



# Release Notes

---

<b>Release</b>	5.4.0.1
<b>Supported Product</b>	StorNext 5
<b>Date</b>	November 2016

---

## Contents

What's New in StorNext 5 release 5.4.0.1 .....	2
Purpose of this Release .....	2
New Features and Enhancements in StorNext 5 release 5.4.0.1 .....	4
Fixed Issues and Enhancements Addressed in StorNext 5 release 5.4.0.1 .....	6
StorNext Compatibility .....	15
Quantum Operating System Upgrade Support Policy .....	15
StorNext and Linux Interoperability .....	15
Supported StorNext Upgrade Paths and Upgrade Considerations .....	16
StorNext Software Upgrade Matrix .....	16
Considerations for the StorNext File System Directories .....	16
Journal Size Guidelines .....	16

---

© 2016 Quantum Corporation. All rights reserved. Artico, Be Certain, DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, GoVault, Lattus, NDX, the Q logo, the Q Quantum logo, Q-Cloud, Quantum, the Quantum logo, Quantum Be Certain, Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo, vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Quantum specifications are subject to change.



Distributed Data Mover (DDM) Guidelines .....	16
Considerations When Upgrading NFS Server Nodes to StorNext 5 .....	17
Database Schema Update During Upgrades .....	17
Compatibility Between StorNext and Other Products .....	18
NAS .....	18
Infiniband .....	18
Lattus .....	18
Partial File Retrieval .....	18
StorNext Web Services .....	19
Apple Xsan .....	19
Supported Browsers .....	19
General Considerations .....	19
Checksum Performance Considerations .....	19
Upgrading Appliances .....	20
Known Issues .....	20
StorNext File System Known Issues .....	20
StorNext Storage Manager Known Issues .....	23
StorNext GUI Known Issues .....	26
StorNext Installation, Replication, HA, and Other Known Issues .....	29
Contacting Quantum .....	32
Getting More Information or Help .....	32

---

# What's New in StorNext 5 release 5.4.0.1

## Purpose of this Release

StorNext 5 release 5.4.0.1 tightens the integration of applications hosted on Quantum Xcellis Workflow Directors, lowers management complexity and costs, and supports the latest available versions of key operating environments. The new **FlexStor** licensing options allow storage administrators to use Storage

Manager as a service, to better manage data residing in a supported public or private cloud more easily – greatly simplifying how cloud storage resources are used. Applications hosted on Quantum Xcellis solutions bring computer resources closer to vital applications, eliminating the need for additional compute servers and network connections. For additional details, see [New Features and Enhancements in StorNext 5 release 5.4.0.1 on the next page](#). StorNext 5 release 5.4.0.1 also resolves several fixed issues listed within the section [Fixed Issues and Enhancements Addressed in StorNext 5 release 5.4.0.1 on page 6](#).

**i Note:** The Stornext Application Programming Interface (SNAPI) will no longer be supported effective with the next major release of StorNext software. Quantum uses SNAPI to integrate legacy libraries, such as the Scalar i500 and Scalar i40/i80 models with StorNext software. StorNext Web Services is included at no additional charge with StorNext, and provides all of the functionality services by SNAPI. For additional information, see [StorNext Web Services on page 19](#).

**i Note:** Advances in disk arrays have made the 32-bit Volume Table of Contents (VTOC) disk labels obsolete. Modern disk arrays use Extensible Firmware Interface (EFI) disk labels capable of supporting physical disks that are larger than 2 TB in size. Beginning with StorNext 5 release 5.4.0.1, VTOC labels are not supported. Prior to upgrading to StorNext 5 release 5.4.0.1, convert arrays that utilize VTOC labels to the EFI format.

To convert arrays that utilize VTOC labels to the EFI format, use the CLI command `cvlabel1`. For additional information about `cvlabel1`, refer to the `cvlabel1` man page in the *StorNext 5 Man Pages Reference Guide* available online at <http://www.quantum.com/sn5docs>.

**i Note:** Beginning with StorNext 5 release 5.4.0.1, the *StorNext and StorNext FX Xsan Compatibility Guide* (part number 6-67068-xx) is deprecated and is no longer maintained. The matrix tables within the *StorNext and StorNext FX Xsan Compatibility Guide* document are now being maintained in the section titled **Xsan Compatibility** within the *StorNext 5 Compatibility Guide* document available on-line at <http://www.quantum.com/sn5docs>.

## Corrections to the StorNext 5 Compatibility Guide

Quantum recommends you download the latest revision of the *StorNext 5 Compatibility Guide* online at <http://www.quantum.com/sn5docs>.

**i Note:** Your software shipped with an incorrect revision of the *StorNext 5 Compatibility Guide*. The following corrections have been made to the *StorNext 5 Compatibility Guide* corresponding to the file name "6-68043-01\_RevAD\_StorNext\_5\_Compatibility\_Guide.pdf":

Section	Section Title	Page	Description of Correction
6	Supported Operating Systems and Platforms	25	In the <b>Red Hat 7 Update 2</b> table, StorNext does support the operating system for Meta-data Controller ( <b>MDC</b> ), StorNext Storage Manager ( <b>SNSM</b> ), and Distributed Data Mover ( <b>DDM</b> ). The table in revision AD of the <i>StorNext 5 Compatibility Guide</i> does not reflect that support. The correct table is provided below.

Red Hat 7 | Update 2  
 See Footnotes: 7, 8, 9, 10  
 Kernel: 3.10.0-327.EL

	MDC	SNSM	DDM	DLS	SAN	DLC	FX
StorNext 5							
5.0.1							
5.1.0							
5.1.1							
5.2.0.x							
5.2.1							
5.2.2							
5.3.0							
5.3.1.x					✓	✓	✓
5.3.2.x					✓	✓	✓
5.4.0	✓	✓	✓		✓	✓	✓

## New Features and Enhancements in StorNext 5 release 5.4.0.1

### Compatibility and Support

The [StorNext 5 Compatibility Guide](#) provides the basic compatibility for StorNext 5 release 5.4.0.1. The *StorNext 5 Compatibility Guide* includes the StorNext components supported, operating systems and service packs, libraries and drives, browsers, virtual machines, and appliance support. Listed below are just a few of the types of information available to you in the *StorNext 5 Compatibility Guide*.

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

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

## Expanded Object Storage

StorNext Storage Manager (SNSM) now has a new storage destination for copies of managed files, in addition to Tape and Object Storage destination. The new destination is known as Object Storage. StorNext 5 release 5.4.0.1 supports multiple Object Storage providers and media types. The [StorNext 5 User's Guide \(Chapter 12: Object Storage\)](#) provides information on the new storage destination functionality, and information on how to configure the new Object Storage providers and media types. The information is also available in the *StorNext 5 Online Help*.

## New Licensing

*FlexStor for Public Cloud* allows end-users to connect their Amazon Web Services or Microsoft Azure managed cloud service accounts to StorNext Storage Manager. The [StorNext 5 Licensing Guide](#) provides information on the new subscription-based **StorNext Storage Manager Subscription (FlexStor)** license, in addition to the new procedure for importing StorNext licenses.

## Improved Capture State

The collection of logs and snapshots has been automated so that all of the important information needed to quickly triage problems can be gathered and sent to Quantum Support with a single button. The [StorNext 5 User's Guide \(Chapter 8: Service Menu Functions\)](#) section **Capture State** provides information on how to create a log file that captures the current state of your system, and automatically upload the log file to Quantum Support. The information is also available in the *StorNext 5 Online Help*.

## Enhancements to StorNext File Permissions

StorNext 5 release 5.4.0.1 has expanded support for Access Control Lists (ACLs). Depending on the system configuration:

- ACLs are now supported for native StorNext Linux clients.
- ACLs are now supported over NFS when using StorNext appliances as NFS servers.

StorNext 5 release 5.4.0.1 also adds support for Unix permission bits on Windows clients. With these changes, StorNext now offers users a choice in cross-platform security models that support more consistent file system permissions for environments containing diverse StorNext client types. The [StorNext 5 User's Guide \(Appendix F: Security\)](#) section **StorNext Security** provides information on considerations when selecting a security model, access-checking functionality and configuring identity mapping.

## Password Policy Enhancement

Enhancements have been implemented to support the configuration and use of strong passwords within your password policy. The [StorNext 5 User's Guide \(Chapter 7: Tools Menu Functions\)](#) section **User Accounts** provides information on how to modify the password policy. The information is also available in the *StorNext 5 Online Help*.

## New Dynamic Application Environment (DAE) for Xcellis Workflow Director

The [StorNext 5 User's Guide \(Chapter 7: Tools Menu Functions\)](#) section **Dynamic Application Environment** provides information on how to display the current status of your virtual machine environments. The information is also available in the *StorNext 5 Online Help*.

# Fixed Issues and Enhancements Addressed in StorNext 5 release 5.4.0.1

[Table 1 below](#) lists fixed issues and enhancements addressed in StorNext 5 release 5.4.0.1.

**Table 1:** Fixed Issues and Enhancements Addressed in StorNext 5 release 5.4.0.1

Operating System	Change Request Number	Service Request Number	Description
All	52225	3468068, 3588642, 3636644, 3684572	An issue was fixed to inform to call Quantum Support if File System expansion operations failed in the GUI. This fix allows users to perform the appropriate File System expansion steps.
All	58981	3588642, 3636644, 3684572	<p>An issue was fixed where the <b>cvupdatefs</b> utility failed to add stripe group with mis-matched size disks.</p> <p>The <b>cvupdatefs</b> utility can add a stripe group to a file system. When doing so, the utility checks to make sure that all disks are the same size. If some disks are larger than others, the file system can only use the space on each disk equal to the smallest disk in the stripe group. The remaining space on the larger disks is wasted. The utility prints out a warning in this situation, indicating how many bytes of disk storage would be wasted if the operation is allowed to proceed. If the administrator chooses to proceed, the issue resulted in the failure of the operation and the stripe group was not added to the file system. The fixed issue allows the operation to go forward and successfully add the stripe group.</p>

Operating System	Change Request Number	Service Request Number	Description
All	60964	3546218	<p>Enhancements were made to the <code>sncmpare</code> utility to log meaningful error messages for the following two conditions rather than core dumping, or getting into an infinity loop:</p> <ul style="list-style-type: none"> <li>The directory depth of a files ystem exceeds a specified value which is user configurable.</li> <li>A full path name exceeds the maximum length of 1024 characters.</li> </ul>
All	61045	3719554	<p>An issue was fixed where some StorNext perl scripts using <code>/usr/adic/perl/bin/perl</code> where calling <code>/usr/bin/perl</code> and failing if the <b>Data::Dumper</b> Perl module was not installed for the operating system provided Perl package. The impacted scripts now explicitly use the Quantum Perl package which provides the <b>Data::Dumper</b> module.</p>
All	61719	3641476	<p>The <b>fsfileinfo</b> command now ensures that it will not insert database values that interfere with the functionality of the <b>fsclean</b> command.</p>
All	61788	3581062	<p>An issue was fixed to reduce downtime duration when upgrading from releases prior to StorNext 5 by deferring pending RPL rebuild requests.</p>
All	61862	3600240	<p>The mechanism used to launch and track distributed data mover processes had some flaws and weaknesses that could cause failures or unnecessary retry operations and a delay or inability to store or retrieve storage manager files. An issue was fixed to clean-up the internal communication so that the operation is more robust and that if errors do occur, the correct error is reported to the administrator.</p>
All	62196	3700172, 3698002, 3692920, 3703154	<p>An issue was fixed in the policy request code to better utilize all available drives.</p>
All	62201	3636874	<p>An issue was fixed to prevent authenticated users from injecting nefarious SQL commands on certain GUI pages. This fix improves application security.</p>
All	62202	3636874	<p>An issue was fixed to prevent authenticated users from injecting nefarious scripts on certain GUI pages. This fix improves application security.</p>

Operating System	Change Request Number	Service Request Number	Description
All	62203	3636874	An issue was fixed that to allow administrative users to define and set a more restrictive password policy for GUI users. This issue improves application security.
All	62204	3636874	An issue was fixed to prevent authenticated users from injecting nefarious scripts on certain GUI pages. This issue improves application security.
All	62206	3636874	An issue was fixed to prevent authenticated users from injecting nefarious hyperlinks on certain GUI pages. This issue improves application security.
All	62208	3636874	An issue was fixed to prevent GUI error messages from revealing private application information. This fix improves application security.
All	62422	n/a	An issues was fixed to allow an Object Storage delete request to timeout in order to prevent processing hangs.
All	62454	3701100	An issue was fixed where asynchronous web services may not be able to start following an upgrade from pre-StorNext 5 release 5.3.0 due to the <b>wsar</b> agent temporary directory not being created on upgrade.
All	62529	3668020	An issue was fixed to support directory names containing spaces when defining File System Quotas. This fix allows directory names to contain spaces when defining File System Quotas.
All	62560	3644724	An issue was fixed to allow StorNext to ignore Spotlight server requests for inode zero in order to avoid filling the system log with error messages.
All	62561	3667834	Enhancements were made to replace the \$VS_DIR string, if it exists, with the real value in the <code>/usr/adic/MSM/internal/config/config_file_&lt;archivename&gt;</code> files. The enhancements allow MSM to start without requiring manual intervention.
All	62797	3680596	An issue was fixed to more accurately track file modification time to ensure files are stored.
All	62814	3658950	An issue was fixed where incorrect errors were being generated when a process doing a BulkCreate received a signal.



Operating System	Change Request Number	Service Request Number	Description
All	62815	3676584	An issue was fixed to improve error handling for the fsfileinfo Web Service call. This fix allows the fsfileinfo Web Service call to properly report Web Service errors.
All	62850	3677566	The WSAR daemon was incorrectly allowing open files to be passed down to the commands it was starting to fulfill web services requests. When the WSAR was restarted, the existence of one or more of those command programs precluded opening the socket port that WSAR needs for communication across the network. The WSAR program was changed to close open files when starting those commands.
All	62891	3648082, 3574008, 3656940	The FSM process dies unexpectedly with the error <code>stripe/X length Y block/Z - Detected bad metadata contents on disk!</code> ; however, running the command <code>cvfscck</code> detects no inconsistent metadata. The FSM process has been fixed so that the race condition that leads to the problem no longer occurs.
All	62991	3681522 3665732, 3672492,3632854, 3669114, 3670060, 3689402	An issue was fixed to prevent the StorNext file system from incorrectly generating a RAS ticket indicating that the metadump's inode count is incorrect and needs to be rebuilt.
All	63040	271072, 1170694, 1170878, 1590718, 1608112, 1611896, 1614798	Enhancements to StorNext file permissions have been implemented. See <a href="#">Enhancements to StorNext File Permissions on page 5</a> .
.All	63043	271072, 1170694, 1170878, 1590718, 1608112, 1611896, 1614798	Enhancements to StorNext file permissions have been implemented. See <a href="#">Enhancements to StorNext File Permissions on page 5</a> .
All	63059	3680048	An issue was fixed to prevent an FSM restart containing the ASSERT message: <code>open_check() hits ASSERT((op-&gt;open_flags &amp; OPEN_FREED) == 0) for ASR helper inode</code>

Operating System	Change Request Number	Service Request Number	Description
All	63060	3669806, 3658080, 3692128, 3697776, 3701602	An issue was fixed to prevent the StorNext file system FSM from crashing when a client attempts to access a deleted file or directory.
All	63095	3683190, 3682876	<p>An issue was fixed to correct the suggestions in error messages which contain suggestions for re-invoking the <b>cvupdate</b> and <b>cvmkfs</b> commands with additional or changed command line options.</p> <p>The <b>cvupdate</b> and <b>cvmkfs</b> utilities issue certain error messages which contain suggestions for re-invoking these commands with additional or changed command line options. Some of these suggestions were not correct and did not coincide with the man pages.</p>
All	63133	3680048	<p>An issue was fixed to prevent an FSM restart containing the ASSERT message:</p> <pre>open_check() hits ASSERT((op-&gt;open_flags &amp; OPEN_FREED) == 0) for ASR helper inode</pre>
All	63149	3683912	Path names with lengths that exceed 1023 characters fail to store, and are listed in an <b>Admin Alert</b> . All store candidates that have full path names with lengths less than or equal to 1023 characters store correctly. Store candidates that have path names with lengths greater than 1023 characters fail to store, and an <b>Admin Alert</b> listing the failed file paths is sent, so that an Administrator can take corrective action for each file (likely by renaming each file to a path whose length does not exceed the 1023 character maximum allowed).
All	63210	3686758	The <b>fsaddrelation</b> command no longer unmounts the target file system, thereby allowing it to execute faster.
All	63263	n/a	An issue was fixed with the Disk Licensing Status (DLS) report generation to include reports from both nodes of a High Availability (HA) system on Red Hat Enterprise Linux (RHEL) 7 systems.
All	63348	n/a	StorNext now prevents a kernel panic when a process running an unmount receives a signal.

Operating System	Change Request Number	Service Request Number	Description
All	63374	3692674	If the FSM process is restarted for any reason (for example, a fail-over occurs), some clients may fail to reconnect. The StorNext reconnect processing has been re factored to avoid the failure.
All	63382	3692674	Incorrect error handling in the FSM may cause attribute updates to be delayed, resulting in stale file attributes being seen on some clients. The StorNext code path in the FSM process that was returning the wrong error code has been corrected.
All	63534	3524236, 3685640, 3644044, 3658054, 3675606, 3682962, 3696610, 3723504	The <b>fsmedcopy</b> and <b>fsclean</b> commands now perform additional up-front validation to ensure media file count integrity.
All	63721	3710720, 3733018, 3730888, 3732348	An issue was fixed to prevent <b>Supervisory Mode Access Protection (SMAP)</b> from being disabled.  Recent improvements in computer chip technology have allowed SMAP functionality to be enabled by host operating systems. A combination of this newer technology and a recent update to Red Hat Linux exposed an improper use of access to user data by the StorNext driver. This resulted in a Linux kernel panic as soon as the StorNext services were started. A workaround was released, disabling SMAP.
All	63873	3705106, 3705890, 3705760, 3705188, 3705138, 3705300	An issue was fixed to update the StorNext metrics Web Service data format to prevent repetitive warning logging statements. This fix keeps the snapshot <code>catalina.log</code> file from growing too large.
All	63911	3706086	An issue was fixed to prevent an unlikely MDC restart when using affinities, autoaffinities, and running low or out of space for an affinity.
All	64098	3712782	An issue was fixed to support Unicode characters in asynchronous web-service calls. This fix allows filenames containing Unicode characters to display properly when using asynchronous web-service calls.

Operating System	Change Request Number	Service Request Number	Description
All	64106	n/a	The <b>fschfiat</b> command now fails if the <b>-A y</b> option is specified for a file with an associated policy class that does not have Alternate Store Location (ASL) enabled.
All	64205	3716838, 3714266	An issue was fixed where GUI images were not able to be seen on Internet Explorer versions 9, 10, and 11.
All	64527	3725120	Executing the command <b>fsclean -c &lt;classid&gt;</b> cleans files associated with that class on all Object Storage and Storage Disk (SDISK) media even when those media are not directly associated to that class. Tertiary Storage Manager (TSM) used to clean Object Storage and SDISK media for <b>fsclean -c &lt;class&gt;</b> only when the media were directly associated to the class. Those media that contain files stored for all classes were not cleaned. In this release of StorNext, media that contain files for all classes are now cleaned for those files associated with that class.
All	64606	3728212	An issue was fixed to provide a better description of the conditions at the client node and includes the file system name.  A client node with an active file system mount remains connected to the file system service (file system manager process) for that file system. If the connection is interrupted, the client attempts to reconnect. Certain conditions such as deleting the file system would cause this service to be terminated. The existing error message did not communicate the situation clearly to the administrator.
All	64615	3729026	The Alternate Store feature dependency on having a file system with an assigned device key equal to one (1) has been removed.
All	64728	3733640, 3729298	An issue was fixed that resulted in an MDC restart due to handling of a recent MDC restart or client reconnects. This issue was uncovered due to network trouble with frequent client reconnects.

Operating System	Change Request Number	Service Request Number	Description
All	64810	3732456	<p>Alternate Store Location (ASL) configuration and status information was added to the set of information collected by <b>pse_snapshot</b>. The following commands are run to collect that information:</p> <ul style="list-style-type: none"> <li>• <b>fsaltnode -l</b></li> <li>• <b>altstoremod -lpc a</b></li> <li>• <b>altstoremod -mlp</b></li> </ul>
All	64819	3732456	The Alternate Store service is now able to recover and reconnect after a prolonged network outage between the local and remote sites.
All	64826	3729298	An issue was fixed to prevent a very unlikely hang during restart handling in the MDC. This would result in a prolonged restart of over 3 minutes. The panic message would be "PANIC: /usr/cvfs/bin/fsm "OpHangLimitSecs exceeded ...". Then, another restart would need to occur after this hang.
All	64879	3733640, 3729298	An issue was fixed that resulted in an MDC restart due to handling of a recent MDC restart or client reconnects. This issue was uncovered due to network trouble with frequent client reconnects.
All	65008	3733390, 3724758, 3729298	An issue was fixed to remove connections from the poll mask when closing the them. When the handler for connections prior to activation exceeds 512 connections, it discards them, but leaves them in the mask used to poll for socket activity on the main listener thread. This causes that thread to spin on a CPU for the lifetime of the FSM.
All	65126	3741040	When invalidating <i>stranger media</i> copies, the <b>fsrmcopy</b> command now clears the FS_STANGER_MEDIA attribute flag from both the on-disk attributes and the database attribute entries associated with the file.
All	65167	3741040	An issue was fixed where a retrieve from <i>stranger media</i> ensures that all data has been flushed to a disk before returning a successful status.
All	65168	3741040	Beginning with StorNext 5 release 5.4.0.1, the <b>dm_util</b> command has been enhanced to support modification of the FS_STRANGER_MEDIA flag setting in the on-disk attributes of a file.

Operating System	Change Request Number	Service Request Number	Description
Linux	32315	1197476	An <b>Admin Alert</b> is now sent when a policy class hard limit or soft limit is reached. In the past, whenever a policy class reaches its hard limit or soft limit, Tertiary Storage Manager (TSM) would not provide any notification that a limit had been reached. TSM now sends an <b>Admin Alert</b> when either the soft or hard limit is reached.
Linux	62557	n/a	The <b>snnas_usage</b> command creates a report file of active and inactive tertiary storage usage per-share per-user for each file system. For more information about the <b>snnas_usage</b> command, see the <b>snnas_usage(1)</b> section in the <i>StorNext 5 Man Pages Reference Guide</i> available online at <a href="http://www.quantum.com/sn5docs">http://www.quantum.com/sn5docs</a> .
Linux	63062	3682876, 3644738, 3666256	An issue was fixed to display the correct and meaningful StorNext GUI error message.  A check was added to <b>cvupdatefs</b> and <b>cvmkfs</b> such that when a file system was being expanded or created, the label on each new LUN was compared with the label on existing LUNs in all active file systems in the cluster. If a match was found, it was likely that a disk was about to be introduced into a file system that was already present in a running file system. If these commands were issued from the CLI, the error message was displayed correctly to the administrator. However, when these commands were run using the StorNext GUI, the error message was not being returned correctly to the GUI, resulting in a meaningless error displayed by the GUI to the administrator.
Linux	63660	n/a	An issue was fixed to allow support for case insensitive mounts on Linux for Samba.
Linux	64066	n/a	ASL remains enabled when moving an ASL enabled file to a non-ASL-enabled directory
Linux	65050	3733306	Enhancements have been made to error messages generated when a few specific metadata consistency problems are detected.

Operating System	Change Request Number	Service Request Number	Description
Windows	62216	3656984	<p>An issue was fixed with the StorNext Windows driver to use the correct calls to the Windows operating system which allows access to files that are created with special characters outside of the standard set of ASCII characters.</p> <p>Special characters are characters outside of the standard set of ASCII characters and are created to allow file names and paths to appear in native languages. Some characters worked, such as German vowels, but others, such as Chinese logograms, would fail to be manipulated correctly. The StorNext file system supports these characters in file names and paths. However, the StorNext file system implementation on Windows used an older version of the interaction with the operating system which resulted in the inability to open or access files containing these characters.</p>

---

## 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 and Partial File Retrieval information is provided in separate documents.

## Quantum Operating System 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 upgrading to an update level or service patch beyond the currently supported levels shown in the *StorNext 5 Compatibility Guide* available online at <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 online at <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.

## Distributed Data Mover (DDM) Guidelines

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



---

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

## Considerations When Upgrading NFS Server Nodes to StorNext 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.

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 StorNext 4.7.x 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 speed and memory availability.

---

**Note:** The database schema update conversion time from StorNext 4.7.x to StorNext 5 is significantly faster than that from StorNext 4.3.x to StorNext 5.

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>;"
```

- 
- i Note:** The query in **Step 2** might require a significant amount of time. Quantum recommends you execute the query before the day of an upgrade.

---

# Compatibility Between StorNext and Other Products

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

## NAS

To view supported StorNext NAS software configurations, review the topic **Supported StorNext NAS Software Configurations** in the *StorNext NAS Documentation Center* available online at <http://www.quantum.com/snnasdocs>.

## 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 5.4.0.1. Object Storage documentation is available online at <http://www.quantum.com/lattusdocs>.

## Partial File Retrieval

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

- 
- i Note:** For Quantum Cloud Storage, PFR is not supported for copies with client-side encryption or compression. It is only supported for copies with server-side encryption or without encryption and compression.

For information about compatibility between PFR and StorNext 5.4.0.1, see the *StorNext Partial File Retrieval Compatibility Guide* available online at <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>.

## Apple Xsan

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

For information about compatibility between Apple Xsan and StorNext 5.4.0.1, refer to the *StorNext 5 Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

## Supported Browsers

For information on browsers supported with the StorNext GUI for this release, refer to the *StorNext 5 Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

---

# General Considerations

This section provides information about items to consider for StorNext 5 release 5.4.0.1.

## Checksum Performance Considerations

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

Current StorNext metadata controller and Mover hardware is able to calculate MD5 checksums at around 300 MB/s to 500 MB/s. For newer generation tape technology, the maximum throughput 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

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

- For Xcellis, see <http://www.quantum.com/xcelliswfd docs>.
- For Artico, see <http://www.quantum.com/articodocs>.
- For M660, M440, M330 Metadata Appliance and Pro Foundation, see <http://www.quantum.com/snmdcdocs>.
- For 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 below](#)
- [StorNext Storage Manager Known Issues on page 23](#)
- [StorNext GUI Known Issues on page 26](#)
- [StorNext Installation, Replication, HA, and Other Known Issues on page 29](#)

---

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

## StorNext File System Known Issues

[Table 2 on the next page](#) lists known issues specific to the StorNext File System process.

**Table 2:** StorNext File System Known Issues

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	54834	3505208, 3516356	<p>If a file is being copied to the StorNext file system using Windows Explorer and Windows Explorer crashes before it finishes copying all the data, the file 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>i Note:</b> This problem can also occur with other programs that set EOF beyond the end of data.</p> <p>This problem does not occur if Windows Explorer encounters an error while writing the file; Windows Explorer will delete the partially written file.</p> <p><b>Workaround:</b></p> <p>To prevent this problem from occurring on StorNext, you can use the StorNext "client configuration" application's advanced mount option "Restrict Pre-allocation API" on Window systems and the "protect_alloc=yes" mount option on Linux systems. This option will set the unwritten parts of the file to zero. When this option is set, non-root users are unable to use the preallocation ioctl. This option also implies sparse=yes.</p> <p>For more information on this option, see the man page <code>mount_cvfs(8)</code>. The sparse option will introduce some overhead when using Windows Explorer. Before setting the <code>protect_alloc</code> option, see the sparse option in <code>mount_cvfs(8)</code> for a description of how it changes StorNext behavior.</p>
All	64963	n/a	See <a href="#">Conditions for Mounting an Xsan Volume as a StorNext Distributed LAN Client (DLC) below</a> .

## Conditions for Mounting an Xsan Volume as a StorNext Distributed LAN Client (DLC)

Beginning with Apple's OS X El Capitan (10.11), support is provided to mount Xsan volumes as StorNext DLC clients. This section provides additional information which should clarify the conditions necessary for this mount to succeed.

When the StorNext services are loaded, the environment is scanned to determine the default mount type - storage area network (SAN) with access to data disks, or DLC with proxy access to data disks through a StorNext gateway node.

If a fibre channel card is detected or internet Small Computer Systems Interface (iSCSI) logical unit numbers (LUNs) are present, the default mode is SAN. If neither of these conditions are true, the default mode is DLC.

In the case that a SAN mount is attempted and not all data LUNs for the file system are present, the mount fails and no error message is logged. If a DLC mount is desired, there are several ways to accomplish this.

First, if the presence of iSCSI LUNs is causing the SAN mount attempt, and these LUNs are not needed, the iSCSI targets can be disconnected. If this is done, the StorNext service must be unloaded and loaded to detect the change.

Example of unload and load:

```
sh-3.2# launchctl unload /System/Library/LaunchDaemons/com.apple.xsan.plist
sh-3.2# launchctl load -w /System/Library/LaunchDaemons/com.apple.xsan.plist
```

If iSCSI LUNs or a fibre channel card are present, the DLC mount is accomplished by indicating this to the driver. There are two ways to achieve this:

- [Set the Xsan Preferences Payload "preferDLC" for the Volume that is to be DLC Mounted below](#)
- [Set the diskproxy=client Mount Options in the /etc/fstab File below](#)

### Set the Xsan Preferences Payload "preferDLC" for the Volume that is to be DLC Mounted

Example of setting preferDLC payload and verifying:

```
sh-3.2# defaults read /Library/Preferences/com.apple.xsan
{
    enableSpotlightServer = 1;
}
sh-3.2# defaults write /Library/Preferences/com.apple.xsan preferDLC '(snfs1)'
sh-3.2# defaults read /Library/Preferences/com.apple.xsan
{
    enableSpotlightServer = 1;
    preferDLC = (
        snfs1
    );
}
```

For more information on the preferDLC payload and Xsan configuration profiles, see <https://support.apple.com/en-us/HT205333>.

### Set the diskproxy=client Mount Options in the /etc/fstab File

An alternative way to indicate to the driver that a DLC mount is desired to set the diskproxy=client mount options in the /etc/fstab file.

Example of the diskproxy=client mount options in /etc/fstab file:

```
sh-3.2# cat /etc/fstab
LABEL=snfs1    none    acfs    rw,diskproxy=client
sh-3.2#
```

In either of the above scenarios, a driver unload and load must be performed to enact the change. Once the Volume is mounted, its status as a DLC client can be verified by using the sub-command **who** of the **cvadmin** CLI. First select the file system, then enter **who**. Look for the **C** after the **CLI** in the client entry.

Example of the `cvadmin who` sub-command:

```
Xsanadmin (snfs1) > select snfs1
Xsanadmin (snfs1) > who
Who (File System "snfs1")

#      acfs  I.D.  Type  Location                Up Time                License Expires
--      -
0>      50569B3B7F  FSM  smw-rhel64.mdh.quantum.com 4d 23h 33m  N/A
4>      50569B3B7F  CLI S smw-rhel64.mdh.quantum.com 4d 23h 33m  *10-27-2016
23:59:59
9>      ADM  smw-elcapitan.mdh.quantum.com 0d 0h 0m  N/A
17>     CLI C smw-elcapitan.mdh.quantum.com 0d 0h 6m  N/A

1 active connection out of 65535 licensed connections
0 active connections out of 65535 licensed proxy connections
1 active gateway connection
1 active client-licensed connection

Xsanadmin (snfs1) >
```

## StorNext Storage Manager Known Issues

[Table 3 on the next page](#) lists known issues specific to StorNext Storage Manager.

**Table 3:** StorNext Storage Manager Known Issues

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	43320	1581004	<p>File retrieves from media to disk can be suboptimal for fast tape drives like the Oracle STK T10K drives. This scenario can occur when the retrieve event is initiated on a host that is different from the host running the mover process, which requires the use of synchronous direct I/O.</p> <p><b>Workaround:</b></p> <p>To work around this issue and achieve optimal performance for both file stores and retrieves with the T10K drives, increase the default I/O size used by the mover process and make the mover process use asynchronous buffered I/O when the use of synchronous direct I/O is not required, using the following steps:</p> <ul style="list-style-type: none"> <li><b>i Note:</b> This workaround may also help improve the performance of the faster LTO drives like LTO-6, and LTO-7 by updating the <b>FS_LTO_BLOCK_FACTOR sysparm</b>.</li> <li><b>i Note:</b> Changes to <b>FS_XXX_BLOCK_FACTOR</b> only affects tapes formatted after the change.</li> </ul> <ol style="list-style-type: none"> <li>1. Change the <b>FS_T10K_BLOCK_FACTOR 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> <ul style="list-style-type: none"> <li><b>i Note:</b> The T10K default I/O block size is 512 KB or 8 * 64 KB. With the block factor changed to 32, the new T10K I/O block size will be 2 MB or 32 * 64 KB. Presently, the <b>FS_T10K_BLOCK_FACTOR sysparm</b> must not be set to a value that exceeds 32.</li> </ul> </li> <li>2. Restart Storage Manager to ensure the change in <b>Step 1</b> goes into effect:  <pre># tsmstop # tsmstart</pre> </li> <li>3. Verify the <b>FS_T10K_BLOCK_FACTOR sysparm</b> contains the new value:  <pre># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> </li> <li>4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients.</li> <li>5. Modify <code>/etc/fstab</code> on the MDCs and the DDM clients to use the <code>auto_dma_write_length</code> and <code>auto_dma_read_length</code> mount options as follows:  <pre>snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0</pre> </li> </ol>



Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>0</p> <ol style="list-style-type: none"> <li>6. Unmount and re-mount your file systems.</li> <li>7. Use new T10K media to store a copy of the file from the disk.</li> </ol> <p><b>i Note:</b> <b>Step 7</b> is very important; when the new copy is made to the new tapes, the new tapes are labeled with a 2 MB block size, which is used for subsequent writes or reads to and from the media. Tapes on which <b>fsformat</b> was run before the change will use the block factor in use at that time. This change will not impact those tapes.</p>
All	46693	n/a	<p>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> <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>When copying files between media using the CLI command <b>fsmedcopy</b>, the file is not re-segmented to match the segment size of the destination media. Rather, the original segments are copied to the target media type and the distribution of segments across destination media will, therefore, be the same as the distribution on the source media.</p> <p><b>i Note:</b> This behavior may cause file data segment distribution to be sub-optimal on the destination media.</p> <p><b>Workaround:</b></p> <p>Currently, a workaround does not exist for this Known Issue.</p>
All	65077	n/a	<p>During StorNext Storage Manager start-up, SCSI Reservation conflicts may be reported by the IBM APFO driver if installed, particularly after involuntary fail-over of the StorNext MDC in High Availability configurations.</p> <p><b>Workaround:</b></p> <p>These reservation conflicts can generally be ignored and cause no issues with StorNext Storage Manager.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	65364	n/a	<p>When Storage Manager stores files to LTFS media, the fs_fmover process for TSM generates warning messages in the trace_05 files (found in the /usr/adic/TSM/log/trace directory) indicating that the fs_fmover process is unable to get the extended attributes on new files on the LTFS formatted media. The following is an example of the warning message:</p> <pre> /usr/adic/TSM/logs/trace/trace_05 Oct 25 13:31:49.151125 northern sntsm fs_fmover[22366]: E1201 (8)&lt;02507&gt;:mdt2l1tfs1605: {1}: mdt2l1tfs_file_get_attr: failed to get attr user.QUANTUM_STORNEXT_PATH on file /.StorNext/objects/00/00/01/0000000000000172: -1040                     </pre> <p><b>Workaround:</b></p> <p>These warning messages are safe to ignore for newly created files on the LTFS media. There may be up to four messages printed per LTFS file, one for each of the StorNext extended attributes on the LTFS media:</p> <ul style="list-style-type: none"> <li>• QUANTUM_STORNEXT_PATH</li> <li>• QUANTUM_STORNEXT_FKEY</li> <li>• QUANTUM_STORNEXT_VERSION</li> <li>• QUANTUM_STORNEXT_META_VERSION</li> </ul>

## StorNext GUI Known Issues

[Table 4 on the next page](#) lists known issues specific to the StorNext GUI.

**Table 4:** StorNext GUI Known Issues

Operating System	Change Request Number	Service Request Number	Description/Workaround
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> <p><b>Workaround:</b></p> <p>There is currently no workaround for <b>Change Request Number 57856</b>.</p> <p><b>i 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>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	65328	n/a	<p>When the MDC is acting as a Distributed LAN gateway for a file system and its configuration is modified using the StorNext GUI or the Connect GUI, a spurious LAN server configuration file named <code>/usr/cvfs/config/dpserver.&lt;fsname&gt;</code> may be created if it does not already exist. This file will have all network interfaces enabled and overrides the main <code>dpserver</code> file, <code>/usr/cvfs/config/dpserver</code>. In some cases, having all network interfaces enabled is undesirable as it may result in sub-optimal DLC performance and cause DLC traffic to interfere with other MDC network traffic.</p> <p><b>Workaround:</b></p> <p>Complete all configurations of the file system prior to setting <code>diskproxy=server</code> in the <code>/etc/fstab</code> file.</p> <p>Or, if the file system configuration requires modification after <code>diskproxy=server</code> has been set in the <code>/etc/fstab</code> file perform the following steps:</p> <ol style="list-style-type: none"> <li>1. Execute the command <code>sndpscfcg -e</code> to verify that <code>/usr/cvfs/config/dpserver</code> exists and contains correct information: <pre># /usr/cvfs/bin/sndpscfcg -e</pre> </li> <li>2. Make <code>/usr/cvfs/config/dpserver.&lt;fsname&gt;</code> be a symbolic link to <code>/usr/cvfs/config/dpserver</code>: <pre># /bin/rm -f /usr/cvfs/config/dpserver.&lt;fsname&gt; # ln -s /usr/cvfs/config/dpserver /usr/cvfs/config/dpserver.&lt;fsname&gt;</pre> </li> <li>3. <b>(Optional)</b> Repeat <b>Step 1</b> through <b>Step 3</b> for all file systems for which the MDC is acting as a gateway and where <code>/usr/cvfs/config/dpserver</code> is used.</li> </ol>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	47954	n/a	<p>The Safari browser becomes unresponsive when you attempt to configure an Email server using the StorNext GUI.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, perform the following procedure:</p> <ol style="list-style-type: none"> <li>1. Shut down the Safari browser window(s).</li> <li>2. Restart the Safari browser, and then retry 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>

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

[Table 5 on the next page](#) lists known issues specific to StorNext installations, data replication, HA systems, and other areas.

**Table 5:** StorNext Installation, Replication, HA and Other Known Issues

Operating System	Change Request Number	Service Request Number	Description/Workaround
------------------	-----------------------	------------------------	------------------------

All	47041	n/a	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 from StorNext 5.0.0 or earlier.
-----	-------	-----	--

**Note:** The `classndxatimeme` index will already be present and not be added again if upgrading from StorNext 5.0.1 or later.

**Workaround:**

To minimize TSM downtime after upgrade, the `classndxatimeme` index can be created prior to performing the upgrade using the `index_tierfiles.pl` PERL script. This file is available in the StorNext installation directory by opening a support ticket and requesting the file. (Quantum service and service partners can obtain this file from the StorNext Metadata Appliances page on CSWeb.) 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.

To manually add the index, you must have the `index_tierfiles.pl` script. Then do the following:

Log in to the primary server node, and access the command line of the system:

1. Log in to the primary server node.
2. Enter the following to 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
```

Operating System	Change Request Number	Service Request Number	Description/Workaround
------------------	-----------------------	------------------------	------------------------

5. Execute the PERL script:

```
./TSM/index_tierfiles.pl
```

The procedure is complete.

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> <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>
All	58814	n/a	<p>When running replication from a deduplicated file system to another deduplicated file system, corruption may occur 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.</p> <p>These areas read as zeroes rather than the appropriate data.</p> <p><b>Workaround:</b></p> <p>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.</p>

---

# Contacting Quantum

## Contacts

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

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

## Getting More Information or Help

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

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

For worldwide support:

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

## Comments

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

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