

StorNext 7.0.1 Release Notes

Contents

What's New in StorNext 7.0.1	2
Supported StorNext Upgrade Paths and Upgrade Considerations	18
Compatibility Between StorNext and Other Products	19
General Considerations	20
Upgrading Appliances	21
Appliance Release Notes	21
Known Issues	21
Contacting Quantum	36

© 2020 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law. ActiveScale, DXi, DXi Accent, FlexSync, FlexTier, iLayer, Lattus, Quantum, The Quantum Logo, QXS, Scalar, StorNext, SuperLoader, Vision, 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.

What's New in StorNext 7.0.1

Purpose of this Release

The StorNext 7.0.1 release provides new features and enhancements listed in the section [New Features and Enhancements in StorNext 7.0.1 below](#), and also provides software fixes listed in the section [Fixed Issues and Enhancements Addressed in StorNext 7.0.1 on page 7](#).

New Features and Enhancements in StorNext 7.0.1

New Package Dependencies

Beginning with StorNext 7.0.1, you must install the following new/additional packages before you install StorNext.

- java-1.8.0-openjdk-1.8.0.121 or later (for the StorNext GUI)
- python-requests, pyxattr, python-dateutil (for the StorNext Primary File System Pooling feature)

See the [StorNext 7.0.1 Compatibility Guide](#) for more information about operating systems, kernel versions, rpm package dependencies, and hardware platforms supported by StorNext.

Support for Primary File System Pooling

StorNext has always supported mixing different types of storage into a file system by using multiple stripe groups. Beginning with StorNext 7.0.1, the primary file system pooling feature introduces a set of mechanisms that allows you to explicitly control the placement and the movement of content between different classes of storage. The aim is to enable you to use workflows where the content has different bandwidth needs at different points in its life cycle and to make it simple for you to relocate content based on those needs.

Use the StorNext User Interface to configure Primary File System Pooling (see [StorNext File System Pooling](#)).

Use Cases

- **Non-linear video editing, especially at higher resolutions can require very significant bandwidth and can also benefit from being done on solid-state storage, which is not subject to seek latencies.** However, the cost of this storage is still significantly higher than spinning disk and for much of its lifespan, content does not need the storage characteristics of solid-state devices. By making it simpler to move specific content between different types of storage, StorNext supports more demanding workflows without the cost of running the storage everywhere.
- **Another potential use is to isolate project content from in-progress work.** By dedicating a pool of storage to a project, you can avoid contention with other projects during times of activity and avoid unexpected contention on devices while still maintaining shared access to the content.

- **High performance ingest where data cannot be dropped.** Migrate content out of a higher performance storage using an age-out policy to make room for new content.

Concepts

Pooling moves data; it does not change any of the user metadata associated with files, nor does it move files in the namespace, or change their attributes when it relocates content. Content is moved between *storage pools* using either pooling *jobs* initiated by a user or administrator, or by a pooling *policy* configured by an administrator and executed automatically by the system. All processing is handled by the pooling daemon in the background.

Storage Pools

A storage pool is a set of stripe groups with an associated name that represent the targets for data movement. Pools are generally associated with stripe groups using specific types of storage; for example, a fast pool for SSD or NVMe and a slow pool for spinning disk. However, there is nothing in the product which requires this be so. You can configure several pools out of the same type of storage.

A storage pool can be exclusive or non-exclusive. You can only use an exclusive pool by content specifically targeted to that pool, and you can only use a non-exclusive pool by both targeted and untargeted content.

You can place both files and directories in a pool. A file in a pool means its contents are placed on the devices in the pool. A directory in a pool means that new file content within the directory will also use the pool. The only limit on the number of pools is the number of data stripe groups in a file system.

Pooling Jobs

A pooling job consists of instructions for the system to perform an action on specified content; the most common job is to move the data content to a specific storage pool. If you are granted access to the pooling system, you can initiate a job, or the job can also be triggered by a pooling policy.

Jobs have different controls to modify its behavior; details are described in a different section. A job can run immediately or at a specified time in the future. When a job is finished, a completion report is held in an internal database allowing the report to be emailed or viewed later.

⚠ Caution: Setting the **iosize** for the job (either directly or through the policy) too big relative to the total memory on the node might have an adverse effect on the system. Each IO thread (default 2) in the Pooling daemon allocates number of buffers (default 4) of the iosize size.

Pooling Policies

A pooling policy consists of a request to run a job on a regular basis looking for work to do and executing the work.

Both jobs and policies define the content to be acted on using the same mechanisms, a set of content can either be explicitly referenced and filtered down by a set of rules, or you can use an internal query mechanism of the file system to look up candidates for processing.

⚠ Caution: Setting **iosize** for the job (either directly or through the policy) too big relative to the total memory on the node might have an adverse effect on the system. Each IO thread (default 2) in the Pooling daemon allocates number of buffers (default 4) of the iosize size.

Pooling Service

All the execution of work for pooling is handled by a daemon; jobs are submitted to the daemon and policies are stored and run by the daemon. You can use the command line to access the service using a REST API.

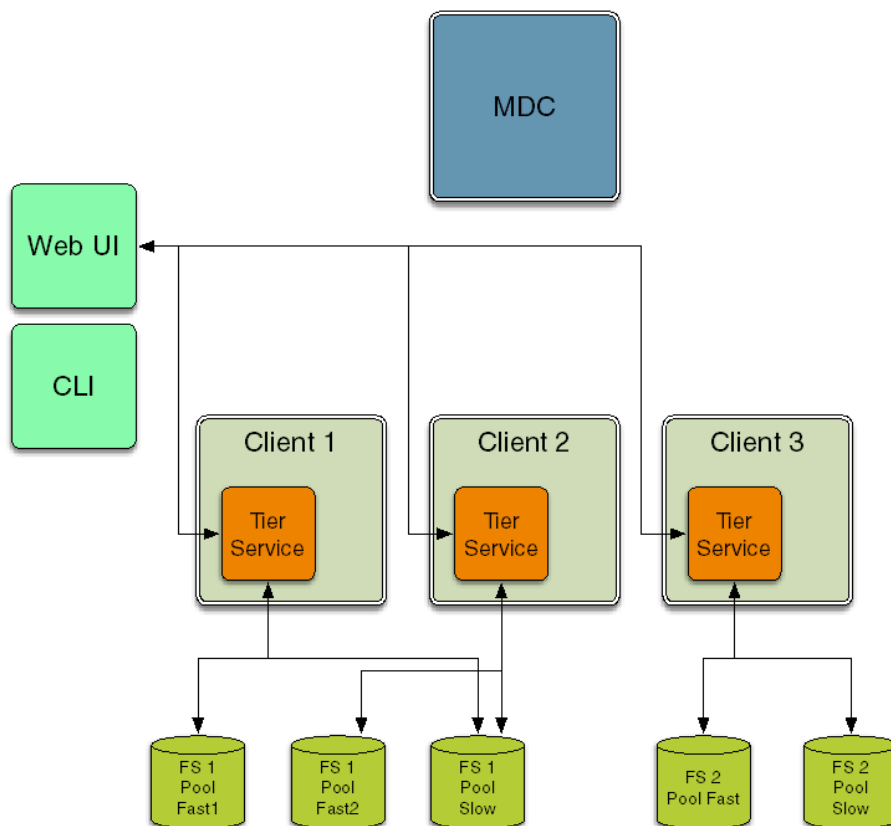
You must run the daemon, **sntierd**, on a Linux-based StorNext client running either the RedHat 7 (or later) operating system or the CentOS 7 (or later) operating system. The system can be a customer supplied system, an MDC, or an Xcellis Workflow Extender.

Note: You must have sufficient storage bandwidth from the system to support the workload.

Multiple instances of the service may be run at the same time. The services are all independent systems and are not aware of each other. Each service can support running policies and jobs, and is capable of managing content in all the StorNext file systems configured with storage pools and mounted on the client host.

Each pooling service maintains a persistent state in a small internal database, which contains configuration information, jobs, and policies. If you restart the service while a job is running, those jobs are re-run after the service is restarted.

Each job runs on a single pooling service. Illustrated below, the CLI and Web UI can access the pooling service on multiple hosts; there are three pooling services, two of the services operate on FS1 and one service operates on FS2.



Support for Object Storage Rate Limiting

Beginning with StorNext 7.0.1, you can throttle the object storage maximum upload and download rates to limit the network bandwidth consumed by Storage Manager. You can define the maximum send and receive rates through system parameters, or set them by command or schedule. Storage Manager will limit the per connection object storage transfer rates, such that the cumulative rates do not exceed the configured values. The rate limits will be applied globally; that is, they will limit the total bandwidth used across all Storage Manager hosts.

Use the StorNext GUI to enable rate limiting. See the **Object Storage Rate Limit** section in **Tools Menu Functions** to manually set the rate limits. See the **Scheduler** section in **Storage Manager Tasks** to schedule a task to set the rates.

Use the command line interface (CLI) to run the **fsobjratelimit** command to set object storage upload and download rate limits. Alternatively use the CLI to run the **fsschedule** command to schedule a task to set the rates. See the **fsobjratelimit** and **fsschedule** commands in the [StorNext 7 Man Pages Reference Guide](#).

Object ID to Path Mappings

Beginning with StorNext 7.0.1, you can use the **fsobjlist** CLI command to generate a report of object ID to path mappings. By default, this report lists the objects for all configured media. Alternatively, options are available for filtering by media, namespace or class.

Use the command line interface (CLI) to run the **fsobjlist** command to generate the report. See the **fsobjlist** command in the [StorNext 7 Man Pages Reference Guide](#).

Locate Files Based on Object IDs

Beginning with StorNext 7.0.1, you can use the **fsobjlocate** CLI command to display the file path for each object ID specified. See the **fsobjlocate** command in the [StorNext 7 Man Pages Reference Guide](#).

Retrieves from Azure Archive Storage

Beginning with StorNext 7.0.1, files in Azure Archive Storage are automatically rehydrated as part of the retrieve processing. By default, the files are rehydrated to the Azure hot pool before the data is downloaded. You can specify the option **-O cool** with the **fsretrieve** command to restore the data to the Azure cool pool. See the **fsretrieve** command in the [StorNext 7 Man Pages Reference Guide](#).

Pathname Interface to the fsazure Command

Beginning with StorNext 7.0.1, when you use the **fsazure** command to change or report on the pool state of an archived file, you can use the file path to specify the file. See the **fsazure** command in the [StorNext 7 Man Pages Reference Guide](#).

Enhancements to the Proxypath Mount Option

Beginning with StorNext 7.0.1, the default proxypath mount option for Linux and Windows is **balance**. For StorNext releases prior to StorNext 7.0, the default is **filestickybalance**.

- For Windows clients, you can configure the proxypath option using the Client Configuration utility (see [Configure the StorNext Client Software on Windows](#)).
- For Linux, macOS, and Solaris clients, see the **proxypath** mount option for the **mount_cvfs** command in the [StorNext 7 Man Pages Reference Guide](#).

Changes to the Default Setting for the Buffer Cache Cap and the Cache Buffer Size File System Mount Options

Beginning with StorNext 7.0.1, the default setting for the buffer cache cap (**buffercachecap**) varies by the amount of physical memory detected on your system, as illustrated below.

Amount of Physical Memory Detected	Default Buffer Cache Cap
Less than 2 GB	256 MB
Greater than or equal to 2 GB and less than 4 GB	512 MB
Greater than or equal to 4 GB and less than 8 GB	1 GB
Greater than or equal to 8 GB	2 GB

The specific cache buffer size (**cachebufsize**) defaults to the following.

Total Buffer Cache Size	Default Cache Buffer Size
Less than or equal to 1 GB	256 K
Greater than 1 GB	1 MB

For more information, see both the **buffercachecap** mount option and the **cachebufsize** mount option for the **mount_cvfs** command in the [StorNext 7 Man Pages Reference Guide](#).

Compatibility and Support

The [StorNext 7.0.1 Compatibility Guide](#) provides the basic compatibility for StorNext 7.0.1, including 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 7.0.1 Compatibility Guide](#).

- **Upgrade Paths:** Provides information on what upgrades to this release are supported.
- **Appliance Support:** Provides information on what StorNext and Lattus appliances are supported with this release or are compatible with it.

- **Operating Systems and Platforms:** Provides information on what StorNext components run on various operating systems and service packs. Also includes which operating systems have been newly added or removed.
- **Client Interoperability:** Provides information on what StorNext clients running other versions of StorNext are compatible with metadata-controllers (MDCs) running this release.
- **Virtual Machine Support:** Provides information on what StorNext components running on selected operating systems and service packs are supported in virtual machines.
- **Compatibility with Other Products:** Provides information on references to additional StorNext sold-separately products that are supported with this release.
- **Browser Support:** Provides information on what versions of browsers are supported with the GUI in this release.
- **Drives and Libraries:** Provides information on what Quantum and 3rd party drives and libraries are supported with this release.

Fixed Issues and Enhancements Addressed in StorNext 7.0.1

Operating System	Change Request Number	Service Request Number	Description
All	36586	n/a	fsmedcopy performance is slow
All	40102	1482282	/etc/init.d/cvfs does not have the "status" option
All	46754	3353006, 521374	vsexport says it worked OK removing a volume ID but it did not remove it
All	49618	1482282	Linux service cvfs status does not work
All	53423	n/a	Cloud Traffic Throttling
All	54261	3499380, 528772	cvfsdb rpl command should report if the inode is free
All	66175	313551, 526472	Default size of MED_SEG_OVER_* parameters for some media types need to be increased
All	67986	337136, 404325, 467638	Having too many qustat files causes pse_snapshot to miss collecting cvgather bundle
All	68864	344944	Add Web Service calls to dump a list of files in a specified file system

Operating System	Change Request Number	Service Request Number	Description
All	69792	451886	Add in logic for LTO-8 for STK SL8500 Library
All	70940	396771	snretrieve does not work in async mode for directories
All	72660	451734, 457348, 482143, 486984, 489143, 501124, 578649	mtime in the storecand table does not get updated by rebuild policy
All	72801	n/a	snaudit should optionally display user/group names instead of RAW UIDs/GIDs
All	72884	538729, 526991	Multi Threading Object PUT and Get.
All	72885	454483	Windows clients have full access to StorNext specific directories.
All	73130	456401	Derby database corruption caused both HA nodes to SMITH when services were started
All	73368	n/a	Support for Google multi-part uploads
All	73576	n/a	NAS beyond 2.2.0 and SNFS work -- Windows Mandatory Locks need to work between SMB clients and native SNFS Windows clients
All	73593	464568	cvfsck -q provides misleading results for the amount of needed memory
All	73634	n/a	sgadd should perform UNMAP for any added stripe groups containing thin-provisioned LUNs
All	74091	527663, 481993, 456474, 491196, 498960, 500798, 533125, 540767	fsm panic w/ ASSERT failed "!set (rptr->iref_flags & IREF_RETURN_ATTRS) == 0" on unmanaged filesystem
All	74239	n/a	New HealthCheck to list failed media for Baidu
All	74251	488616	When SNFS SAN client releases lock, NFS3 waiter takes long time (30 secs) to get lock
All	74278	562153	FSM should error if metadataArchiveSearch = false and storageManager = true
All	74309	538729	Delete object error after copy from tape to object storage fails

Operating System	Change Request Number	Service Request Number	Description
All	74545	501127, 493912, 501127	Ubuntu Kernel 4.4.0.145 new get_user_pages () signature, cvfsbuild fails error: too many arguments to function 'get_user_pages'
All	74599	n/a	Items added to sl_email_notification table are never removed
All	74604	529382	Incorporate LTFS Performance enhancements into Storage Manager
All	74699	n/a	Provide optional functionality in fspolicy for only sending complete clusters
All	74784	500752	Need better error messaging in fsobjcfg
All	75113	490242	Unable to retrieve file using xreps (Lattus object not found) and file state is incorrect
All	75143	525455	ENOENT returned for lookups on linux 7.6 with storage manager
All	75175	n/a	TSM checkMediaAvailabilityTsm health check script must detect non-tape store availability
All	75213	526974, 555991, 566835, 570951, 570482	Misleading message on MDC w.r.t san_client DOWN disk devices
All	75270	n/a	setfattr for special attribute names not always working as expected
All	75307	511169	Support for CentOS/RHEL 7.6 - MDC, SNMS, DDM and DLS
All	75393	517863, 438714, 439502	install.stornext should not allow to select a directory on a SNFS as install dir for wsar_agent
All	75403	n/a	GUI: Cloud Traffic Throttling
All	75411	535900	P100 media remained 100% used after fsclean removed objects
All	75455	517831, 522496	Need notification or more retries before marking Object Storage media as write-protected

Operating System	Change Request Number	Service Request Number	Description
All	75468	512218, 569557	fsimport fails to import media with the hardware write lock set
All	75469	518684	sgoffload fails due to autoAffinity with Invalid argument - Error comparing config files
All	75613	543550	Large segmented files retrieve very slowly
All	75636	514500	pse_snapshot gets stuck for a long time when looking for the "filelist" files
All	75637	514500	snbackup gets stuck for a long time when looking for the "filelist" files
All	75641	495735	Tracking bug: Problem with SM using lin_tape driver during 100 million file store test
All	75684	396771	Allow more than 1 recursive retrieve to work at same time
All	75688	517608	amend REST API to get media ID, in order for flexsync task report to list media by ID not by media UUID
All	75723	520733	snbackup script does not check if getpwuid (/)getgrgid() calls return valid data
All	75735	492515	snprobe not detecting secondary node as an mdc
All	75770	513514	configuration settings are lost when unmapping a file system in Windows Client Configuration Tool
All	75773	n/a	Merge code from 6.3.0LFR for TS1155 & TS1160 drives to trunk
All	75826	496940	Windows rename not working if the Parent Dir has not the ACL FILE_WRITE_ATTRIBUTESpermission
All	75889	529772	renametracking enabled causing issues with office files

Operating System	Change Request Number	Service Request Number	Description
All	75897	524283	GUI on RYO installation does not show the serialnumber from the license.dat when cvfsid has leading zeros
All	75904	530604	snrecover caused FSM segmentation fault because lost+found directory not found
All	75939	527400	Tape drives taken offline due to excessive source file read errors during store operations
All	75952	525455	nfsd hung waiting for cvfs_access()
All	75976	529382	Tapes are not ejected from the drives when DDM is enabled
All	76209	529757, 530493	CVFS client hung after directory rename on managed file system under Linux 4.15
All	76259	534915	fs_restd fails to check for allocation errors due to oversized COPY_REQ_IPC_T objects
All	76277	526535	Use of a proxy server causes fsobjcfg to timeout
All	76309	531647, 571720	snmsm calls sync_config_files -f <lib> reports a Error: unable to rename config_file_xyz when there are 2 archive with same prefix in the name
All	76335	534439	sgoffload command fails with I/O error to SG being vacated
All	76338	536060, 558426	GUI should warn of consequences if a user tries to set the number of read streams to the value of max streams
All	76341	536060	TSM needs to log a descriptive message instead of just "No controllers found" when no streams are available
All	76347	533735	FlexSync can cause the fsm of the file system its replicating from to use up all its memory
All	76348	n/a	Changing number of backup copies does not affect all files

Operating System	Change Request Number	Service Request Number	Description
All	76352	537146	DDM: file system internally disabled and Admin Alert when fsretrieve from P100
All	76376	528985, 545298	AWS retrieve from Glacier not optimized and can take a very long time if multiple jobs are started
All	76405	570207	cvfs assert f_rwlck->rw_readers > 0 on heavily contested lock
All	76414	531558, 530339, 538946	Data loss due bad error handling with specific store scenario
All	76444	533787, 563472	StorNext Client gets errors when trying to get a directory listing
All	76473	538245	snquota has an incorrect check of the hard and soft limit maximum values
All	76501	540085, 566270	An exception 0xc0000005 has occurred on line 459 of source file S:/00273715_win1064_pkg/snfs/client/vfs/windows/ntif_fastops.c.
All	76543	541186	Installation of Ubuntu 18.04.1 LTS security patches breaks build of cvfs kernel module
All	76561	534146; 546243	SM should not store files with too many segments
All	76563	542031	fsrecover -u -d error message is misleading when trying to recover a file, will help for GUI too
All	76615	560280	Auditing generates a nodechange to every client on a close
All	76694	544346	Need to have later Ubuntu releases available to install on DAE VM's
All	76726	545381	partial retrieval hang with multi-stream enabled with big start offset number
All	76735	537346, 541579	Line of null characters in a WRF_ file spins fs_policyd and fills tac_00 with errors

Operating System	Change Request Number	Service Request Number	Description
All	76783	545249	sgoffload fails with extent_swap error when SG being vacated
All	76792	538729	fsmedcopy terminates on S3 multipart uploads and fails all files in the request
All	76793	538729, 556379	Files marked as failed even though store to S3 bucket was successful
All	76820	547478, 588379	quota limits report 16x greater in size after upgrade
All	76831	543055, 561204	activefl - Database table size grows very large due to high fragmentation
All	76876	547556	mdarchive: sgoffload leaks qtree query transactions causing mdarchive to grow without bound
All	76885	547452	FSM Panic in qtree_cache_reclaim() due to empty clean list
All	76898	548973	Enhancement - need a utility to look up file name based on object ID
All	76916	538729 542765	fsmedcopy terminates on curl errors during S3 multipart uploads and fails remaining files in the request
All	76922	548498	snbackup fsstore and store policy conflicts result in many backup failures
All	76987	546138	metadb rest call files_by_media rest api seems to hang forever
All	76997	543055, 571150, 576036, 576037, 577537, 569740, 578514, 578812, 580923, 578146, 586085, 585472	Long running transaction in UI logic - this MAY result in growing ibdata1 file
All	77077	492515	snprobe not detecting secondary node as an mdc

Operating System	Change Request Number	Service Request Number	Description
All	77115	547071, 559738	fsmpm: PANIC: fsmpm ASSERT failed "'dim2 >= *dim3" in cvfs_square_box(), snfs\fsmlib\disks.c
All	77118	547071	cvpaths "directory" keyword recommended for Windows is producing odd/unexpected results
All	77159	542323, 554123, 548115	fspolicy rebuild metadb query timing out
All	77215	548237	fscclean logs:"rdb1service_del_files_from_mediaMSR_NOT_FOUND therefore fcnt is suspect" on all media which have nothing to delete
All	77217	538729	add sysparms to tune for Object Storage connectivity outages
All	77283	555650	GUI: modify policy "Copy Expiration" warning to be more explicit about file deletion
All	77286	550176, 586656	sgoffload: occasionally gets stuck looping on the same line of output: polling for extents, set xx, yy<= -1 out of yy extents retrieved and doesn't proceed
All	77350	547295	kernel panic w/ exception RIP: strcmp+24 upon RefreshDiskList() invoked upon filesystem expansion
All	77391	557521	ibdata1 file growing very large due to the undo log within
All	77403	557311	OpHangLimit triggered by thread hung in DoClose() with InodeFlagTrimClose set and VOP_RETRY
All	77464	557790	fsm panic: new coherency model and global share mode enabled hit fsm ASSERT failed "(optr->open_flags & OPEN_HAVE_BOTH) == OPEN_HAVE_BOTH" upon InodeOpenUpgrade()
All	77474	558258	On a managed file system fsCapacityThreshold does not trigger a RAS event

Operating System	Change Request Number	Service Request Number	Description
All	77515	560062, 549992	FSM: PANIC: /usr/cvfs/bin/fsm ASSERT failed "IP_IS_SPACE_TREE(ip) this.idiext_frblock + extsize <= next.idiext_frblock"
All	77547	549992	FSM: PANIC: /usr/cvfs/bin/fsm ASSERT failed "IP_IS_SPACE_TREE(ip) this.idiext_frblock + extsize <= next.idiext_frblock"
All	77646	559566, 570899	Is of a directory failing on a client with "Input/output error" while mdc is creating many files in it.
All	77653	563234, 563898	Mysqld fails to start after upgrading to StorNext version 6.4.0
All	77668	557790	Global Share mode with iotokens enabled sometimes displays msg: Sharemode change forced
All	77671	563430	Object Storage: metadata not stored on AWS when "V4 Full-Payload Signing" enabled on bucket
All	77678	563234, 563898	file-per-table conversion needs to account for sys_config.ibd file
All	77679	555576, 562683, 572047	Evidence for memory leak in send_ha_reset_timer_rest_req() causing FSMPM high memory usage
All	77705	554829	fsfilecopy -n generates failures for partially copied multi-segment files
All	77730	564850	Snapshot: Include contents of files in the /usr/adic/TSM/internal/locks directory
All	77741	562683, 572047	Evidence for memory leak in fsm_pm_ha_reset_timer_req_handler() causing FSMPM high memory usage
All	77771	567119	Reduce the processing time of request api files_by_media with small number of truncated files

Operating System	Change Request Number	Service Request Number	Description
All	73496	465444, 577160	ArcDisp hit OOM condition and got killed, leading to MSM abnormal termination
All	76696	543550	Retrieve performance degrades for multi-segment files when movers don't run on the same system
All	76970	546119, 554609, 576026, 578397, 576097, 571413	fsm panic upon ASSERT failed "pclient->cl_rsvd_counted == 0"
All	76971	546118, 551599, 586977	fsmpm segmentation fault due to unhandled null pointer exception in send_delayed_activation()
All	77221	553837, 580930	wsar_agent is running into heap exhaustion due to MariaDB mem_root implementation
All	77402	554829	Abnormally terminated TSM CLI's can leave daemons in a bad state.
All	77746	n/a	Red Hat (RHEL) 7 update 8 Client Support
All	77812	n/a	GUI: Red Hat (RHEL) 7 update 8 Client Support
All	77819	566647	fs_copymand throws SISEGV in TSM::copymand::ReqInfo::logInternals() for canceled request
All	77861	569557	The fsmedscan util is producing extraneous fields that causes fsimport to fail
All	77862	567325	TSM spams tac_00 log with user or group does not exist
All	77873	568177	fsconfig fails silently when adding a drive if a SCSI reservation release fails
All	77879	561495	Disks and Path of mulitpath devices set to standby
All	77916	569657	TSM excludes.store should be updated to exclude FLEXSYNC_TMP_PREFIX(.__flexsync_tmp__) files, by default

Operating System	Change Request Number	Service Request Number	Description
All	77964	567462, 576880, 578833	Resilience issue: shared FS should have dedicated buffer cache to avoid interference with production FS's
All	77966	572076	Race for ASR helper inode leading to OpHangLimit
All	78022	563167	RPM script change needed to mod inotify max_user_watches to 32768 in sysctl.conf
All	78031	574044, 581478, 578897	fs_moverd segmentation fault occurs on multipart upload to Google Cloud Storage, with full payload disabled
All	78051	577213	install.stornext will not upgrade components if previous install did not fully complete
All	78070	577099	Ubuntu: mounting an Image as loopback result in "mount: /dev/loop0: can't read superblock"
All	78105	579144	snrecover process deadlocks restoring files that have been deleted
All	78116	578751, 583673	fsretrieve -R fails if the files are stored on SAMFS tapes
All	78117	570899	CVFS 6.4.0 still causes d_splice_alias() return -EIO
All	78146	580579	fs_moverd segmentation fault occurs during Google delete batch request, due to null-string boundary
All	78188	581921, 589950	fs_resourced crashing frequently after reporting many "No controllers found" errors
All	78218	546118, 551599, 586977	fsmpm cored when fsnameserver changed
All	78243	582016	Need a fix for CVE-2019-11043
All	78244	582016	Nginx 1.15.x is no longer supported and will not be getting regular patches.
All	78292	583924	fsimport can fail to import tapes if a directory has a single quote in the name

Operating System	Change Request Number	Service Request Number	Description
All	78315	585257	fsrecover is unable to recover files from a long directory path

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 7.0.1 Compatibility Guide](#).

Journal Size Guidelines

StorNext 7.0.1 requires a minimum journal size of 4 MB (Quantum recommends you configure your journal size to 64 MB). If your file system has a journal size less than the 4 MB minimum, you must resize your journal size before you upgrade to StorNext 7.0.1.

⚠ Caution: If you upgrade and the journal size is less than the 4 MB minimum, then the FSM does not start and an error is logged.

Use the `cvupdatefs` utility (see the [StorNext 6 Man Pages Reference Guide](#)) or the GUI (see [Edit a File System](#)) to resize your journal size. When you resize your journal size, the new size must be 16 MB or greater. File systems with journals between 4 MB and 16 MB run with StorNext 7.0.1, but Quantum recommends you configure your journal size to 64 MB.

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 7.0.1

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 7.0.1, special consideration must be made for Linux NFS servers.

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

Compatibility Between StorNext and Other Products

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

Appliance Controller

To view supported Appliance Controller software configurations, see the [StorNext 7.0.1 Compatibility Guide](#).

Infiniband

Infiniband installations require assistance from the Quantum Professional Services team, a Service Partner, or a Quantum Service Provider. For additional information, contact [Quantum Technical Support](#).

Lattus (AXR, S3) or P100/X100

See the [StorNext 7.0.1 Compatibility Guide](#) in the [StorNext Documentation Center](#) for information about compatibility between Lattus (AXR, S3) or P100/X100, and StorNext 7.0.1.

i Note: Object Storage documentation is available online at <https://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 7.0.1, see the *StorNext Partial File Retrieval Compatibility Guide* in the [StorNext Documentation Center](#).

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 Web Services Guide](#) in the [StorNext Documentation Center](#).

Apple Xsan

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

For information about compatibility between Apple Xsan and StorNext 7.0.1, refer to the [StorNext 7.0.1 Compatibility Guide](#) in the [StorNext Documentation Center](#).

Supported Browsers

For information on browsers supported with the StorNext GUI for this release, refer to the [StorNext 7.0.1 Compatibility Guide](#) in the [StorNext Documentation Center](#).

For all other components and features, see the [StorNext 7.0.1 Compatibility Guide](#) in the [StorNext Documentation Center](#).

General Considerations

This section provides information about items to consider for StorNext 7.0.1.

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 might exceed the rate at which the system can generate checksums. In this case, the MD5 checksum calculation will define the throughput of a single data movement operation. With multiple movement streams, MD5 calculations will be done in parallel across the streams and aggregation of performance will be seen.

Upgrading Appliances

⚠ Caution: If you have a Replication, Deduplication, or Object Storage license, see Change Request [Known Issues below](#) in the [StorNext Installation, Replication, HA, and Other Known Issues on page 33](#) section before you upgrade.

For instructions on upgrading your firmware, see [Upgrade the System \(Upgrade Firmware\)](#) on the *Appliance InfoHub Documentation Center* (www.quantum.com/ApplianceInfoHub).

Appliance Release Notes

Refer to the respective Release Notes document for important information you should know about your system.

- [Xcellis Workflow Director](#)
- [Xcellis Workflow Extender](#)
- [Xcellis Foundation](#)
- [Artico](#)

Known Issues

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

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

StorNext File System Known Issues

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

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	54834	3505208, 3516356	<p>If a file is being copied to the StorNext file system using Windows Explorer and Windows Explorer crashes before it finishes copying all the data, the file might contain data blocks from old, deleted files. This problem occurs because Windows Explorer sets EOF to the size of the file before it writes the data to the file. This leaves a gap of uninitialized data in the file.</p> <p>i Note: 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>Workaround:</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>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	67363	n/a	<p>StorNext 5.4.0.x incorrectly allowed the Unix ID Mapping type to be set to none when the Security Model is set to acl. As a result, file systems fail to start when the Unix ID Mapping type is set to none when the Security Model is set to acl.</p> <p>Beginning with StorNext 6, the FSM does not start when this invalid combination of settings is used.</p> <p>Workaround:</p> <p>To prevent this issue, set the Unix ID Mapping to either winbind or algorithmic for any file system where the Security Model is set to acl. You can make the adjustment before or after upgrading.</p>
All	75140	n/a	<p>Exporting an SNFS file system on Ubuntu releases 16.04.2 or later is not supported.</p> <p>Workaround</p> <p>There is currently no workaround for this issue. If you experience this issue, contact Quantum Technical Support.</p>
All	75633	n/a	<p>A StorNext NAS client cannot rename a file if the file has the read-only attribute set. This problem only affects StorNext NAS clients.</p> <p>Workaround</p> <p>A StorNext NAS client must remove the read-only attribute before it can rename the file.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	78082	575600	<p>If you set the parameter audit=true in the file system configuration file, the result causes all I/O activity, including reads, to be recorded in the mdarchive. This can greatly increase the number of updates applied to the mdarchive which, in turn, increases the amount of mdarchive compaction activity by the FSM.</p> <p>There is one instance where this increased compaction activity caused the mdarchive to grow beyond the capacity of the HA shared file system.</p> <p>Workaround</p> <p>To work around this issue, Quantum recommends you increase the parameter metadataArchiveCache from the default 2 GB to at least 4 GB, but preferably larger when setting audit to true in the configuration file.</p>
macOS	66948	322824, 336945	<p>If you access StorNext file systems from Apple Xsan clients, then you might encounter I/O error messages in the system log that do not contain details about real I/O errors detected on the Xsan client.</p> <p>Workaround</p> <p>If you encounter the errors on an Xsan client, contact Apple.</p>

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

macOS	75819	n/a	An Xsan client cannot mount a StorNext File System volume when a cluster number is included in the local fsnameservers file. An unexpected EOF reading reply error is displayed.
-------	-------	-----	---

```
# xsanctl mount snfs3
xsanctl: unexpected EOF reading reply
```

If you encounter the EOF error on your Xsan client, do the following workaround to prevent the issue.

Workaround

If your MDC **fsnameservers** file includes a cluster number, remove the cluster number (**@_cluster_xx**) from the **mysan.configprofile** file **before** you copy the file to the Xsan client.

For example, change:

```
10.65.181.158@_cluster0
```

to

```
10.65.181.158
```

See [Mount the StorNext File System on Xsan 5.0 \(or later\)](#) for additional information.

StorNext Storage Manager Known Issues

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

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	43320	1581004	<p>File retrieves from media to disk can be suboptimal for fast tape drives like the Oracle STK T10K drives. This scenario can occur when the retrieve event is initiated on a host that is different from the host running the mover process, which requires the use of synchronous direct I/O.</p> <p>Workaround:</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"> i Note: This workaround might also help improve the performance of the faster LTO drives by updating the FS_LTO_BLOCK_FACTOR system parameter. i Note: Changes to FS_XXX_BLOCK_FACTOR only affects tapes formatted after the change. <ol style="list-style-type: none"> 1. Change the FS_T10K_BLOCK_FACTOR system parameter from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <pre style="background-color: #f0f0f0; padding: 5px;">FS_T10K_BLOCK_FACTOR=32;</pre> <ul style="list-style-type: none"> i 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. Presently, the FS_T10K_BLOCK_FACTOR system parameter must not be set to a value that exceeds 32. 2. Restart Storage Manager to ensure the change in Step 1 goes into effect: <pre style="background-color: #f0f0f0; padding: 5px;"># tsmstop # tsmstart</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>3. Verify the FS_T10K_BLOCK_FACTOR system parameter contains the new value:</p> <pre data-bbox="818 436 1458 541"># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> <p>4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients.</p> <p>5. Modify <code>/etc/fstab</code> on the MDCs and the DDM clients to use the <code>auto_dma_write_length</code> and <code>auto_dma_read_length</code> mount options as follows:</p> <pre data-bbox="818 814 1458 953">snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_ dma_read_length=16m 0 0</pre> <p>6. Unmount and re-mount your file systems.</p> <p>7. Use new T10K media to store a copy of the file from the disk.</p> <p>i 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. Tapes on which fsformat was run before the change will use the block factor in use at that time. This change will not impact those tapes.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	46693	n/a	<p>Executing the command snbackup -s while a full or partial backup is running might result in a message that <code>/usr/adic/TSM/internal/locks/backup.1f</code> is in an invalid format.</p> <p>This is due to the snbackup -s process reading the backup.1f status file while the backup process is updating it.</p> <p>Workaround:</p> <p>Ignore the message; to clear-up the process, re-execute the command snbackup -s (provided that the backup is not writing to the backup.1f status file while snbackup -s is trying to read it again).</p>
All	47833	n/a	<p>When copying files between media using the CLI command fsmecopy, 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>Note: This behavior might cause file data segment distribution to be sub-optimal on the destination media.</p> <p>Workaround:</p> <p>Currently, a workaround does not exist for this known issue.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69265	n/a	<p>Your DDMs might experience a timeout if you try to connect to the database. The issue is identified by an error log in /usr/adic/TSM/logs/tac which contains the text:</p> <div data-bbox="750 470 1458 779" style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <pre>Process fs_moverd on <host> timed out trying to connect to the database. This usually indicates network connectivity trouble. Try increasing the timeout value by setting the connect_timeout value in /usr/adic/mysql/my.cnf. The default setting is 10 seconds so the new value should be larger.</pre> </div> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Increase the database connection timeout value by adding the following line to /usr/adic/mysql/my.cnf under the section labeled [mysqld] connect_timeout=240. 2. Cycle the Storage Manager in order to pick up the updated timeout value.
All	69341	n/a	<p>If you have the IBM APFO driver installed and configured, then when you perform an fsmedread operation of a partial tape block from a full tape block, the operation can fail with errno=12.</p> <p>i Note: This issue affects all IBM APFO versions 3.0.19 and earlier, and has an impact primarily on disaster recovery procedures.</p> <p>Workaround:</p> <p>To correct this, perform an fsmedread operation without the IBM APFO driver.</p>

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

StorNext GUI Known Issues

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

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

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69958	373823	<p>There is a known issue where stripe group expansion using the StorNext GUI can fail and requires manual intervention to restore file system operation.</p> <p>Stripe group expansion allows an additional disk (LUN) to be added to an existing stripe group, growing the file system user data space.</p> <p>The other way to add space to a file system is to create a new stripe group and add this to the file system.</p> <p>When the GUI executes stripe group expansion, it stops the file system, modifies the configuration file and runs cvupdatefs. If cvupdatefs takes more than 5 minutes to complete, the GUI kills the cvupdatefs utility and reports an error. At this point the file system does not start because the configuration file does not match the current state of the metadata.</p> <p>Workaround:</p> <p>To address this issue, you can do one of two things:</p> <ul style="list-style-type: none"> • Optimally, you would not attempt the stripe group expansion at all. Instead add a new stripe group to the file system. • If stripe group expansion is deemed necessary, use the cvupdatefs CLI directly instead of through the GUI. This method does not have any time limitations. <p>In the case that the StorNext GUI attempt was made and hit the 5 minute timeout, file system operations can be resumed by performing the following:</p> <ol style="list-style-type: none"> 1. Restore the previous version of the configuration file. This can be found in the following directory: <pre style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;">/usr/cvfs/data/<fs>/config_history</pre> 2. Run the cvfsck utility to verify and potentially correct the metadata. 3. Start the file system.

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>Workaround:</p> <p>To workaround this issue, perform the following procedure:</p> <ol style="list-style-type: none"> 1. Shut down the Safari browser window(s). 2. Restart the Safari browser, and then retry the operation. 3. Uncheck the Verify SMTP Server Connectivity box, and then retry the operation. 4. Set Authentication to NONE, and then retry the operation. 5. Disable the Safari User names and passwords AutoFill under Safari > Preferences > AutoFill, and then retry operation.

StorNext Installation, Replication, HA, and Other Known Issues

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

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

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	78382	n/a	<p>Quantum installs a self signed certificate (valid for 365 days) in /usr/cvfs/config/certs/ only when installing a snfs-common RPM and if a preexisting certificate file does not exist.</p> <p>If you use monitoring software (for example, Zabbix), an expired self signed certificate is flagged after it expires.</p> <p>Workaround:</p> <p>To workaround this issue, do the following to update an expired self signed certificate:</p> <ol style="list-style-type: none"> 1. Stop CVFS on the affected client: <pre># service cvfs stop</pre> 2. Update the certificate: <pre># cd /usr/cvfs/config/certs # mv server.crt server.crt.orig # mv server.key server.key.orig # openssl req -x509 -newkey rsa:2048 -keyout server.key -out server.crt -days 365 -subj "/C=US/ST=California/L=SanJose/O=Quantum/OU=Demo/CN=`hostname`" -nodes</pre> 3. Start CVFS on the affected client. <pre># service cvfs start</pre>

Contacting Quantum

Contacts

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

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

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

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

For worldwide support:

<https://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

Quantum®

ABOUT QUANTUM

Quantum technology and services help customers capture, create, and share digital content—and preserve and protect it for decades. With solutions built for every stage of the data lifecycle, Quantum's platforms provide the fastest performance for high-resolution video, images, and industrial IoT. That's why the world's leading entertainment companies, sports franchises, researchers, government agencies, enterprises, and cloud providers are making the world happier, safer, and smarter on Quantum. See how at www.quantum.com.

www.quantum.com • 800-677-6268