870, 3437466, 3438678	**ls -l** is showing ? instead of file attributes in case of FSM services being cycled
Linux	50849	3371316	Setting the elevator to deadline on the kernel line increase the boot speed
Linux	50878	3437876	FSM panic when VopFullPathReverse() is given invalid inode number.
Linux	51412	3442480, 3449424, 3456412, 3451726	Losing a single path to a LUN can send the fsmpm in a rescan loop
Linux	51486	3420306	**sncompare** generates a script which may delete filecomp db records that should be kept
Linux	52240	3443766	LTFS drives exceed their failure threshold and are placed offline if invalid characters are used in filenames
Linux	52469	3470544	StorNext upgrade gateway client script did not fully clean up stale upgrade remnants from a previous upgrade causing the upgrade to fail
Linux	52768	1319898, 3469594	**fsretrieve -B** option does not work with copy2-copy4 media
Linux	52848	3453464	Upgrade: ntp.conf too restrictive on Appliance causing server to be ignored.
Mac OS	50891	3437384	RADAR 18156105 :: CVFS ASSERTION FAILED: f_rwlck->rw_state & RW_WRITE line 1861 file lock.c @nomad/md_debuglog.c:358
Windows	50539	3303762	Win: CliCfg: Remove the Hard Mount option
Windows	50621	1577606, 3400798	There is a race between InitSockHandlers/SockInput() starting and KillSockHandlers()
Windows	53333	3439086, 3380044, 3441922, 3459502, 3475416	Windows Server 2012 Crashes when CIFS sharing a StorNext volume

What's New in StorNext 5 release 5.1.1

Purpose of this Release

StorNext 5 release 5.1.1 is a maintenance release of StorNext that provides support for StorNext Connect, and also addresses several customer reported issues listed in the section Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.1.1 on page 13 in this document.

Note: If you have StorNext 5A release 5.1.1, StorNext 5E release 5.1.1, or StorNext 5G release 5.1.1 installed and are upgrading to the same build number but a different edition (StorNext 5A, StorNext 5E, or StorNext 5G), the installation will fail. An example of the error displayed by the `install.stornext` command is shown below:

```
Product StorNext version 5.1.1 (47518) is already
installed.
```

To workaround this issue, use the **-force** option with the `install.stornext` command, for example:

```
install.stornext -force
```

Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.1.1

The following table lists the customer reported issues and enhancements addressed in StorNext 5 release 5.1.1.

Operating System	Change Request Number	Service Request Number	Description
All	34555	1292936, 1203510, 1244184, 1252666, 1256238, 1264130, 1283186, 1284344, 1284436, 1329830, 1478788, 3436562, 3436750, 3424342	Save disk labels for possible recovery if the need to restore the labels occurs.
All	45423	3434504, 3434562	Files with InodeFlagPerfectFit set have the potential to lose blocks when freeing inode

Operating System	Change Request Number	Service Request Number	Description
All	45693	3357910, 3432000	UNKNOWN VPINST: UNKNOWN EVENT: 85 TEXT: Failed to roll event file /usr/adic/TSM/internal/event_dir/ 0x0004ea0ad4bd3bd9.create_d, errno: 2 fs: /stornext/ cvfs1
All	47973	3422014, 3443890	The /etc/init.d/init_wait script should skip down stripe group.
All	49068	3413358	pse_snapshot too large because it includes multiple compressed copies of all qustats
All	49153	3409452, 3436972	Restarting MSM causes the tape drives to be disassociated from their DAS library
All	49416	n/a	Hardware status -> Storage arrays should not use terms "node 1" and "node 2"
All	49755	3429246, 3430904	Large ntsecurity descriptors can cause FSM crash
All	49764	3430876	RPL_upgrade broken in StorNext 5
All	50195	n/a	GUI Hardware status Node 2 serial number is blank
All	50489	3442896	fsobjcfg -r -f invalid usage in SN 5.1
Linux	48916	3399452, 3440750, 3441078	5.0.1: ASSERT(IP_IS_SPACE_TREE(ip) ||this.idiext_frblock + extsize <= next.idiext_frblock)
Linux	49940	3319200, 3429870, 3437466, 3438678	ls -l is showing "?" instead of file attributes in case of FSM services being cycled"
Linux	50142	3429870	VopLookup() deadlock in the LOCK_ACTIVE_CLIENT_WR() call.
Windows	48208	3357938	Windows: StorNext 5 SMBv1 share does not display all files in a directory
Windows	49505	3397428	Windows Explorer will fail to open the root of a StorNext file system
Windows	49968	1577606, 3400798	There is a race between InitSockHandlers/SockInput() starting and KillSockHandlers()

What's New in StorNext 5 release 5.1

Purpose of this Release

StorNext 5 release 5.1 for StorNext Metadata Appliance customers adds support for moving data from tape to Lattus Object Storage with the **fsmedcopy** command, Lattus Object Storage to use the S3 protocol, and introduces a new Web Service for StorNext. This release also addresses several customer reported issues listed in the section Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.1 on page 18 in this document.

> **Note:** If you have StorNext 5A release 5.1, StorNext 5E release 5.1, or StorNext 5G release 5.1 installed and are upgrading to the same build number but a different edition (StorNext 5A, StorNext 5E, or StorNext 5G), the installation will fail. An example of the error displayed by the **install.stornext** command is shown below:
>
> ```
> Product StorNext version 5.1 (46950) is already
> installed.
> ```
>
> To workaround this issue, use the **-force** option with the **install.stornext** command, for example:
>
> **install.stornext -force**

New Features and Enhancements in StorNext 5 release 5.1

Move Data from Tape to Lattus

The Storage Manager CLI command, **fsmedcopy**, has been enhanced to support the following new source to destination media combinations to move data to Lattus Object Storage:

- Tape to Object Storage
- Object Storage to Object Storage
- Storage Disk to Object Storage

For additional details, see the *StorNext MAN Pages Reference Guide*.

Lattus Object Storage Enhancements

This release adds Simple Storage Service (S3) protocol including Multipart Upload support for Lattus Object Storage.

The list below provides a synopsis of the Wide Area Storage label change, and the new support for Lattus use of the S3 protocol.

- All system parameters prefixed with **FS_WASTORAGE_** have been changed to the **FS_OBJSTORAGE_** prefix. All system parameters with the old **FS_WASTORAGE_** prefix in your existing **fs_sysparm** and **fs_sysparm_override** files will automatically be converted to the new **FS_OBJSTORAGE_** prefix.

- The **snpolicy was.conf** file has been renamed to **objs.conf**, which includes new parameter values to support both AXR and S3 protocols. Your existing **snpolicy was.conf** file will automatically be converted to the new **objs.conf** file.

- All **snpolicy** option values and attributes prefixed with **was_** have been changed to the **objs_** prefix. Any parameters prefixed with **was_** in your existing configuration files **was.conf** and **snpolicyd.conf** will automatically be renamed with the **objs_** prefix.

- The command **fswascfg** has been deprecated by command **fsobjcfg**, which provides new configuration options and values to support both AXR and S3 protocols.

- This release adds new options to the commands **fsmedread** and **fsmedwrite** to support both AXR and S3 protocols.

- This release adds a new Lattus Object Storage media type labeled **S3**, which describes a Lattus S3 namespace (or bucket), while the existing media type **LATTUS** describes an AXR namespace. The following commands now support both LATTUS and S3 media types:

  - **fsaddclass**

  - **fsfilecopy**

  - **fsmedcopy**

  - **fsmodclass**

  - **fsstore**

- The commands **snbackup**, **snrestore**, and **snbkpreport** now support both AXR and S3 protocols.

- Upon upgrading to StorNext 5 release 5.1, all existing SNSM Lattus Object Storage configuration will be automatically updated to support both AXR and S3 protocols. For additional details, refer to the **Upgrading of SNSM Lattus Object Storage Configuration to StorNext 5 release 5.1** section in the *StorNext and StorNext FX 5 Upgrade Guide*.

Object Storage License Enhancements

Starting with StorNext 5 release 5.1, a new **Object Storage** (or object_storage) license is required for customers who wish to access Lattus Object Storage using StorNext Storage Manager. The new license is capacity based and only applies to the capacity used with Lattus Object Storage.

Note: There is no license change on the A10 snpolicyd systems which will continue to use the snfs_was license.

The old **SNSM Wide Area Storage** (or snsm_was) license has been deprecated in this release of StorNext. If you have an existing snsm_was license, you will need to update your StorNext license file to include the new object_storage license, in order to continue using StorNext Object Storage features in StorNext 5 release 5.1. For additional details, refer to section **Upgrading from snsm_was License to object_storage License** in the *StorNext and StorNext FX 5 Upgrade Guide*.

Web Services version 2

StorNext 5 release 5.1 introduces StorNext Web Services version 2.

StorNext Web Services version 2, enhance web services with better command coverage and support for XML, JSON, and TEXT responses. The web services are categorized according to StorNext functional areas. For example, directories, files, media, etc.

Version 2 adds a user interface to turn the web service on or off, set the protocol (HTTP or HTTPS), and control authentication. By default, version 2 web services are turned off. Use the StorNext GUI to make changes (refer to **Web Services (V2)** in the *StorNext 5 User's Guide*).

The table below provides a description for each category of a web service. The web services are organized into the following high-level categories.

Category	Description
Archive	Return information about an archive, query an archive port, or change the state of an archive.
Directory	Modify the class attributes of a directory or retrieve or recover files from media.
Drive	Report or change the state of drive components and storage subsystems.
File	Report, retrieve, and store files to tiered storage.
Media	Manage media – copy, clean up, move, and report.
Object Storage	Report Object Storage components.
Policy	Manage and report policies.
Quota	Manage and report quotas.
Report	Return information about subsystem resource requests.
Schedule	Manage and report schedules.
System	Get the status of system and Tertiary Storage Manager (TSM) components. Manage and report backups.

For more information, see the *StorNext Web Services Guide*.

Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.1

The following table lists the customer reported issues and enhancements addressed in StorNext 5 release 5.1.

Operating System	Change Request Number	Service Request Number	Description
All	35413	1484626	**fsname** in **cfgx** file not validated with **sb fsname**.
All	37948	1410974, 1461558, 1572000, 3357910, 3401348	Excessive RAS messages when manager license capacity exceeded.
All	38028	1412936, 3301594	Directories associated with a policy display incorrectly in the GUI if renamed on the FS.
All	40237	1480756	**cleanup** still leaves files under a cleaned up sub-directory.
All	41494	1524540, 1523344, 1551214, 3303296	**OpHangLimitSecs** exceeded when running quota reports and **quotas=true** on HA Shared FS.
All	42459	1554048, 1591144, 1625316	**fsactivevault** does not accept vault name containing a "**-**".
All	43267	1564556, 1633820	**showc** performance is very poor with large number of store candidates.
All	44346	1601286	Corrections needed for **cvcp** man page.
All	44887	1610030	The Storage Manager CLI command, **fsmedcopy**, has been enhanced to support the following new source to destination media combinations: • Tape to Object Storage • Object Storage to Object Storage • Storage Disk to Object Storage For additional details, see the *StorNext MAN Pages Reference Guide*.
All	44964	1327034	The drive type is not being printed out in message.
All	45338	1646652	**fsformat** with **-c** and **-q** flags not working as documented in the man page.
All	46448	3303748	**fsmedcopy** fails when multiple media specified and **FS_MAX_ACTIVE_TAPECOPIES > 1**.
All	46472	1653036	**cvcp** man page mentions nonexistent command options.

Operating System	Change Request Number	Service Request Number	Description
All	46499	1648070, 3395062	File attributes may not match metadata if store policy and truncation policy manipulate file at same time.
All	46505	3405854	Return ENOENT to the client instead of calling PANIC when given an invalid inode number.
All	46566	3413358, 3431998, 3438210	PANIC: /usr/cvfs/bin/fsm "FSM MetaData READ: ERROR stripe/0x0 length 4096 block/0x5ffa - Detected bad metadata contents on disk!"
All	46740	3307950	**fsmedinfo** output displays negative mount count value.
All	46938	3356884	64M or larger DMA I/O fails with ENOMEM on RHEL5 on stripe groups having only one LUN.
All	46998	3361560	GUI performs automatic eject from vault and LOI displays blank **Current Library**.
All	47154	3399694, 3401610	Truncation (**Inode_punch_hole**) does not limit the transaction size.
All	47312	3371126	**fs_cpyresp** dies and caused TSM to terminate due to a tape failure which lead to a database error.
All	47362	3374136	Better error handling in closing socket in FSM (REGRESSION in 5.0)
All	47421	3372300	**cvfsck** creates duplicate RPL attributes for unconverted inodes.
All	47502	3372300	FSM workaround needed for 47421.
All	47510	1546270, 1628796, 1627156, 3363406	GUI needs to expose method for increasing the allowed number of slots within a library.
All	47532	3374136	**snhamgr** status errors can cause problems for syncha and **DSM_control**.
All	47685	1612198, 3361724	Eliminate false alarm admin tickets for ports that are not connected.
All	47694	1610030	GUI Support for 44887 - need method to copy (**fsmedcopy**) between all media types.
All	47980	n/a	r**andom_in_range()** causes **segfault**.
All	48004	3384334	Unable to upgrade M-Series to 5.1.0 from 4.7.0.1 due to failure to enter and exit config mode reliably to set the **ha_smith_interval** file on both nodes.
All	48135	3393136	GUI should allow upgrades to FW released before maintenance license expiration date.

Operating System	Change Request Number	Service Request Number	Description
All	48340	3398958, 3342200, 3402564	Optimize boot speed by reducing log level.
All	48362	3397844, 3406632, 3393506, 3399000, 3413356	Invalid Admin Alert, about no space on disk, when tape becomes full during **fsmedcopy** or **fsstore**
All	48363	3399694	StorNext journal can corrupt itself leaving the file system unusable and resulting in metadata corruption.
All	48426	3403012	GUI: space usage data is not consistent between file system and SN metric report pages.
All	48496	3402564	A few drivers need to be blacklisted to avoid library problems.
All	48577	3408948	kernel heap corruption in **LookupCvp()**.
All	48682	3404598	FSM panic with **Segmentation Fault Signal**.
All	48696	3396622	**snstatd** leaks memory
All	48756	3399002	**fpit(0)**: Cannot deallocate space of inode [blah] from Pending Free List! -Invalid argument.
Linux	39026	1445834, 1444982, 1554654, 1605450, 3399930	**cvupdatefs** allows label and config file capacity mismatch.
Linux	40800	3359158	GUI is not updating /usr/adic/TSM/config/ filesize.config correctly.
Linux	47965	3378182	Files corrupted on StorNext Client/NFS Server with multiple open and writes to same file from NFS clients.
Linux	48507	3401610, 3409512	Kernel stack overrun on CVFS doing mmaped I/O
Linux	48937	3409934	RamSan multipath.conf entry needs different path_checker value.

What's New in StorNext 5 release 5.0.1

Purpose of this Release

StorNext 5 release 5.0.1 is a maintenance release for StorNext Metadata Appliances.

StorNext 5 release 5.0.1 adds support for several new platforms, enhances the StorNext Web Services API, and the StorNext command line.

Finally, this release addresses several customer reported issues listed in the section Customer Reported Issues and Enhancements Addressed in StorNext 5 Release 5.0.1 on page 23 in this document.

Note: StorNext 4.7.1 **cannot** be upgraded to StorNext 5 Release 5.0. This upgrade **is not** supported.

StorNext 4.7.1 **can** be upgraded to StorNext 5 Release 5.0.1. This upgrade **is** supported.

New Features and Enhancements in StorNext 5 release 5.0.1

New Platforms

This release adds support for:

- Windows 8.1
- Windows 2012 R2
- RHEL 6.5
- RHEL 5.10
- SuSE 11 SP3
- Debian 7 through 7.2
- HPUX 11i v3 Update 12

For additional information, refer to the *StorNext 5 Compatibility Guide* available here:

http://www.quantum.com/sn5docs

Web Services

The following web services API has been enhanced:

- The **getMediaReport** web services API no longer reports extraneous data that was previously appended at the end of the report. It now lists all fields in fully parse-able JSON or XML.

Details on this web services API are in the *StorNext 5 User's Guide*, **Appendix D**.

CLI Commands

The following Storage Manager CLI commands have been enhanced to provide optional XML or JSON output. The affected commands are:

- fschmedstate
- fschstate
- fsclassinfo
- fsclean
- fsconfig
- fsmedcopy
- fsmedout
- fsmodclass
- fsschedule
- showsysparm
- vsarchiveeject
- vsarchiveenter
- vsarchiveqry
- vsarchivevary

The commands listed above now support an optional [-F type] command line option, where type can be TEXT, XML, or JSON.

For additional details, see the *StorNext MAN Pages Reference Guide*.

Additionally, three Storage Manager commands have been enhanced to report additional information. These commands and the changes are listed below.

- The **fsconfig** command now reports the Vendor ID, Product ID, and Product Revision for tape drives.
- The **fsmedinfo** command now reports additional information on storage media. The following additional fields are now part of the report:
  - Media Class
  - Current State
  - Assignment
  - Action State
  - Location State
  - Current Archive
  - Pending Archive
  - Import Date
  - Move Count

Furthermore, the report fields appear in newly organized logical groupings.

- The **showsysparm** command now contains a header and footer like the other standard Storage Manager commands.

CPU Power Saving States Notifications

Beginning with StorNext 5 release 5.0.1, StorNext 5 detects potential performance problems caused by CPU power saving features on the StorNext MDC or client. The detection is performed during software installation and at system start-up. If CPU power saving states are detected, a message is displayed.

On Windows, this causes the installer to display a new dialog box during new installations or upgrades.

Also, the Linux `install.stornext` script prompts when CPU power saving states are detected.

While not a requirement, Quantum recommends that CPU power saving states are disabled on systems running StorNext. See the section titled **Disabling CPU Power Saving States** in the *StorNext File System Tuning Guide* (part number 6-68046-01) for additional information.

Customer Reported Issues and Enhancements Addressed in StorNext 5 Release 5.0.1

The following table lists the customer reported issues and enhancements addressed in StorNext 5 Release 5.0.1.

Operating System	Change Request Number	Service Request Number	Description
All	27692	963346, 1602338, 1638158, 1561708	Write event generated for private memory mapping
All	33440	1262314	GUI: Failure to add tapes does not specify "library is full"
All	33974	1576244	Too easy to choose wrong file system for the HA shared
All	39465	1459366	StorNext read/write error messages misinterpreted as metadata problems
All	43934	n/a	StorNext GUI scan of Lattus Namespace to use Username/Password
All	44220	1597622	Message "dm_process_destroy: no xattr for inode" should show file name and mention snfsdefrag
All	44605	1637692, 1638158, 1561708	fs_cpyresp core dump due to segmentation fault

Operating System	Change Request Number	Service Request Number	Description
All	44747	1603008, 1622612	Files on replication target should contain data but are sparse files with no data blocks
All	45072	1609850	fsrelocate core dumps
All	45510	n/a	Display LUN number in the GUI cvlabel
All	45645	1640524	Support RHEL 6.5 (expected by end-of-2013)
All	45649	n/a	Support Debian 7.2, released 2013-10-12
All	45679	1615756, 1580686	cvfs umount causes panic on HP-UX B.11.31 Update 12
All	45758	1602338	getCandidates uses "select max(mtime) from storecand" query
All	45903	1625772	Client performance degradation after 3.0.2 client upgraded to 4.3.2
All	46027	1631446	StorNext should warn if it detects CPU Power Saving States are not disabled
All	46116	1638828, 1640408	Deadlock between quota inode i_lock and inode freelist root node b_lock
All	46174	1644654, 2200188, 3302354	TSM: snbackup: DBI error: Duplicate entry '2-vsn001-1-1-68' for key 'PRIMARY'
All	46466	1621520	TSM processes use of "Well-Known" ports can cause other services to fail
All	46473	1646978	Data overrun for fsmpm PORTMAP_GET_CLIENT_CCTLS reply with > 7 file systems
All	46750	3353490, 3304856, 3351214	Journal replay of pre-5.0 journal fails in detection of oldest log record
All	46770	n/a	Wrong dir stored in bucket leads to failure of bulk unlink and create
All	46825	n/a	Do more checking on replication target and pull data if the tgt file's content is missing
All	46896	n/a	Fixing sparse file due to "open for write" after running replicateforce
All	47021	n/a	FSM PANICs during conversion because inode_buffer_validate fails
All	47046	2200556, 3304342	OS issue causing a tape file corruption (duplicate I/O block written) when **pse_snapshot** is executed

Operating System	Change Request Number	Service Request Number	Description
Linux	42111	1444740	RAS:0602:53 tickets generated when StorNext snapshot is taken
Linux	44114	3352724	Storage Manager: DDM client: WAS iopath has exceeded its failure threshold and is being taken offline
Windows	45748	1621728	Windows BSOD (PAGE_FAULT_IN_NONPAGED_AREA) in ExtApiGetExtlist()

What's New in StorNext 5

Purpose of this Release

StorNext 5 is a new generation of Quantum StorNext that performs faster, scales farther, and expands flexibility. StorNext 5 has been built from the ground up with a new architecture designed to meet the needs of today's evolving digital workflows.

New Features and Enhancements in StorNext 5

Redesigned to be Faster

While many of the features in StorNext 5 improve the usability and capabilities over previous generations, StorNext 5 is all about performance. Whether it's better performance in larger scale deployments of more users or more files, or mixed topologies where users share Fibre Channel and IP networks, StorNext 5 just works better. StorNext unleashes hidden performance in existing metadata controllers, and improves overall storage network performance by lowering latency. When combined with Quantum appliances, highly optimized and efficient StorNext 5 deployments can be achieved.

Operations perform faster with StorNext 5, because of its improved underpinnings, with better memory usage efficiency and caching techniques, more efficient multi-threaded operations for improved CPU usage, and highly optimized on-disk metadata.

Small file performance improves by up to 5-10x plus significant performance improvements for all file operations. In addition, large-scale file creation can occur concurrently with high performance access. With StorNext 5, customers will not need to set up a "delete window" of limited file system access.

No downtime upgrade from StorNext 4.7

StorNext 4.7 customers with StorNext Metadata Appliances or High Availability can upgrade to StorNext 5 without having to take the file system offline.

Metadata Controller upgrade with no downtime and no interruptions will be part of all future releases.

Improved Stability by Eliminating Fragmentation

StorNext 5 addresses fragmentation challenges by virtually eliminating the problem from the start. By re-architecting the file system's on-disk structures, metadata tracking capabilities, diagnostic and report capabilities and caching techniques, StorNext 5 not only reaches new heights of performance, but the majority of file fragmentation is eliminated, improving the stability and long-term performance of all file system operations.

Reduced Latency

Availability of flash memory and solid state disks has promised a revolution in file storage performance. In most cases, although the disk performance is better, the benefits are not as great as promised. This is simply because flash memory

operation is different from that of spinning disk, and system and controller changes must be made to harness the maximum power of this technology.

With StorNext 5, Quantum engineers considered all possibilities and requirements and designed StorNext 5 to take full advantage of the unique capabilities of flash storage (SSD). Improvements from flash are seen in nearly every application, and with StorNext 5, metadata controllers and primary storage will benefit greatly from the higher throughput and reduced latency of flash. As a complete end-to-end solution, Flash based MDCs and Flash-enabled Quantum storage can greatly enhance the performance of primary storage and often reduce the capacity needs allowing users to deploy lower cost nearline and long-term storage, providing a more cost effective and better balanced storage infrastructure.

Enables New Metadata Appliances

StorNext 5 enables new StorNext 5 Metadata Appliances, specifically designed to leverage the modern unpinnings of StorNext 5, to meet the growing throughput and capacity needs of Quantum storage offerings. Carefully tailored memory and CPU selections and associated components take full advantage of the new and improved StorNext metadata format, journaling, and caching.

Evolved beyond the SAN

Just a few years ago, the only way to ensure performance and stability was to use a pure Fibre Channel network. But the workflow model has changed, it's now about mixed performance and dispersed workflows.

No single topology is ideal for today's requirements. In many cases faster IP network technologies such as one and ten Gb Ethernet are good enough for certain applications and more cost effective than dedicated Fibre Channel networks. Small file sizes are often associated with IP networks and traditionally have not performed well with Fibre Channel or StorNext. That all changes with StorNext 5 which now delivers superb performance and efficiency in both Fibre Channel and Ethernet with greatly enhanced small file throughput.

In addition, today's workflows extend beyond a single facility to incorporate users and applications in geographically dispersed locations. These workflows need a central content or data repository that allows efficient access to all team members regardless of location or client platform. StorNext 5 retains SAN support and also better supports the increasing number of SAN-free workflows. In addition, StorNext 5 is integrated with Lattus object storage for policy-managed, scalable, long-term central repositories that also efficiently serve geographically dispersed teams.

Up to Five Billion Files

StorNext 5 will allow customers to scale to 5 billion files per Metadata Controller, five times the number of files previously supported, and this support comes with better performance than previous versions of StorNext. Metadata Controllers with StorNext Storage Manager can scale to 1 Billion files.

Since the supported file count is per-MDC, this could be a single file system of 5 Billion files (1 Billion files with StorNext Storage Manager), or eight file systems

of 625 Million files (125 Million files with Storage Manager), or something in between.

Linear Tape File System (LTFS) Media Format Support

StorNext 5 supports both the open standard LTFS format and Quantum ANTF formats for writing data to LTO tape.

LTFS offers a standard format that is readable outside of StorNext, suitable for long-term vaulting, content delivery and collaboration with non-StorNext environments. Quantum ANTF format offers higher performance, and the ability for a single file to span multiple tapes, something the LTFS standard does not offer.

Customers select the tape format when defining their Storage Manager policies. Policies with more than one copy can be configured to use both LTFS and ANTF.

For additional information, see **Linear Tape File System (LTFS) Media Format Support** in the *StorNext 5 User's Guide*.

More Flexible Stripe Group Affinities

StorNext Affinities allow customers to tie files to stripe groups so special files can use specific disks within the same file system. StorNext 5 makes affinities easier to use by allowing setting affinity by file extension.

Simplified Tuning

In previous StorNext releases, users set the average files per directory during file system creation to optimize file system block size (`FsBlockSize`). With StorNext 5, `FsBlockSize` is now set at an optimal value of 4 KB and is no longer tunable. This simplifies the tuning for users.

Additionally, StorNext now provides default parameters based on RAID type.

The recommended Journal size is now 64 MB.

The Maximum Connections configuration item, Thread Pool Size, and Data Migration Thread Pool Size have all been eliminated from the GUI because they are now sized dynamically in the FSM.

The configuration file parameter Directory Warp has been eliminated from the GUI since it no longer has meaning in StorNext 5.

The configuration parameter **abmFreeLimit** has been deprecated as it no longer has meaning in StorNext 5.

See the *StorNext 5 File System Tuning Guide*.

Improved cvfsck

The Check File System functions available from the GUI and the check available from the command line (**cvfsck**) have been improved for StorNext 5.

In StorNext 5, **cvfsck** runs much faster because of metadata and caching enhancements.

In previous releases, **cvfsck** would directly repair the file system when problems were found. This led to a best-practice of running **cvfsck** in read-only mode first to assess the extent of the problems, then running it again to fix the file system. Having to run the utility twice was time consuming.

With StorNext 5, the default command **cvfsck** with no options will accumulate file repairs needed in the /tmp directory. At the end of the file checking process, the user is asked whether to accept or reject the changes. By In StorNext 5, `cvfsck` runs much faster because of metadata and caching enhancements. the need to run the read-only check first, the file system repair process is completed much faster, with user control over whether or not to make the fixes. Best practice remains to run the **cvfsck** command in an unmounted, idle file system.

Deleting Filtered RAS Tickets

Customers will appreciate the new StorNext 5 enhancements for managing RAS tickets. Using the filter options, customers can view a list of the tickets they are interested in, including filtering by a date range. Customers can now easily delete all tickets in a given date range. Previously, customers who had a large number of tickets to delete could have difficulty deleting them all without calling support.

MetaDump Replaced

The StorNext 5 release eliminates the **snmetadump** command and `restore_journal` files completely. The FSM now handles all restore journal processing in real time without the need to generate sequenced journal files.

Archive Compare

Under **Configuration > Storage Destinations**, selections have been added to perform some of the library maintenance functions of the archive compare utility through the GUI. These new features can be run for a single library or for all libraries at one time.

New Media Report Filters

There are new filters available on the **Reports > Media** page. Now there are filters for:

- Destinations
- Library
- Media Class
- Policy Class
- Media ID

These filters should help customers find the reports they want more easily.

StorNext Metrics

With StorNext 5, the Advanced Reporting features have been integrated into the StorNext GUI. The previous Advanced Reporting item from the Tools menu was

linked to a third-party application. These tools can now be found under the **Reports > StorNext Metrics**.

Web Services

Two new web services have been added for StorNext 5:

- The **getSNAPIVersion** API returns the Web Service version and the output from the **cvversions** command.
- The **getSystemStatus** API returns the software version and the status of DSM, TSM, MSM, MySQL and SRVCLOG.

Details on these new web services are in the *StorNext 5 User's Guide*, **Appendix D**.

Two web services have been removed for StorNext 5:

- **getMountPoints**
- **setMediaMoveInfo**

The *StorNext CLI Reference Guide* was deprecated as of release 4.7; the link to that document has been removed from the **Help > Documentation** list for StorNext 5.

Security

Under **Tools > User Accounts**, users with the **Manage Users** privilege can disable any GUI accounts, including **admin** and **service**.

New snupdate Service for Linux Clients

The new **snupdate** service allows RHEL 5 and 6 and SUSE 10 and 11 users to install, upgrade, remove, and roll back StorNext client software packages for those platforms. These packages can be applied immediately, or they can be staged and updated at a later time. Multiple versions can be copied into the snupdate repository and snupdate will track the different versions.

For more information on **snupdate**, see the *StorNext 5 and StorNext FX 5 Upgrade Guide*.

Customer Reported Issues and Enhancements Addressed in StorNext 5

Addresses over 100 Customer-Reported Issues

StorNext 5 continues steady improvements by addressing over 100 customer reported issues.

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

（不適用）

Operating System	Change Request Number	Service Request Number	Description
All	14630	588194, 1202400, 1400624, 1369968, 1283890, 1485518, 1518636, 1546344	The `sl_event_details` table can grow large. Add functionality to the StorNext GUI to allow old and closed tickets to be deleted.
All	27479	1210294, 1260244	cvfsck does not detect directories with out-of-order btree keys
All	27888	951446, 1221556, 1242372, 1242518, 1262896, 1281998, 1302038, 1392598, 1418606, 1573296	MSM does not detect when its drive configuration gets out of sync with Archives internal configuration
All	29740	1062606 1259404	cvfsck doesn't verify/reconstruct directory hash type in superblock
All	31072	1119366, 1234184, 1297634, 1375442	cvfsck can't repair corrupt iel chain
All	31547	1148100, 1482826, 1486212	Failed snbackup due to existence of metadump .gz files in /usr/adic/database/metadumps includes (errant?). "Could not get media list" messages
All	31623	1153036, 1503532	snbackup took 22 hours to run because of snmetadump execution time
All	32079	n/a	"phdist exists in electronic downloads, not released media"
All	32844	1407132, 1471604	4011LCR ASSERT failed "(((alloc_blockp[byteoffs] >> nextbit) & 1) && 1) == 0", trans_abm.c, line=1689
All	33440	1262314	GUI: Failure to add tapes does not specify "library is full"
All	33803	1261402, 1275304, 1404170, 1449740	Warning message generated from CiReply about failure to read from admin file
All	35125	n/a	Punch hole not disabling extent warp in many cases

Operating System	Change Request Number	Service Request Number	Description
All	35248	1289510, 1288908, 1468858	Space allocation transaction appears missing in restore journal
All	35546	1323664	cvfsck doesn't detect free list problems
All	35872	1330586, 1482384, 1581236, 1591714, 1594108	fs_resource segfault causing TSM to shut down
All	35876	1338082, 1383514, 1395610, 1426164, 1426168, 1554418, 1592298, 1563290	ls command hangs - directory node has stale hash value
All	36106	1157710, 1183706, 1341218	2nd Heavy buffered write activity completely stalls buffered reads
All	36256	1357646, 1381510	Confusing error when cvfsck runs out of space in temp directory it's using
All	36412	1360670, 1361986, 1373274, 1449138, 1548032	Conversion of HA secondary fails with "Error invoking action"
All	36545	1372090, 1367328, 1482682	No fsm restart with 2 core's in 1 hour
All	36730	1375442	cvmkfs needs to be logged to trace directory like cvupdatefs and cvfsck
All	36872	1363824	Upgrade fails if file bp_stop exists but blockpool executable does not
All	36940	1323664	cvfsck requests run after free list clobbered

Operating System	Change Request Number	Service Request Number	Description
All	37005	1375808, 1399648, 1365752, 1528828, 1545648, 1549526, 1559314, 1588006, 1485970	fspolicy timeout does not release MSM mount request
All	37087	1386784	FS hung because only buffercache for fsm was 128k with 64k blocksize which is 4 buffers
All	37183	1364412	usurp from active peer in one pass saying it was idle
All	37184	1364412	nss way too chatty to easily determine causes of changes, reason for usurps
All	37490	1392410	snmetadump -a fails if temporary directory cannot be removed
All	37704	1406528	GUI cannot enable/disable DDM SDISKs on a host
All	38702	1434230	cvfsck fails to correct abm after first run
All	38735	1435124	No errors in StorNext logs when no drives available in drive pool
All	38775	1385240	Balance allocation strategy not always allocating from stripe group with most free space.
All	38948	1438176	Look at addressing the flood of not accessible and soft mount RAS from fsmcom.c
All	38973	1470288, 1470988, 1509568	snfsdefrag man page suggests cvfsck command that is wrong
All	38992	1444512	FSM must be restarted when moving from temporary to permanent license
All	39103	1476304	cvgather - no /usr/cvfs/qustats/MDC directory
All	39169	1428860	snmetadump insists on consecutive allocation of directories when space not available
All	39217	1456410	DDM does not recognize fsddmconfig'ed MDC hostname reliably
All	39464	1459366, 1526280	Need more info in error message about failure to respond to proxy client
All	39473	1419636	Rebuild policy fails due to XattrSearchSNEA() bug

Operating System	Change Request Number	Service Request Number	Description
All	39579	1468688	cvlog SUMMARY messages make it difficult to grep for important messages
All	39820	1475960	GUI "Purge Media" fails with "Empty CLI command" error when its the mediaclass that is in error
All	39841	1469732	SNAPI CvApi_GetSgInfo displays incorrect value for free blocks
All	40237	1480756	cleanup still leaves files under a cleaned up sub-directory
All	40433	1575178, 1588104, 1600246, 1572686	Same Request Ids being used for multiple commands
All	40495	1490680	Adding media to S10K with LTO5 drives fails due to "unknown medium type" error
All	41647	1429488, 1528528	fsm fails after upgrade from 4.2.1 with ASSERT failed (ip->i_idinode.idi_flags & InodeFlagPendFree) == 0
All	41941	1374782	Updates needed for rebuild policy and mapping code to support new metadump structure.
All	41961	1526850 1608946	fsmedcopy -r fails with "Internal processing failure"
All	41964	1515914	Cancelled fsretrieve didn't cancel, instead requested copy 2 tapes on very large file retrieve.
All	42305	1551542	Proxy clients fail to mount if .auth_secret file is present on them
All	42403	1536424	Ability to mix generations of drives/media in the same library (at least for fsmedcopy operations), e.g. LTO4 & LTO5
All	42460	1554048	fsschedule "--" option for activevault can easily exceed the 128 character limit"
All	42508	1553148	Incorrect debug message in mdt1dev.c module
All	42884	1556646	cvfsck hang forever on directory with 2 hash for the same block
All	42887	n/a	Add "-R directory	-B batchfilename" options of fsretrieve to fsstore, fsrmdiskcopy and fsfileinfo
All	43028	1493552	Implement HA fail-over upgrades for StorNext
All	43046	n/a	Add "-R directory	-B batchfilename" options of fsretrieve to fsstore, fsrmdiskcopy and fsfileinfo - GUI changes

Operating System	Change Request Number	Service Request Number	Description
All	43066	1553148, 1342716, 2373206, 1272914	fsmedinfo displays 0 available space remaining for an SDISK even though df -h shows available space
All	43105	1526692	double file close issue leads replication failure
All	43189	1571554	fsstate man page should indicate when the DISMOUNT status will be seen
All	43191	1568720	The cvfsctl driver advertises as a disk device - causes udev activity on open
All	43238	1569522	Existence of nss_cctl.xml prevents from setting quotas in cvadmin
All	43294	n/a	Webservices: Implement webservice API calls for directory quota.
All	43316	n/a	fsstore (new multiple ipc messages version) consumes cpu cycles instead of blocking on queue
All	43355	1577094	snfsdefrag -r SIGSEGV's when running on deeply nested file system
All	43445	1547876	health_check fails on "Database" test via GUI
All	43587	1561322	fs_store daemon uses "select count(*)..." query
All	43725	1330892, 1331826, 1363418, 1394636, 1512732	Additions to the archive_cmp.pl utility to support interaction with GUI ability to use it
All	43947	n/a	GUI Lattus-M manual namespace mode not available with empty scan result set
All	44099	1595024	Thousands of RAS messages got generated and sent as email notifications
All	44165	1593118, 1595470	Failed event file processing can lead to data loss
All	44172	1594202	fsrminfo should warn user about potential data loss when used incorrectly
All	44221	1578760	cnvt2ha.sh should use unique logfile names to preserve logging of failures
All	44673	1614178	RAS emails may have incorrect "SR Problem Code"

Operating System	Change Request Number	Service Request Number	Description
All	44742	1511868, 1530470, 1561322, 1568238, 1589984, 1602338	MySQL Reconnect code - Update sndbapicxx/sndbapi versions (4.3.5/4.3.3)
All	44797	1610828	fsclean on sdisks fails after 4.7 upgrade
All	45071	1609850	fsrelocate gets an error trying to move a file to an affinity
All	45471	1608946	fslocate fail with valid pfile_key
All	45595	1627394	`innodb_buffer_pool_size` in my.cnf is overwritten when upgrading
All	45663	1626338	Update the manpage for the command **snhamgr** to provide documentation on how to switch to the primary HA node.
Linux	23349	725254, 1154372	ABMFreeLimit needs better fix
Linux	23351	1048584, 1154372	Free space fragmentation needs work
Linux	25318	1582578	Man pages missing for the following MSM CLI commands
Linux	29735	1235990	Unable to delete Tape Library from the GUI
Linux	30210	1156698	Make it easier to remove library from the GUI - update vsarchiveconfig
Linux	34970	1304484, 1331828	TSM does not pick new sg device for a replaced tape drive without a restart
Linux	36129	1204952, 1316996, 1517194	Client process trying to open an "exclusive" metadata/ journal LUN
Linux	36334	1326764, 1558768, 1555748	HA conversion of primary MDC fails because vsping executing when rpc service is stopped
Linux	36626	1377158, 1433956, 1514658, 1526064, 1536138, 1611808, 1613000, 1620378	GUI fails to expand a file system that has some stripe groups down

Operating System	Change Request Number	Service Request Number	Description
Linux	38433	1417768, 1454282, 1551416	Fsync causing PANIC: /usr/cvfs/bin/fsm "OpHangLimitSecs exceeded"
Linux	39111	1523966, 1379394, 1440578, 1482240	StorNext: Storage Manager: System backup failed: Failed to optimize dump for fs
Linux	41775	1537406, 1535352	snmetadump failed due to Transaction out of order after MDC failover
Linux	41908	1531432, 1538246, 1474334, 1558768, 1555748	Check for portmap rpm depenedancy at install time required by MSM
Linux	43414	1304484, 1331828, 1598210	GUI: Ability to set a new /dev/sg device path for a replaced tape drive without a restart of TSM
Linux	43449	1584402	Creating a new mountpoint, The GUI limits the number of available characters in a filesystem mount name to 25.
Linux	43450	n/a	Storage Manager support for native LTFS tape format
Linux	44222	n/a	MAS Agent can be disrupted by temporary absence of its config file -- Web Services
Other	44349	1589858	HP-UX: panic due to sleep while in interrupt context
Other	45940	1634032	Exposed improper settings for multipath devices in a clustered environment.
Windows	42756	n/a	SNFS 4.x not caching case insensitive file name matches in name cache
Windows	43933	1598718, 1603770, 1614240, 1624040, 1618496, 1624040, 1625022	LDAP floods logs when SID is not found in AD
Windows	44266	1574936	StorNext installer fails with 'lodctr cvperf.ini' error 2001 and installation aborts
Windows	44925	1611896	SNFS does not support Set-GID on Directories correctly on Windows
Windows	45527	1633428, 1632848	Windows SMB clients sometimes cannot mount StorNext shared file systems on Windows 2012

Visit www.quantum.com/ServiceandSupport for additional information and updates for StorNext.

64-bit inode Numbers are Exposed to Linux Stat Operations

Beginning with StorNext 5, 64-bit inode numbers are exposed through the stat() system call on Linux platforms. (Prior releases exposed 32 bits of the full 64 bit inode number.)

While most applications should not be affected by this change, some legacy applications could experience compatibility issues. If compatibility issues are expected, a compatibility mount option is available.

See the `compat32` option in the `mount_cvfs` man page for more details.

Note: Pre-StorNext 5 Linux clients will not expose 64-bit inode numbers, even after the MDC is upgraded to StorNext 5.

StorNext Limits

• The recommended maximum number of mounted StorNext file systems on a single host is 16. For additional information, see the section titled **StorNext Limits** in the *StorNext 5 User's Guide* available here:

http://www.quantum.com/sn5docs

• StorNext software on non-StorNext metadata appliances has been tested up to eight file systems per metadata controller.

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

A database schema update is applied to Storage Manager when upgrading from StorNext 4.3.x to StorNext 5. This update takes approximately one hour for every hundred million entries in the filecomp tables.

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.

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 5.2.2.

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 (PFR) is not compatible with Lattus.

Refer to the *StorNext Partial File Retrieval Compatibility Guide* at this location for information about compatibility between PFR and StorNext 5.2.2:

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 5.2.2:

http://www.quantum.com/sn5docs

Apple Xsan

Xsan is software that enables multiple Mac computers to concurrently access hundreds of TBs of content on Xserve RAID or Promise RAID storage over high-speed Fibre Channel so creative artists can 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 5.2.2:

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

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 44
- StorNext Storage Manager Known Issues on page 46
- StorNext GUI Known Issues on page 50
- StorNext Installation, Replication, HA and Other Known Issues on page 52

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	47709	1531430	If a managed file system is deleted, a TSM configuration file will not be properly updated. If this occurs, startup of the TSM may result in failures.
			Workaround
			To correct the situation (following the deletion of a managed file system), edit the `/usr/adic/TSM/config/filesystems` file and remove the entry that matches that of the managed file system that was used.

For example:

• A system has 2 managed StorNext file systems:

  `/stornext/cvfs1`

  `/stornext/cvfs2`

• The `/usr/adic/TSM/config/filesystems` file contains two entries similar to:

  `/stornext/cvfs1       75    85    true       1 true`

  `/stornext/cvfs2       75    85    true       1 true`

• The user deletes the `/stornext/cvfs1` file system.

• The `/usr/adic/TSM/config/filesystems` file will need to have the matching line deleted such that the only remaining file system in the file is:

  `/stornext/cvfs2       75    85    true       1 true`

• Restart **StorageManager**. |

Operating System	Change Request Number	Service Request Number	Description
All	54834	3505208, 3516356	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. **Note:** This problem can also occur with other programs that set EOF beyond the end of data. This problem does not occur if Windows Explorer encounters an error while writing the file; Windows Explorer will delete the partially written file.
			Workaround
			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. For more information on this option, see the man page mount_cvfs(8). The sparse option will introduce some overhead when using Windows Explorer. Before setting the protect_alloc option, see the sparse option in mount_cvfs(8) for a description of how it changes StorNext behavior.
All	56162	n/a	**Description**
			When **snquota** is used on systems running SLES12, Debian7, or RHEL7, the command may fail with a message such as: snquota: error marking namespace: error finding mountpoint for file system 'snfs1'
			Workaround
			To workaround this error, execute **snquota** on a system running a different operating system such as the metadata controller.

Operating System	Change Request Number	Service Request Number	Description

All	57304	3561252	Due to a limitation in the Linux **rpc.mountd** 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 `/stornext/snfs1`, then the following export is exposed:
			
			`/stornext/snfs1/myshare *(rw,sync)`
			
			Whereas this export entry is not:
			
			`/stornext/snfs1 *(rw,sync)`
			
			Also, the NFS server must be running Linux and have StorNext 5.x installed and the version of **rpc.mountd** on the system must correspond to version of nfs-utils less than 1.9.0. For example, RHEL6 is exposed but RHEL7 is not.
			
			Finally, the directory being exported must have an inode number that is greater than or equal to 4294967296 (2^{32}). This can be checked by executing the following command:
			
			`ls -id`
			
			StorNext file systems having large inode numbers will usually have one or more of the following attributes:
			
			1 A file count greater than 16 million.
			2 Multiple metadata stripe groups.
			3 Stripe groups containing mixed data and metadata.

			Workaround
			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 **compat32** SNFS mount option.

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 47.

Table 2 StorNext Storage
Manager Known Issues

Operating System	Change Request Number	Service Request Number	Description
All	43320	1581004	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.
			Workaround
			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:
			1 Change the **FS_T10K_BLOCK_FACTOR sysparm** from 8 to 32 by adding the following entry to /usr/adic/TSM/config/fs_sysparm_override:
			`FS_T10K_BLOCK_FACTOR=32;`
			Note: 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.
			2 Restart Storage Manager to ensure the change in **Step 1** goes into effect:
			`# tsmstop`
			`# tsmstart`
			3 Verify the **FS_T10K_BLOCK_FACTOR sysparm** contains the new value:
			`# showsysparm FS_T10K_BLOCK_FACTOR`
			`FS_T10K_BLOCK_FACTOR=32`
			4 Save the current copies of your /etc/fstab on the MDCs and the clients.
			5 Modify /etc/fstab on the MDCs and the clients to use the auto_dma_write_length and auto_dma_read_length mount options as follows:
			`snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0 0`
			6 Unmount and re-mount your file systems.
			7 Use new T10K media to store a copy of the file from the disk.
			Note: **Step 7** 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.

Operating System	Change Request Number	Service Request Number	Description
All	46815	n/a	Executing the command **snbackup -s** while a full or partial backup is running may result in a message that /usr/adic/TSM/internal/locks/backup.lf is in an invalid format.
			This is due to the **snbackup -s** process reading the backup.lf status file while the backup process is updating it.
			Workaround
			Ignore the message; to clear-up the process, re-execute the command **snbackup -s** (provided that the backup is not writing to the backup.lf status file while **snbackup -s** is trying to read it again).
All	47833	n/a	**Description**
			When copying files between media using the CLI command **fsmedcopy**, 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.
			Note: This behavior may cause file data segment distribution to be sub-optimal on the destination media.
			Workaround
			Currently, a workaround does not exist for this Known Issue.
All	52548	n/a	**Description**
			When executing the command **fsretrieve -B batchfile**, the socket will hang when sending too large of an IPC message.
			Workaround
			For the command **fsretrieve -B batchfile**, limit the number of file names in the batchfile to 3000 or less. This limitation is to prevent the socket from hanging in the event a retry of all files in the batchfile is necessary.

Operating System	Change Request Number	Service Request Number	Description
All	54978	n/a	The StorNext Web Service interface GetFileAttribute propagates an incorrect value for the location field. Prior to StorNext 5 release 5.2, locations could include TAPE and DISK AND TAPE, in StorNext 5 release 5.2 and StorNext 5 release 5.2.0.1, the locations are returned as ARCHIVE and DISK AND ARCHIVE. **Note:** This can also impact Web Service Commands.
			Workaround
			In order to work around this issue, add the following line to the /usr/adic/.profile file. export SNAPI_COMPAT_MODE=TRUE Then restart the GUI by executing the following command: **/usr/adic/tomcat/bin/service stornext_web restart** This will return the output of the GetFileAttribute to respond with TAPE as a location rather than ARCHIVE.
All	41413	1504258, 1635952, 3471000, 3526376	**Description**
			In certain conditions when Storage Manager receives an end of tape indication early, SCSI sense messages may be logged by the **fs_fmover** process. These the messages will have a format similar to the following: Dec 28 20:38:00 MDC-Hostname sntsm fs_fmover[29837]: E1201(8)<1034815852>:fsScsi1311: {2}: Check condition: op=0Ah key=00h asc=00h ascq=02h END OF PARTITION/MEDIUM DETECTED
			Workaround
			These messages do not cause operational problems and can be safely ignored.
Linux	43246	1568166	**Description**
			If connectivity faults occur when transferring data from StorNext Storage Manager to and from Lattus, consider using less sensitive offline of iopaths and media.
			Workaround
			The workaround for this issue is to add the following lines to /usr/adic/TSM/config/fs_sysparm_override and then restart StorNext Storage Manager: FS_THRESHOLD_INC_NUM=0; MEDIA_SUSPECT_THRESHOLD=32767; **Note:** If you store data to tape, changing these settings will affect the normal behavior for dealing with unreliable media and tape drives.

Operating System	Change Request Number	Service Request Number	Description
Linux	45718	n/a	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. This problem is described by RedHat Bugs 663068 and 672981. https://bugzilla.redhat.com/show_bug.cgi?id=663068 https://bugzilla.redhat.com/show_bug.cgi?id=672981 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.
			Workaround
			Contact your Linux vendors for guidance on which OS versions you should run to avoid the issue.

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	The Safari browser becomes unresponsive when you attempt to configure an Email server using the StorNext GUI.
			Workaround
			To workaround this issue, perform the following procedure: 1 Shut down the Safari browser window(s). 2 Restart the Safari browser, and then retry operation. 3 Uncheck the **Verify SMTP Server Connectivity** box, and then retry the operation. 4 Set **Authentication** to **NONE**, and then retry operation. 5 Disable the Safari **User names and passwords** AutoFill under **Safari > Preferences > AutoFill**, and then retry operation.

Operating System	Change Request Number	Service Request Number	Description
All	53986	n/a	If your Q-Cloud Archive™ account has been discontinued, you are not able to delete Q-Cloud Access ID using the StorNext GUI.
			Workaround
			The StorNext GUI provides the capability to configure and delete Q-Cloud devices; however, to completely deconfigure Q-Cloud on your system, you will be need to manually remove the files by executing the following commands:
			`rm /opt/quantum/qcc/instance/products`
			`rm /opt/quantum/qcc/instance/secret`
			To remove the **Q-Cloud Access ID** and clean-up the rest of the deconfiguration, perform the following:
			1 Using the StorNext GUI, navigate to the **Configuration > Storage Destinations > Q-Cloud** page.
			2 On the **Configurations > Storage Destinations > Q-Cloud** page, click **Manage Keys...** to display the **Manage Keys** page.
			3 In the **Q-Cloud Access ID** field, select the entire string, and then press **BACKSPACE**.
			4 Click **Reset**.
All	57627	n/a	**Description**
			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).
			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.
			Workaround
			There is currently no workaround for Change Request Number 57627.

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	A database index named `classndxatimeme` will be automatically added to the tmdb.tier000files% and tmdb.tier001files% tables upon starting TSM for the first time after upgrading to specific StorNext 5 releases.

StorNext 5 Releases affected: 5.0.1, and 5.1

Upgrading from: 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.

> **Note:** 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.

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.

Workaround

To minimize TSM downtime after upgrade, the `classndxatimeme` index can be created prior to performing the upgrade using the `index_tierfiles.pl` 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.

> **Caution:** The following script should only be executed against StorNext release as indicated in the previous section.

To manually add the index:
1 Login to the primary MDC.
2 Source the profile:

`. /usr/adic/.profile`

3 Change to the directory where `install.stornext` resides on the installation media. For example:

`cd /tmp/stornext/stornext_full/RedHat60AS_26x86_64`

4 Verify that the database is up by running:

`mysql_control start`

5 Execute:

`./TSM/index_tierfiles.pl`

Operating System	Change Request Number	Service Request Number	Description
All	53933	n/a	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. 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 **Workaround**.
			Workaround
			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. **Note:** In Lattus 3.6.X this WAR is not required. 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. 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.
All	56135	n/a	**Description**
			The StorNext GUI does not show the RHEL7 and SLES12 clients in the list of clients available for download. **Note:** This issue only affects systems running StorNext 5 Release 5.2.1.
			Workaround
			To manually download the client installers for RHEL7 and SLES12: 1 Open an SSH connection to the MDC Node (either MDC Node will work for StorNext Metadata Appliances or HA systems) using the IP address assigned to that Node on the MDC/Metadata network. Manually copy the Redhat7 or SuSE12 **.bin** file from **/usr/cvfs/ CLIENTS** to an external USB thumb drive or copy over the network to the client system. 2 Continue with the installation procedure for the client for your operating system as described in the *StorNext Installation Guide* or the online help.

Operating System	Change Request Number	Service Request Number	Description
Linux	53601	n/a	The Q-Cloud configuration files are not automatically propagated in an HA pair via `/usr/adic/util/syncha.pl` from the primary system to the secondary system.
			Workaround
			Any time a Q-Cloud storage destination is added or modified, or anytime that a new secondary node is created via the HA conversion process, the Q-Cloud configuration files `/opt/quantum/qcc/instance/product` and `/opt/quantum/qcc/instance/secret` must be manually copied from the MDC that is primary to the MDC that is secondary.

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

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/

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

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

For worldwide support:

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

Worldwide End-User Product Warranty

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

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