

StorNext 6.x Xcellis Foundation Release Notes

Contents

Revision History	2
About the Xcellis Foundation System	3
Training and Documentation Resources	3
About StorNext 6	3
What StorNext Releases are compatible with my system?	4
Upgrade StorNext Software and System Firmware	4
General Notes	4
Quantum Appliance Licenses	4
Fixed Issues, Notes, and Enhancements – StorNext 6.x	5
Known Issues – StorNext 6.x	19
Known Issues Workarounds and Additional Information	41
Contacting Quantum	58

© 2021 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law. Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, DXiV-Series, FlexSync, FlexTier, Lattus, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners. Quantum specifications are subject to change.

Revision History

The following StorNext Releases are supported by these Release Notes:

StorNext 6 Releases Supported	Release Notes Date
StorNext 6.4.1	February, 2021
StorNext 6.4.0	May, 2020
StorNext 6.3.1.1	December, 2019
StorNext 6.3.1	November, 2019
StorNext 6.3.0	August, 2019
StorNext 6.2.1	June, 2019
StorNext 6.2.0	March, 2019
StorNext 6.1.1	February, 2019
StorNext 6.1.0	September, 2018
StorNext 6.0.6.1	May 2018
StorNext 6.0.6	March 2018
StorNext 6.0.5.1	February 2018
StorNext 6.0.5	January 2018
StorNext 6.0.1.1	September 2017

About the Xcellis Foundation System

The Quantum Xcellis Foundation system is an affordable, easy-to-install, simple-to-manage scale-out NAS storage appliance that provides the capacity and performance you need for both today and tomorrow. It is a turnkey solution with the features, performance, and scalability of an enterprise-grade system, sized and priced for smaller workgroups and organizations. Integrated storage tiering efficiently moves less-frequently used data to the cloud or other archive tiers to drive down the cost of long-term storage and preserve valuable free space in Xcellis Foundation.

For more information about the Xcellis Foundation system, see [About the Xcellis Foundation System](#) on quantum.com.

Training and Documentation Resources

[Xcellis Foundation training and documentation](#)

If you are unsure about which type of system you have visit the "[Identify My System](#)" page to determine which system you have, and to find the training and documentation resources for that specific system.

About StorNext 6

What is StorNext?

StorNext® is high-performance, multi-tier shared storage designed for large, data-intensive workloads. It includes Xcellis workflow storage, extended online storage, and tape archives—all powered by StorNext, the industry's fastest streaming file system and policy-driven data management software.

At the core of all Quantum scale-out storage is StorNext advanced data management—engineered to tackle the world's most demanding workloads, with the performance and efficiency needed to cost-effectively achieve desired business results.

StorNext 6 delivers the unique combination of high performance and advanced data management, providing cost-effective scalability and access for a wide variety of workloads and use cases, including media production, genomics research, video surveillance, geospatial imaging, VR content, and more.

What StorNext Releases are compatible with my system?

Learn More on the Web...

[Learn more](#) about StorNext 6 and Scale-Out Storage and take a look at the [StorNext 6 Documentation Center](#), which includes training and documentation resources for StorNext 6, including Release Notes.

What StorNext Releases are compatible with my system?

See the "StorNext Upgrade Matrix" and "StorNext Appliance Compatibility" sections of the [StorNext 7.x Compatibility Guide](#) on quantum.com for information about system compatibility.

Upgrade StorNext Software and System Firmware

Before upgrading your system, refer to [Known Issues](#). This section contains important information you need to know before upgrading.

See [Upgrade the System \(Upgrade Firmware\)](#) for the steps necessary to upgrade to the latest StorNext release and hardware firmware for your system.

General Notes

Refer to the [General Notes](#) page on quantum.com for important information you should know about your system.

Quantum Appliance Licenses

See [Quantum Appliance Licenses](#) page on quantum.com.

Fixed Issues, Notes, and Enhancements – StorNext 6.x

This section lists the fixed issues, enhancements and notes for different StorNext Releases supported for your system.

- [Fixed Issues, Notes, and Enhancements for StorNext 6.4.1 below](#)
- [Fixed Issues, Notes, and Enhancements for StorNext 6.4.0 below](#)
- [Fixed Issues, Notes, and Enhancements for StorNext 6.3.1.1 on the next page](#)
- [Fixed Issues, Notes, and Enhancements for StorNext 6.3.1 on the next page](#)
- [Fixed Issues, Notes, and Enhancements for StorNext 6.3.0 on page 7](#)
- [Fixed Issues and Enhancements for StorNext 6.2.1 on page 8](#)
- [Fixed Issues and Enhancements for StorNext 6.2.0 on page 10](#)
- [Fixed Issues and Enhancements for StorNext 6.1.1 on page 12](#)
- [Fixed Issues and Enhancements for StorNext 6.1.0 on page 12](#)
- [Fixed Issues and Enhancements for StorNext 6.0.6.1 on page 14](#)
- [Fixed Issues and Enhancements for StorNext 6.0.6 on page 15](#)
- [Fixed Issues and Enhancements for StorNext 6.0.5.1 on page 16](#)
- [Fixed Issues and Enhancements for StorNext 6.0.5 on page 17](#)
- [Fixed Issues and Enhancements for StorNext 6.0.1.1 on page 18](#)

Fixed Issues, Notes, and Enhancements for StorNext 6.4.1

There are no new appliance-specific code fixes or enhancements for this release. However, you can reference the following:

- See the [StorNext 6.4.1 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues, Notes, and Enhancements for StorNext 6.4.0

Table 1: Fixed Issues/Enhancements for StorNext 6.4.0

CR Number	SR Number(s)	Description
77096	n/a	New entries in the <code>/etc/multipath.conf</code> file now support new Quantum storage array models.

CR Number	SR Number(s)	Description
76484	540318	DAE VMs are now correctly managed from the StorNext GUI, and "No content" error messages are no longer generated.
75394	517419	DAE now correctly handles multiple configured StorNext file systems when VMs are started

- See the [StorNext 6.4.0 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues, Notes, and Enhancements for StorNext 6.3.1.1

StorNext 6.3.1 has been superseded and is replaced by StorNext 6.3.1.1. All of the fixes and enhancements for StorNext 6.3.1 still apply to 6.3.1.1.

- i Note:** Quantum recommends you upgrade to StorNext 6.3.1.1 if you have a system with Mellanox Connect CX4-Lx cards configured for 25G connections and running 6.3.1. These cards can fail to operate when running StorNext 6.3.1.

Table 2: Fixed Issues/Enhancements for StorNext 6.3.1.1

CR Number	SR Number(s)	Description
76354	n/a	Firmware upgrade now reports and detects Mellanox LSPCI device ID correctly.
76353	n/a	StorNext 6.3.1.1 upgrade includes the Mellanox CX4-Lx 25 GbE firmware for 25 GbE networks.

- See the [StorNext 6.3.1.1 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues, Notes, and Enhancements for StorNext 6.3.1

StorNext 6.3.1 has been superseded and is replaced by StorNext 6.3.1.1. All of the fixes and enhancements for StorNext 6.3.1 still apply to 6.3.1.1.

- i Note:** Quantum recommends you upgrade to Appliance Controller 2.3.2 if you have one or more F2000 storage arrays. 2.3.2 includes updates to the `qstorage` drive attach/detach commands to work with the F2000 storage array RAID configuration. See the [Appliance Controller upgrade page](#).

Table 3: Fixed Issues/Enhancements for StorNext 6.3.1

CR Number	SR Number(s)	Description
75477	518583	Upgraded to iDRAC version 2.61.60.60 from 2.21.21.21 due to security vulnerabilities in the earlier release.
75627	520429	Health output is now provided when the qstorage show health command is run from the Appliance Controller console.

- See the [StorNext 6.3.1 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues, Notes, and Enhancements for StorNext 6.3.0

- i Note:** Quantum recommends you upgrade to Appliance Controller 2.3.1 if you have one or more QXS 12G 484 or F2000 storage arrays. 2.3.1 includes RAS alert support for these storage arrays. See the [Appliance Controller upgrade page](#).

Table 4: Fixed Issues/Enhancements for StorNext 6.3.0

CR Number	SR Number(s)	Description
75158	n/a	A new SSL certificate is now included in the 6.3.0 release for Cloud-Based Analytics.
75094	n/a	Firmware upgrade will now include a status of "Incomplete" if a step in the StorNext Connect upgrade times out or is incomplete.
74888	499350, 494119, 488583	"reg show" lines in the "messages" log are no longer captured every 4 hours after upgrading to 6.3.0. For 6.2.x, this was known issue.
74838	n/a	Managed configuration files due to system time issues. This fixes an issue where, after an upgrade to 6.2.1, NTP changed the system time which resulted in "FC HBA 0 is missing" RAS created for server node1..

CR Number	SR Number(s)	Description
74792	n/a	The Admin Alert column in the validation table in the StorNext GUI now correctly shows only the important information about the alert during validation of a firmware upgrade file. Previously, a few unnecessary items were included in the validation table when firmware upgrade files with incorrect checksums were uploaded to the system, and Validate was clicked.
74649	n/a	The Appliance Controller no longer throws an exception when it is unable to reach a QXS array.
74526	n/a	Enhancement: iSCSI is now supported for multiple QXS RAID arrays. i Note: This fix requires upgrading to Appliance Controller 2.3.0.
74316	n/a	Firmware upgrade to StorNext 6.3.0 upgrades OMSA (OpenManage server administrator) to 9.2.0.
74269	n/a	Complete RAS messages are now created for unsupported PCIe cards installed in the system.
74258	n/a	In the presence of bad drive, the qxsarray show health command results in correct output. i Note: This fix requires upgrading to Appliance Controller 2.3.0.
74245	n/a	Firmware upgrade process no longer changes slot configuration information; card slot configuration settings are kept.
74201	n/a	CX-3 card ports mode settings no longer changed after firmware upgrades; original configuration settings are kept.

- See the [StorNext 6.3.0 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues and Enhancements for StorNext 6.2.1

i Note: StorNext 6.2.1 replaces StorNext 6.2.0, and includes all fixes issues and enhancements completed for StorNext 6.2.0. StorNext 6.2.0 firmware upgrade files are no longer available.

This release addresses the following vulnerability:

CentOS7 Docker (runc) vulnerability – versions prior to 18.09.2. (see [74017 on the next page](#))

Table 5: Fixed Issues/Enhancements for StorNext 6.2.1

CR Number	SR Number(s)	Description
74780	500550	The Appliance Controller now restarts correctly after applying the firmware upgrade, and snmp and Cloud-based Analytics function normally.
74724	498766, 496820	The system will now boot correctly after upgrading firmware.
74721/ CON-4344	498766	Upgrades to StorNext 6.2.1 are successful even if there are duplicate entries in /etc/hosts.
74697	498571, 497610	The StorNext Connect Docker installation during upgrade no longer fills up the /var directory on the server, so upgrades will no longer fail due to this issue.
74667	496820	<p>StorNext Connect failures occurring during firmware upgrades now generate an admin alert, and the upgrade will complete successfully. Admin alert message displayed on StorNext Connect docker upgrade failure:</p> <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p>Connect upgrade is incompleteFailed to upgrade to Connect dockers.</p> </div>
74626	0489653	The NAS controller in the Appliance Controller is now restarted correctly after firmware upgrades, fully supporting pre-upgrade checks of the system.
74543	501460	The StorNext ConnectGUI now appears correctly, and Discover complete properly and stats are displayed correctly.
74303	n/a	Autoflexconfig now supports either type of port mode setting (Ethernet or Infiniband) that may be used with the ConnectX-3 card.
74017	n/a	Docker (runc) upgraded to CVE-2019-5736 to address a security vulnerability in the CentOS7 Docker (versions prior to 18.09.2) where the host runc binary could be overwritten and allow host root access.
74010	n/a	The connecthostname is no longer set to 10.1.1.2 on server node1 and server node2 after firmware upgrades for systems that previously was not configured for StorNext Connect.

CR Number	SR Number(s)	Description
72147		For systems already upgraded to StorNext 6.2.0 or later, the "Validate" button on the Upgrade Firmware page in the StorNext GUI now executes pre-upgrade checks before initiating the upgrade process. This is the first release where this button initiates this pre-upgrade check.

- See the [StorNext 6.2.1 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues and Enhancements for StorNext 6.2.0

i Note: StorNext 6.2.1 replaces StorNext 6.2.0, and includes all fixes issues and enhancements completed for StorNext 6.2.0. StorNext 6.2.0 firmware upgrade files are no longer available.

The new appliance-specific code fixes or enhancements for this release are:

Table 6: Fixed Issues/Enhancements for StorNext 6.2.0

CR Number	SR Number(s)	Description
73712		Added time-stamp logging for each of the steps in the upgrade process to more easily identify which steps take longer to run.
73485	469076	The qxsarray feature in the Appliance Controller console (CLI) now correctly handles both Fibre Channel and iSCSI-connected QXS array information when the qxsarray show command is run.
73476		snupdate.log no longer consumes most of the root file system, previously caused when the Appliance Controller ran the /opt/quantum/snupdate/bin/snupdate show versions command every minute.
73194		Enhanced the Network Menu within the Service Menu to allow reconfiguring of previously-configured interfaces
72983		Updated the StorNext GUI Networking page (Configuration > System) to remove metadata-specific selections, and to add active/backup bonding support.
72977		Better handling of manually-configured VLANs in the StorNext GUI.

CR Number	SR Number(s)	Description
72919		Simplified and enhanced the Network Menu within the Service Menu with an active/backup bonding option.
72800	443258, 443259, 443260, 433261, 470240	Fixed CVE-2014-3566 Nexpose scan for SSLv3 and POODLE warnings – port 33777 (CentOS7)
72624		Added the PERC controller log to the RAS emails for server boot drive failures.
72586		The Appliance Controller qxsarray show iscsi command executes correctly for QXS arrays with firmware release names containing a "GFS" prefix (for iSCSI-connected systems ONLY).
72421	424025	Security Scanner now leaves the FSM operating correctly on the secondary (failover) server node (typically node 2).
72240 (HYDRA-3672, 3716, 3717, 3740)		RAS tickets for QXS drive failures are now generated with the same detail that is provided in RAS tickets for system boot drive failures.
72147		Add a pre-upgrade check feature on the GUI and in the upgrade scripts enables the "Validate" button on the Upgrade Firmware page in the StorNext GUI (for systems upgrading from releases after StorNext 6.2.0) to execute pre-upgrade checks before beginning the upgrade process.
72006		Updated the firmware upgrade code to include fixes for Spectre/Meltdown vulnerabilities (R630 models ONLY).
71981		Background DAE jobs that the StorNext GUI initiates, now show success for the job status even though the command may fail.
71809	431853	Fixed the QUANTUM-STORNEXT-APPLIANCE-MIB for snaFileSystemActive OID , so it no longer returns a bad data type for SNMB.
71299	456036, 459515, 460901	Running a Nexpose scan no longer results in a fs_copymand/fs_moverd core dump.
71147	451310	Clients using NSS2 can now see file systems that are not on same subnet as coordinators.

CR Number	SR Number(s)	Description
70915	409963, 455314	For upgrades to 6.2, /tmp has been changed to a symbolic link which points to /scratch/tmp , which prevents the root file system (/) from getting filled up. This prevents processes like pse_snapshot or fsmesdscan (see 72726 on page 22) from overwhelming the root file system with data.
70212		Integrate QXS and metadata storage array hardware events into StorNext RAS system, to generate RAS tickets for array drive failures.
69255 (HYDRA-2004)	349873	Implemented a fix to address the SSL/TLS server factoring RSA export keys (FREAK) vulnerability.
53486	425247	Srvclg RAS tickets are created even if the hafs (HA shared file system) is not mounted.

- See the [StorNext 6.2.0 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues and Enhancements for StorNext 6.1.1

- i Note:** StorNext 6.1.1 replaces StorNext 6.1.0, and includes all fixes issues and enhancements completed for StorNext 6.1.0, including the Spectre and Meltdown vulnerabilities addressed in that release. StorNext 6.1.0 firmware upgrade files are no longer available.

There are no new appliance-specific code fixes or enhancements for this release. However, you can reference the following:

- See the [StorNext 6.1.1 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues and Enhancements for StorNext 6.1.0

- i Note:** StorNext 6.1.1 replaces StorNext 6.1.0, and includes all fixes issues and enhancements completed for StorNext 6.1.0, including the Spectre and Meltdown vulnerabilities addressed in that release. StorNext 6.1.0 firmware upgrade files are no longer available.

This release addresses Spectre and Meltdown vulnerabilities as follows:

CentOS Version	Kernel	Vulnerability Variant Addressed
CentOS7 Update 4	3.10.0-693.21.1	1,2,3

Table 7: Fixed Issues/Enhancements for StorNext 6.1.0

CR Number	SR Number(s)	Description
71708	440392	Enhanced card discovery to reflect when PCI slot addresses change. Update configuration files with the new addresses.
71599	425749	Firmware upgrade now ensures autoconfig is enabled.
71507	421334	Firmware upgrade now upgrades PHP on server nodes.
71469	421581	Firmware upgrades previously failed during OS Update due to insufficient space on local/root filesystem. Available space is now checked during the system pre-upgrade check, and, if not enough free space is available, the firmware upgrade does not attempt to begin. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i Note: You will still need to make 5 GB free space available for the upgrade and restart the upgrade process.</p> </div>
71328	403647	Systems with large LUNs no longer take excessive time to boot.
71230	414352, 429898	Nexpose security scan previously reported HTTP Header was missing on port 443.
71125	410540	Rsyslog stopped logging into /var/log/messages but all logs were included in systemd-journal.
71012	0399206	The StorNext Connect GUI showed the status of an appliance server node being undefined, and was displayed as a gray dot on the dashboard, as a result of a Fibre Channel card being part of the expected hardware installed for the system. The Fibre Channel requirement has been removed, so for systems upgraded to StorNext 6.1 and later, refresh the dashboard page, and it will correctly display the system as a green dot.
70854	383195	<code>blacklist_exceptions</code> and entries for EMC/NetApp arrays are now included in the reference <code>multipath.conf</code> file provided on appliances to support EMC VNX7600, EMC VNX8000, EMC VMAX 20k and NETAPP FAS6240 arrays. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i Note: See also the new How to Modify the multipath.conf File for External Storage Arrays page released for StorNext 6.1.0.</p> </div>

CR Number	SR Number(s)	Description
70649	398896	The Promise array path_checker definition has been updated in the reference multipath.conf file provided on appliances. i Note: See also the new How to Modify the multipath.conf File for External Storage Arrays page released for StorNext 6.1.0.
70588	396789	tsunami.log and service.log files are now rotated to prevent these logs and snapshot files becoming too large, and helps snapshot files to be generated more quickly.
70526	398061, 396903, 398206, 397045, 399232, and 401331	Kernel Page Table Security Problems addressed for CVE-2017-5753, CVE-2017-5715 and CVE-2017-5754.

- See the [StorNext 6.1.0 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues and Enhancements for StorNext 6.0.6.1

i Note: StorNext 6.0.6.1 replaced previous 6.x StorNext releases, and firmware upgrade files will only be provided for StorNext 6.0.6.1 or later.

i Note: Spectre and Meltdown vulnerabilities were not addressed in this release.

Table 8: Fixed Issues/Enhancements for StorNext 6.0.6.1

CR Number	SR Number(s)	Description
71394	0421219	Can now successfully install a Mellanox ConnectX-3 VPI InfiniBand/Ethernet card in the system. The MLNX_OFED HBA driver was replaced for StorNext 6.0.6.1. The new driver allows the Mellanox HBA card to be configured.
71350	n/a	For dual-server node and for NAS clustered systems, after upgrading the system to 6.x, StorNext file systems (snfs) mounts correctly on server node 2 when over a billion files are present on the snfs.

CR Number	SR Number(s)	Description
71346	n/a	For systems running StorNext 5.3.x upgrading to StorNext 6.0.6.1 and configured as a single-node NAS cluster (a "cluster of one") running NAS 1.2.1, the NAS/Controller software should first be upgraded to 1.4.1 or later, but will not cause the firmware (.fw) upgrade to fail.
		<p>i Note: Before applying firmware upgrades for NAS-clustered systems, Quantum recommends upgrading to the latest version of the Appliance Controller.</p>

- See the [StorNext 6.0.6.1 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues and Enhancements for StorNext 6.0.6

- i Note:** StorNext 6.0.6.1 replaced previous 6.x StorNext releases, and firmware upgrade files will only be provided for StorNext 6.0.6.1 or later.

This release includes the following software updates:

- CentOS 7.4 for the server node(s)
- CentOS 7.4 SAN and DLC clients
- Appliance Controller 2.0.1 (and supports CentOS 7.4)

The following new appliance-specific code fixes and enhancements are included in this release:

Table 9: Fixed Issues/Enhancements for StorNext 6.0.6

CR Number	SR Number(s)	Description
71002	n/a	RAS ticket for an IPMI controller failure is no longer shown after the firmware upgrade.
70902	n/a	The system prevents firmware upgrades when NAS clusters are configured for Appliance Controller/NAS versions prior to 1.4.1. You will still need to upgrade the Appliance Controller/NAS prior to applying the firmware upgrade. See Upgrade Required for Clustered StorNext NAS Systems .
70773	n/a	The <code>/etc/resolv.conf</code> configuration file is now included when collect files are created (in the tar.gz file at: <code>usr/adic/tmp/platform/hw-info/etc/resolv.conf</code>).

CR Number	SR Number(s)	Description
70770	n/a	aiWARE agent failures are no longer included in system log files for non-aiWARE, CentOS7-based systems.
70455	n/a	For systems with multiple Ethernet NIC ports active, when an Ethernet cable is unplugged from a port, the admin alert now contains the affected NIC port number.
70340	n/a	PHP upgraded to php56w-5.6.33-1.
70338	n/a	Dracut updated to dracut-033-502.el7; included with CentOS7.4.
70335	n/a	YUM help output is no longer displayed in <code>/var/log/messages</code> during fw upgrade failures.
70291	n/a	A root mailbox junk filter is included in this release. A new daily cron job trims mailboxes with over 1000 email messages & improves functionality of the logrotate cron job.
70208	n/a	Admin alerts are now created in the StorNext GUI when NIC ports go down.
70194	n/a	For systems using snmpwalk , the time taken by SNMP Agent MIB processing has been greatly reduced.
70178	n/a	During the system firmware upgrade, the current <code>/etc/multipath.conf</code> file now remains on the system as-is.
69478	n/a	Update server code to CentOS 7.4.

- See the [StorNext 6.0.6 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues and Enhancements for StorNext 6.0.5.1

i Note: StorNext 6.0.6.1 replaced previous 6.x StorNext releases, and firmware upgrade files will only be provided for StorNext 6.0.6.1 or later.

There are no new appliance-specific code fixes or enhancements for this release. However, you can reference the following:

- See the [StorNext 6.0.5.1 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues and Enhancements for StorNext 6.0.5

i Note: StorNext 6.0.6.1 replaced previous 6.x StorNext releases, and firmware upgrade files will only be provided for StorNext 6.0.6.1 or later.

Because of security vulnerabilities found in testing (highlighted below in CR [69866 below](#) and CR [69869 below](#)), Quantum recommends upgrading to 6.0.5.1 or later.

The following new appliance-specific code fixes and enhancements are included in this release:

Table 10: Fixed issues/Enhancements for StorNext 6.0.5

CR Number	SR Number(s)	Description
70300	385040	Simultaneous Fibre Channel-connected QSX and Netapp arrays attached to Quantum appliances no longer prevent system boot after firmware upgrades to 6.0.5 from supported 5.x releases. The base/default SCSI enclosure drivers provided by CentOS which previously caused this issue are no longer referenced in system configuration files. The specific drivers used for QXS and NetApp array FC connections are used instead.
69869	n/a	Updated Apache HTTPD to the latest CentOS 6.9 and 7.4 versions to resolve some HTTPD security vulnerabilities found by Nexpose.
69866	n/a	Updated php56w from version 5.6.29-1 to version 5.6.31-2 to resolve some PHP security vulnerabilities found by Nexpose.
69506	n/a	RAS tickets are no longer created when the system is unable to obtain fan status from the iDRAC. i Note: This fix applies to CentOS7 (Xcellis Workflow Director (R630), Xcellis Foundation, aiWARE for Xcellis Embedded (DAE), aiWARE for Xcellis Standard and High Performance, Artico (R630), and Xcellis Workflow Extender (R630) systems) only.

CR Number	SR Number(s)	Description
69492	n/a	<p>The RAS ticket for a failed server drive now contains several new attributes, including the drive size, as shown in this example of a failed drive:</p> <pre>Node 1 The Drive PHYSICAL_DISK_0:1:1 has failed. Capacity: 465.25GB Media: HDD Encrypted: NA Vendor: DELL(tm) VendorPN: MY01KWKJ1255523P091GA0 VendorSN: WD-WMAYP3387029 Model: WDC WD5003ABYX-18WERA0 PowerStatus: Spun Up Status: Offline</pre>
69473	n/a	<p>Updated the build to include a new CentOS 7 (7.3) kernel with an additional patch. This corrects an issue introduced with the RHEL7.3 kernel that failed to display and manipulate ACLs from NFS NAS clients.</p> <p>i Note: This fix applies to CentOS7 (Xcellis Workflow Director (R630), Xcellis Foundation, aiWARE for Xcellis Embedded (DAE), aiWARE for Xcellis Standard and High Performance, Artico (R630), and Xcellis Workflow Extender (R630) systems) only.</p>
67895	n/a	<p>Added elapsed time for firmware upgrades in the StorNext GUI upgrade firmware status page, and in the upgrade_progress log file.</p>
66799	322817, 342909, 362343	<p>In the StorNext GUI, the Reports > Hardware page now correctly shows system board information for node1 and node2 after StorNext 6.0.5 installation StorNext.</p>

- See the [StorNext 6.0.5 Release Notes](#) for information about StorNext software updates for this release.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Fixed Issues and Enhancements for StorNext 6.0.1.1

- i Note:** StorNext 6.0.6.1 replaced previous 6.x StorNext releases, and firmware upgrade files will only be provided for StorNext 6.0.6.1 or later.

The following new appliance-specific code fixes and enhancements are included in this release:

Table 11: Fixed Issues/Enhancements for StorNext 6.0.1.1

CR Number	SR Number(s)	Description
69706	363354	After rebooting the system configured with iSCSI connections, the iSCSI connection will now re-establish correctly, and StorNext file systems will also restart.

i Note: * The StorNext 6.0.1.1 release only applied to R630 server models (which run on CentOS7), and not R520 server models (which run on CentOS6).

- See the [StorNext 6.0.1 Release Notes](#) for information about StorNext software updates.
- See the [Upgrade the System \(Upgrade Firmware\)](#) for the instructions for upgrading system firmware.

Known Issues – StorNext 6.x

This section lists the known issues that could potentially affect your system.

CR Number	SR Number	Description
77518	n/a	Mellanox ConnectX-3 (CX-3) cards in G300Pro Foundationsystems is not updated by the firmware upgrade.
77317	n/a	<p>Adding a DAE virtual machine to a file system created on a QXS or P-Series LUN can create a RAS alert stating:</p> <pre>Excessive fragmentation detected in file "/DAE_Boot_Disk_location/pk-DAE-vm.qcow2", inode 0x4e80000ef0b19 (decimal inum 15665945)'</pre> <p>Workaround:</p> <p>While this issue is not currently fixed, the message itself may be disabled. To disable the RAS alert, change the value of the extentCountThreshold parameter in the file system configuration file to "0". For information about file system configuration files, see:</p> <p>https://gsupport.quantum.com/kb/flare/Content/stornext/SN6_DocSite/Default.htm#Guide_Tuning/Topics/Example_FSM_Configuratio.htm</p>

CR Number	SR Number	Description
76675	543978	<p>If your system is running as a FlexSync host and you upgrade StorNext, it also installs the latest version of FlexSync. If there are other systems configured as FlexSync hosts, FlexSync will not function on those other systems until they are upgraded to the latest available FlexSync version.</p> <p>Error message displayed:</p> <div data-bbox="505 527 1458 596" style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> <p>Host Unreachable, Upgrade Required</p> </div> <p>The same FlexSync version must be installed on ALL FlexSync hosts (see Can different versions of FlexSync be installed or upgraded within a configuration? on the <i>FlexSync Documentation Center</i>).</p> <p>Workaround:</p> <p>Upgrade FlexSync on ALL FlexSync hosts. See Upgrade FlexSync in the <i>FlexSync Documentation Center</i>.</p>
76353, 76354	n/a	<p>Mellanox Connect CX4-Lx cards can fail to operate if configured for 25 GbE connections and running 6.3.1.</p> <p>i Note: This issue has been fixed in 6.3.1.1. See 76353 on page 6 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <p>Upgrade to StorNext 6.3.1.1.</p>

CR Number	SR Number	Description
74024	470777, 479905	<p>The following applies to Xcellis Foundation systems ONLY. It does not apply to other server model names/generations.</p> <p>Upgrade between StorNext releases prior to 6.2.1 removes custom Infiniband configuration settings.</p> <ul style="list-style-type: none"> i Note: This issue was Fixed in 6.3.0. Upgrades to StorNext 6.3.0 DO NOT experience this issue. If this has not happened to your system, you can upgrade to 6.3.0 without this issue affecting your system. i Note: Because the workaround is required after every upgrade to StorNext releases prior to 6.2.1, Quantum strongly recommends you to upgrade the system to 6.2.1 or later prior to contacting Quantum Support about this issue, so the fix will be applied to a release that no longer removes the Infiniband settings, and it only has to be applied once. <p>After upgrading to 6.2.1 or later (prior to 6.3.0):</p> <p>If you have this problem and have upgraded the system to 6.2.1 or later, contact Quantum Support and reference TSB (Technical Service Bulletin) 3065.</p> <p>Once the system is upgraded to 6.2.1, and the fix is applied, future upgrades will perform successfully, and will not remove the Infiniband configuration settings.</p>
74571	n/a	<p>For StorNext 6.3.0 and Appliance Controller 2.3.0, there is a situation where the firmware upgrade process pre-stages Appliance Controller upgrade file(s) on server node 1, but the same files are not pre-staged to server node 2, which can create a failure for firmware upgrades. This can happen if you click "Validate" on the firmware upgrade page, or the function that checks for upgrade dependencies when you click "Activate" on the firmware upgrade page.</p> <p>If this happens, and Node 2 was the master node of a NAS cluster, you will not be able to run the Appliance Controller upgrade local command on node 2 to upgrade to the current controller release pre-staged on the server, and you will have to manually move the nas rpms from the /var/upgrade or /scratch/saved/NAS directory on node 1 to the /var/upgrade directory on node 2, and upgrade the NAS cluster manually from Node 2.</p> <p>See: http://qsupport.quantum.com/kb/Flare/Content/appliances/ACC/DocSite/Upgrades/ACC_Upgrades.htm</p> <ul style="list-style-type: none"> i Note: See About firmware upgrades and pre-staged Appliance Controller upgrades for server nodes on page 41 for information about how these files are pre-staged to enable you to upgrade Appliance Controller manually, separate from the firmware upgrade process.

CR Number	SR Number	Description
74202, 67955, HYDRA- 4326	n/a	<p>If NAS is configured on the system and StorNext services are stopped, StorNext services can fail to stop completely. This can block server fail-over and potentially leaves managed file systems without running FSMs.</p> <p>See 74202 – StorNext services can fail to stop completely when NAS is Configured and StorNext services are stopped on page 41 for an in-depth explanation of the issue and several alternate workarounds.</p>
73688	n/a	<p>When the <code>/var</code> directory on the appliance gets full, the following serious errors can occur:</p> <ul style="list-style-type: none"> • Network communication errors on the server node • iDRAC is not accessible • The server node gets stuck during reboot and require a physical power cycle <p>See also Leave Space on Appliance File Systems on the <i>Appliance InfoHub</i>.</p> <p>Workaround:</p> <p>Keep additional space available on <code>/var</code>.</p>
73485	469076	<p>The <code>qxsarray show</code> feature in the Appliance Controller console (CLI) no longer returns the following error: "Operation failure: Failed to retrieve all QXS array configuration data. Space required after the Public Identifier, line 1, column 47 (E-2003)".</p>
73036	n/a	<p>Systems running Java 1.7.1, 1.8.1 or later are no longer able to launch the iDRAC Virtual Console using StorNext 6.2.0 or later, due to an increase in Java security.</p> <p>Workaround:</p> <p>Uninstall the current Java version, and install Java 1.7.0. iDRAC should work properly again.</p>
72726	3734274, 336772, 409963, 430039, 455314	<p>If you run <code>fsmedscan -R</code>, the root file system on the active server node (typically node 1) can fill up and the node will become unresponsive. This is because <code>fsmedscan</code> creates copious log files and, by default, places them in <code>/tmp/logs</code>, which will eventually fill up the root system drive <code>/</code>.</p> <p>Note: This issue was fixed in StorNext 6.2. See 70915 on page 12 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workarounds:</p> <ul style="list-style-type: none"> • Upgrade to StorNext 6.2. • Prior to 6.2, you could create a symbolic link from <code>/tmp/logs</code> to something like <code>/scratch/fsmedscan_logs</code> since <code>/scratch</code> should contain plenty of available space.

CR Number	SR Number	Description
71809	431853	<p>snaFileSystemActive OID can return a bad data type for SNMB information.</p> <p>i Note: This issue was fixed in StorNext 6.2. See 71809 on page 11 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <p>Upgrade to StorNext 6.2.</p>
71708	440392	<p>Issue:</p> <p>When slot PCI addresses change, need to reflect their new location in the system.</p> <p>i Note: This issue was fixed in StorNext 6.1.0. See 71708 on page 13 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <p>Upgrade to StorNext 6.1.0. There is no other workaround available.</p>
71469	421581	<p>Issue:</p> <p>Firmware upgrade failed during the OSUpdate step of the upgrade process, due to insufficient space on local/root filesystem.</p> <p>i Note: Free space is now checked before continuing with the firmware upgrade in StorNext 6.1.0. See 71469 on page 13 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <ol style="list-style-type: none"> Determine which directory on the root directory has files that can be removed: <pre>df -h</pre> <p>Notice any directories you have created that consume GBs of data.</p> Remove files so you ensure the system has at least 5 GB of free space for the firmware upgrade to use before attempting the upgrade. Restart the failed upgrade.

CR Number	SR Number	Description
71328	403647	<p>Issue:</p> <p>Systems can a very long time to get from POST until the filesystems are mounted with large LUNs present, since all devices are added to the sysfs and on udev. During boot, a message similar to the the following is displayed:</p> <pre>mam-node-1 udevd-work[39397]: inotify_add_watch(6, /dev/dm-2699, 10) failed: No space left on device</pre> <p>Note: This issue was fixed in StorNext 6.1. See 71328 on page 13 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <p>Upgrade to StorNext 6.1.0.</p> <p>OR</p> <p>Increase the amount of notify watches. Do the following:</p> <ol style="list-style-type: none"> 1. Run the following command: <pre>cat /proc/sys/fs/inotify/max_user_watches</pre> <p>This should return a value similar to:</p> <pre>10240</pre> 2. Edit the /etc/sysctl.conf file to allow for a longer boot time: <pre>fs.inotify.max_user_watches=32768</pre>
71299	456036, 459515, 460901	<p>Running a Nexpose scan on systems running StorNext 6.0.6.x can result in a fs_copymand/fs_moverd core dump.</p> <p>Note: This issue was fixed in StorNext 6.2. See 71299 on page 11 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workarounds:</p> <ul style="list-style-type: none"> • Exclude the system from security scans. • Upgrade to StorNext 6.2.

CR Number	SR Number	Description
71230	414352, 429898	<p>Issue: Nexpose security scan reports HTTP Header is missing on port 443.</p> <p>i Note: This issue was fixed in StorNext 6.1. See 71230 on page 13 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround: Upgrade to StorNext 6.1.0. There is no other workaround available.</p>
71125	410540	<p>Issue: /var/log/messages can omit expected information beginning with a specific point in time.</p> <p>i Note: This issue was fixed in StorNext 6.1. See 71125 on page 13 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround: To correct this issue: Upgrade to StorNext 6.1.0.</p> <p>OR</p> <ol style="list-style-type: none"> 1. Remove imjournal.state: <pre>rm -f /var/lib/rsyslog/imjournal.state</pre> <p>Restart rsyslog:</p> <pre>/sbin/service rsyslog restart</pre> <p>Then rotate the log:</p> <pre>\$logrotate /etc/logrotate.conf</pre> 2. If the logging stops again, disable the imjournal input. To do this, edit the /etc/rsyslog.conf file, and comment out <code>OmitLocalLogging on</code> as shown: <pre>#\$OmitLocalLogging on</pre>

CR Number	SR Number	Description
70854	383195	<p>Issue:</p> <p>blacklist_exceptions and entries for EMC/NetApp arrays are not included in the default multipath.conf file to support EMC VNX7600, EMC VNX8000, EMC VMAX 20k and NETAPP FAS6240 arrays.</p> <p>i Note: This issue was fixed in StorNext 6.1. See 70854 on page 13 in the "Fixed Issues" section. Upgrade recommended.</p> <p>i Note: See also the new How to Modify the multipath.conf File for External Storage Arrays page released for StorNext 6.1.0.</p> <p>Workaround:</p> <p>Upgrade to StorNext 6.1.0 and modify the new EMC and/or NetApp array entries in the multipath.conf file and move it to the correct directory on your system.</p> <p>OR</p> <p>Manually update the multipath.conf file by hand for the array settings you need and move it to the correct directory on your system.</p>
70649	398896	<p>Issue:</p> <p>The Promise array path_checker definition needs to be updated in the reference multipath.conf file provided with our systems and firmware upgrades.</p> <p>i Note: This issue was fixed in StorNext 6.1. See 70649 on page 14 in the "Fixed Issues" section. Upgrade recommended.</p> <p>i Note: See also the new How to Modify the multipath.conf File for External Storage Arrays page released for StorNext 6.1.0.</p> <p>Workaround:</p> <p>Upgrade to StorNext 6.1.0.</p> <p>OR</p> <p>Edit the multipath.conf file, and change the path_checker value from its current setting (such as readsector0) to tur for the Promise Array definition.</p>

CR Number	SR Number	Description
70588	396789	<p>Issue:</p> <p>tsunami.log and service.log files can become large. are now rotated to prevent these logs and snapshot files becoming too large, and can significantly increase the time it takes to generate snapshot files.</p> <p>Note: This issue was fixed in StorNext 6.1. See 70588 on page 14 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <p>Upgrade to StorNext 6.1.0.</p> <p>OR</p> <p>Manually trim data from these log files when needed.</p>
70526	398061, 396903, 398206, 397045, 399232, and 401331	<p>Issue:</p> <p>Kernel Page Table Security Problems: CVE-2017-5753, CVE-2017-5715 and CVE-2017-5754.</p> <p>Note: This issue was fixed in StorNext 6.1. See 70526 on page 14 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <p>Upgrade to StorNext 6.1.0. There is no other workaround.</p>
71093	n/a	<p>Issue:</p> <p>Unused NIC ports can be reported as "down" in an admin alert during appliance firmware upgrades, which creates about 20 RAS emails.</p> <p>Note: This issue was fixed in StorNext 6.1. See Fixed Issues, Notes, and Enhancements – StorNext 6.x on page 5 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <p>Upgrade to StorNext 6.1.0. There is no other workaround.</p>


CR Number	SR Number	Description
72287, 72286, 72277	n/a	<p>Issue:</p> <p>After an appliance is upgraded to StorNext 6.1.0 (see Upgrade Firmware), sometimes one of the nodes on the NAS cluster is not able to re-join the cluster.</p> <p>Log on to the Appliance Controller and issue the following command:</p> <pre>nascluster show</pre> <p>If the node has not joined the NAS cluster, the following message is displayed in the controller output:</p> <pre>stornext initiated stop</pre> <p>Workaround:</p> <ol style="list-style-type: none"> From the command line of the system, issue the following command, where <myfs> is the mount path of the file system you would like to join to the cluster: <pre>nascluster join <myfs></pre> <p>The server should be rejoined to the cluster.</p> <ol style="list-style-type: none"> Run the following command: <pre>nascluster show</pre> <p>You should see the following in the output (in bold, green):</p> <pre>NAS Cluster IP: 10.10.100.100/eth0, Master: No, SNFS Root: <myfs>, Joined: Yes ...</pre>

CR Number	SR Number	Description
-----------	-----------	-------------

71644

Issue:

After the **Apply** button is clicked during a firmware upgrade on the StorNext GUI, there are certain instances where the firmware upgrade will fail to complete, and server node 2 may be upgraded, but node 1 will fail to upgrade causing a mismatch of StorNext releases between the nodes. When this happens, the **Firmware Upgrade** page in StorNext, an error message similar to the following may display when the firmware upgrade is being applied:



i Note: This issue was fixed in the upgrade code for StorNext 6.1.0. See [Fixed Issues, Notes, and Enhancements – StorNext 6.x on page 5](#) in the "Fixed Issues" section. Upgrade recommended.

Workaround:

The secondary node web service timeout needs to be increased from the default value 5 minutes/300 seconds) to 10 minutes/600 to account for additional processing in the pre-upgrade script. See [Change the component.properties Timeout Setting on page 47](#) for how to change this setting, and begin the firmware upgrade process again.

CR Number	SR Number	Description
-----------	-----------	-------------

71599

425749

Issue:

For systems being upgraded from StorNext 6.0.6 to 6.0.6.1 from a fresh install of the server with a StorNext 6.0.6 image, the firmware upgrade to support some expansion cards is not applied. Future card changes in the system will not get automatically configured since the firmware and autoconfig services which support some cards were disabled after the upgrade to StorNext 6.0.6.1.

You can verify this issue by logging into the command line of the server node (or on both servers for dual-server node systems) with the **stornext** user account and **stornext** user password, and issuing the following command from the command line of the system. If you see the "Virtual" (as shown here in green), your system is affected by this issue.

Example:

```
# /opt/DXi/util/seer VendorModel
# Virtual
```

i Note: This has been fixed in StorNext 6.1.0. See [71599 on page 13](#) in the "Fixed Issues" section. Upgrade recommended.

Workaround:

If the system has not been upgraded yet:

1. Log in to the command line of a server node. See [Launch the Command Line](#).
2. Issue the following command. This will correct the initial fresh install issue and avoid a subsequent upgrade issue:

```
cp -f /opt/DXi/theSeer/vendor.info.conf.13G
/opt/DXi/theSeer/vendor.info.conf
```

If the system was fresh installed with 6.0.6 and then upgraded to 6.0.6.1 the following needs to be done:

- a. Remove the three touch files:

```
rm /var/opt/DXi/fwboot-disable
```

```
rm /var/opt/DXi/hwdetect-onboot-disable
```

CR Number	SR Number	Description
		<pre>rm /var/opt/DXi/startup_autoconfig_disable</pre>
		<ol style="list-style-type: none"> b. Reboot the node so the Firmware Bundle updates will be performed and <code>autoconfig</code> and <code>hwdetect</code> will be enabled and run. <ol style="list-style-type: none"> 3. Repeat these steps for the other server node (dual-server node systems only).
71507	421334	<p>Issue:</p> <p>Firmware upgrade did not upgrade PHP on server nodes.</p> <p>i Note: This has been fixed for StorNext 6.1.0. See 71507 on page 13 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Log in to the command line of a server node (see Launch the Command Line.), and upgrade php56w: <pre># cd /scratch/saved/PHP/ # ls php56wUpgrade-5.6.33-1.tgz php56wUpgrade.tgz # tar xzf php56wUpgrade-5.6.33-1.tgz # rpm -Uhv php56w*.rpm ... Updating / installing... 1:php56w-common-5.6.33-1.w7 ##### [4%] ...</pre> 2. Remove the upgrade warning touch file: <pre># rm -f /var/DXi/upgrade_notify *.rpm *.tgz # rpm -q php56w php56w-5.6.33-1.w7.x86_64</pre> 3. Repeat for the other server node (for dual-node systems).

CR Number	SR Number	Description
71464	n/a	<p>Issue:</p> <p>Platform upgrade is currently allowed to the StorNext 6.0.6.1 release for systems currently configured and running StorNext 5.3.x and as a clustered NAS system running NAS 1.2.x. Even though StorNext upgrades, the NAS software does not upgrade. NAS is left running 1.2.5, and NAS/Controller software 2.0.1 is incorrectly staged for upgrade in <code>/var/upgrade</code>.</p> <p>Workaround:</p> <p>To upgrade NAS after the StorNext firmware upgrade:</p> <ol style="list-style-type: none"> 1. Delete the contents of <code>/var/upgrade</code> on both nodes (for dual-node systems). 2. Download the offline 1.4.1.3 NAS upgrade .rpm. 3. Copy the 1.4.1.3 NAS upgrade file to <code>/var/upgrade</code> on the NAS master and other node as well. 4. Upgrade NAS to 1.4.1.3 (see Upgrades for the offline upgrade procedure). 5. Delete the contents of <code>/var/upgrade</code> on both nodes (for dual-node systems). 6. Copy the offline 2.0.1 NAS/Controller software upgrade .rpm saved in <code>/scratch/saved/NAS</code> to <code>/var/upgrade</code>. 7. Upgrade the NAS/Controller software to 2.0.1 (see Upgrades for the offline upgrade procedure).
71394	0421219	<p>Issue:</p> <ul style="list-style-type: none"> Note: This issue affects R630 (CentOS7) systems only, not R520 (CentOS6) systems. Note: This issue was fixed in StorNext 6.0.6.1. See 71394 on page 14. Upgrade recommended. Note: For StorNext 6.0.6, the included the <code>MLNX_OFED</code> HBA driver for Mellanox Connect X-3 VPI InfiniBand/Ethernet cards was incompatible with the system kernel version. Because of this, the software failed to install, resulting in missing Mellanox software components needed to configure the card. <p>Workaround:</p> <ul style="list-style-type: none"> • For Systems prior to StorNext 6.0.6, the correct driver was shipped with StorNext, so no further action is necessary for Mellanox HBA card configuration. • For systems running StorNext 6.0.6, upgrade to StorNext 6.0.6.1 for the correct driver version for this card.

CR Number	SR Number	Description
71350	n/a	<p>Issue:</p> <ul style="list-style-type: none"> i Note: This issue applies to dual-server node systems only. i Note: This issue was fixed in StorNext 6.0.6.1. See 71350 on page 14 in the "Fixed Issues" section. Upgrade recommended. <p>For 6.0.6 systems, StorNext file systems (snfs) can potentially fail to mount on the secondary server node (typically server node 2) when over a billion files are present on the snfs. This could occur when cvfs is stopped and then restarted, after a hard stop of the system, HA failovers, or when the primary server node restarts/reboots.</p> <p>If the snfs had file locking "enabled" and was used to create a NAS cluster, then the server node will also be shown in the "not ready" state on the NAS cluster.</p> <p>Workaround:</p> <p>If your system was upgraded to a StorNext release earlier than 6.0.6.1 and your system exhibits this issue, here's the workaround.</p> <p>Do the following, on the server node that did not mount the file system:</p> <p>Mount the File System on the Server Node</p> <ol style="list-style-type: none"> Log in to node 2 (or node 1 if currently acting as the secondary node). See Launch the Command Line. Manually mount the file system on the server. <p>Example:</p> <pre>mount snfs1-fs</pre> <p>Re-Join the Server to the NAS Cluster, if Needed</p> <p>If the server is part of a NAS cluster:</p> <ol style="list-style-type: none"> Log into the NAS controller (see Access the Appliance Controller Console), and join that node to the cluster. <p>Example:</p> <pre>nascluster join /stornext/snfs1-fs</pre> <ol style="list-style-type: none"> Ensure this is joined by running the following command: <pre>nascluster show</pre>

CR Number	SR Number	Description
		<p>Look for the following (in green, bold) in the output (ellipses indicate additional information not shown). Note that snfs1-fs is the file system that did not mount before:</p> <pre data-bbox="565 443 1458 869"> ... NAS Cluster IP: 10.65.166.215/em2, Master: No, SNFS Root: /stornext/snfs1-fs, Joined: ... 1: 10.65.166.216 (Joined, MDC) test- bh3a.mdh.quantum.com 2: 10.65.166.215 (Joined, MDC) test- bh3b.mdh.quantum.com ... </pre>
71346	n/a	<p>Issue:</p> <p>i Note: This was fixed in StorNext 6.0.6.1. See 71346 on page 15 in the "Fixed Issues" section. Upgrade recommended.</p> <p>For systems running StorNext 5.3.x, and NAS 1.2.x, upgrading to 6.0.6 or later</p> <p>Workaround:</p> <p>Upgrade to Appliance Controller/NAS 1.4.1 or later prior to upgrading to StorNext 6.x or later.</p>

CR Number	SR Number	Description
71333	n/a	<p>Issue:</p> <hr/> <p> Caution: A system running StorNext 6.0.6 or later with a clustered NAS environment should not run NAS/Appliance Controller 1.3.0 or earlier. Running these configurations may produce unpredictable results.</p> <p>If you are running StorNext 6.0.6 with NAS 1.3.0 or earlier, you might experience one of the following scenarios:</p> <ol style="list-style-type: none"> 1. Controller configuration commands are not available for iSCSI and SNMP configuration, even though the Appliance Controller Doc Center states that these commands are available. 2. If your system is running with NAS enabled, NAS may not work properly and you will have to upgrade NAS locally with the latest Appliance Controller release. <p>Workaround:</p> <p>An upgrade to the Controller software is REQUIRED when running StorNext 6.0.6. See Appliance Controller Upgrades.</p>
71002	n/a	<p>Issue:</p> <p>After upgrading to StorNext 6.0.5 and StorNext 6.0.5.1 systems, a RAS ticket might be created for an IPMI failure.</p> <p> Note: This was fixed in StorNext 6.0.6. See 71002 on page 15 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <p>This RAS message is benign and can be ignored. Upgrade to StorNext 6.0.6 to correct this issue.</p>

CR Number	SR Number	Description
70902	n/a	<p>Issue:</p> <p>StorNext 6.0.6 requires the appliance controller to be upgraded for clustered NAS environments running Appliance Controller/NAS versions earlier than 1.4.1 before upgrading appliance firmware.</p> <p>If you attempt to upgrade the system to StorNext 6.0.6 prior to upgrading the controller, the appliance upgrade will fail, and the system will generate an "Operation failure" RAS ticket in the StorNext GUI on the server (for both server nodes on dual-node systems), the upgrade log will include a message about the upgrade attempt, and an admin alert will be generated.</p> <p>Example (dual-node system with NAS 1.3.0 installed):</p> <pre data-bbox="505 699 1458 1094"> 10 2018-02-22 15:12:22 PST 136 Upgrade Minimum NAS version not satisfied for a clustered system. Node 2: NAS clustering is enabled but installed NAS version is 1.3.0, NAS must be upgraded to version 1.4.1.x before this upgrade is allowed. A NAS 1.4.1.3 offline upgrade bundle has been placed in /scratch/saved/NAS/. </pre> <p>Workaround:</p> <p>Upgrade Appliance Controller/NAS to 1.4.1 or later prior to upgrading the system. See How to Upgrade on the Appliance Controller Documentation Center.</p>

CR Number	SR Number	Description
70338	n/a	<p>Issue:</p> <p>During a firmware upgrade, the secondary server node (typically Node 2) will upgrade successfully, but the other server node (typically Node 1) can become stuck on reboot, and unable to restart. If you don't have an SSH session open when this happens, you will notice that the server node will not respond for a long time. If you happen to have an SSH session open during the upgrade, you will see an error message similar to the following shown repeatedly on-screen:</p> <pre>[189.398035] dracut-initqueue[402]: Warning: dracut-initqueue timeout - starting timeout scripts</pre>
		<p>Workaround:</p> <p>The server node will have to be manually upgraded. Contact Contacting Quantum on page 58, and refer them to TSB# 3015.</p>

CR Number	SR Number	Description
70300	385040, 385789	<p>Issue:</p> <p>After a firmware upgrade, the system cannot boot.</p> <p>This does not need to be applied to all systems connected to storage arrays, only impacted systems, as described here.</p> <p>Systems Impacted:</p> <ul style="list-style-type: none"> This ONLY affects systems upgrading from StorNext 5.x releases to 6.x (prior to 6.0.5) This ONLY affects systems upgrading that are Fibre Channel-connected to BOTH NetApp and QXS arrays. <p>i Note: This issue was fixed in StorNext 6.0.5. See 70300 on page 17. If you are upgrading to StorNext 6.0.5, you will not see this issue, and do not have to perform this workaround.</p> <p>Messages on the boot screen and in the KDUMP (the trace of server boot activities on the sysfs) shows messages similar to:</p> <pre>Wrong diagnostic page; asked for 7 got 0</pre> <p>OR</p> <pre>Wrong diagnostic page; asked for 10 got 0</pre> <p>Workaround:</p> <p>DO NOT upgrade your system to StorNext 6.x before first applying the workaround. See Server Boot (QXS/NetApp).</p>

CR Number	SR Number	Description
70287	n/a	<p>Issue:</p> <p>StorNext upgrades can fail due to a slowdown of the StorNext HA file system (HAFS) when attempting to apply upgrade firmware to StorNext 6.x from supported StorNext 5.x and 6.0.x releases in the StorNext GUI. This only occurs when attempting to upgrade the system while jobs with a lot of metadata changes are running. This can happen when:</p> <ul style="list-style-type: none"> • The system is running with StorNext 6.x file systems with metadata archive enabled and heavy client activity • The system is running with StorNext 5.0.x file systems with the restore journal enabled and heavy client activity <p>Workaround:</p> <p>DO NOT upgrade your system to StorNext 6.x before first applying one of the workarounds below if your system is processing a lot of metadata jobs.</p> <ul style="list-style-type: none"> • To fix this issue, edit the <code>component.properties</code> file on one node of the system and increase the <code>"ws.apply.fw.upgrade.timeout"</code> value to 2 hours. <p>See Change Upgrade Timeout Setting.</p>
70178		<p>Issue:</p> <p>The <code>multipath.conf</code> file is overwritten for StorNext 5.3.1 through 6.0.5.x. Any custom-configured settings added to the file for those StorNext releases are available in <code>/etc/multipath.conf.bak</code>. The upgrade code added correct multipath settings to the <code>/etc/multipath.conf</code> file required for SAS-attached NetApp metadata arrays.</p> <p>Note: The overwrite issue was changed in StorNext 6.0.6. See 70178 on page 16 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <ol style="list-style-type: none"> 1. If you don't have a backup of your <code>multipath.conf</code> file, you will have to add the settings into the file. 2. If you have a backup of the custom multipath settings, copy the settings into the new <code>multipath.conf</code> file using a command-line editor like <code>vi</code>.

CR Number	SR Number	Description
70067	n/a	<p>Issue:</p> <p>CentOS 7 systems display log files with the journalctl command. Any commands executed as the rootsh user are logged to the system log file. A syslog filter sends all rootsh log commands to /var/log/rootshell. This causes excessive information to be captured in the journalctl log file, and makes the log file difficult to read.</p> <p>Workaround:</p> <p>There is a script to use which removes unnecessary command line output from the syslog file for CentOS7 systems running StorNext 6.0.5:</p> <p>See the journalctl Filter Script</p> <p>For StorNext releases prior to 6.0.5, contact Contacting Quantum on page 58.</p>
66799	322817, 342909, 362343	<p>Issue:</p> <p>On the StorNext GUI Reports > Hardware page, node1 system board information is shown but information for node2 is not after an upgrade to StorNext 5.x and later.</p> <p>i Note: This issue is fixed in StorNext 6.0.5. See 66799 on page 18 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <p>To fix this issue, create the /usr/adic/HAM/shared/HwMonStatus directory on the primary server node (typically node 1), and restart the hwmond service.</p> <p>See Create the HW Monitor Directory.</p>
53486	425247	<p>Srvclg RAS tickets are lost if the hafs (HA shared file system) is not mounted on the system.</p> <p>i Note: This issue was fixed in StorNext 6.2. See 53486 on page 12 in the "Fixed Issues" section. Upgrade recommended.</p> <p>Workaround:</p> <p>Upgrade to StorNext 6.2.</p>

Known Issues Workarounds and Additional Information

About firmware upgrades and pre-staged Appliance Controller upgrades for server nodes

When NAS is enabled, detected, and is configured on a sever node as part of a NAS cluster, and a firmware upgrade is attempted, an offline *.tar.gz file is copied into the /scratch/saved/NAS directory and is then unzipped and untarred into the /var/upgrade directory. Based on the installed version of NAS on the system, the appropriate valid upgrade bundle for NAS will be unpacked to /var/upgrade for the user to perform the upgrade. If a multi-version upgrade is required, then the next applicable version will be unpacked into /var/upgrade and you will be instructed to perform the cluster upgrade first and then retry the firmware upgrade. If the cluster state of the node cannot be determined then it looks like no versions are unpacked on the system and the upgrade is aborted (in this case, you would have to upgrade the cluster manually using instructions in the Appliance Controller Doc Center (see [Appliance Controller Upgrades](#)), and then retry the firmware upgrade.

If NAS clustering is not configured, the firmware upgrade will perform the direct upgrade to the current version of Appliance Controller shipped with the firmware upgrade.

74202 – StorNext services can fail to stop completely when NAS is Configured and StorNext services are stopped

How might this be seen?

There are two basic scenarios where this might be exposed:

Scenario # 1

When StorNext stops, it signals NAS to vacate the StorNext file systems that NAS is using. However, if these file systems remain available for too long, the Appliance Controller will attempt to automatically restart NAS services. This is what happens, in order:

1. Restarting NAS services can resume use of StorNext resources.
2. Resuming StorNext resources interferes with attempts to stop StorNext.
3. Continued use of these StorNext resources can result in mounted StorNext file systems that do not have a running fsm process which blocks access to the StorNext file system.

If experience this on your system, temporarily stop NAS services:

1. Log in to the command line of the server (as "sudo rootsh"):

Launch the Command Line With an SSH Utility

- a. Open an SSH connection to the appropriate server and use the IP address assigned to the node on the Management/LAN Client network.

a. Log in to the command line using the following credentials:

- User name: **stornext**
- Password: **<stornext user account password>**

i **Note:** "password" is the default password for the stornext user account. If the password has been changed, use the current password.

b. Enter **sudo rootsh** to gain root user access.

c. Enter the password for the **stornext** user account again.

2. Stop NAS by executing the following command:

```
/usr/cvfs/lib/snnas_control stop
```

3. Stop the Appliance Controller by executing the following command:

```
/usr/local/quantum/bin/sml_service_tool stop snnas_controller
```

4. When you have finished Restart NAS and the Appliance Controller:

a. Start NAS by executing the following command:

```
/usr/cvfs/lib/snnas_control start
```

b. Start the Appliance Controller by executing the following command:

```
/usr/local/quantum/bin/sml_service_tool start snnas_controller
```

Scenario # 2

If you need to stop the system during a maintenance window. This is an extension of **Scenario #1**, but NAS must vacate the file systems for an extended period of time. If NAS services are resumed when they should be suspended, those service processes may interfere with maintenance operations.

To affect both scenarios, timers exist that you can adjust to extend the amount of time required between the **stornext stop** operation and before NAS services resume. The timers are controlled by the following **Controller Registry** values:

- **nas.heartbeat.check_state_secs**
- **stornext_service.stop_period**

To see what values are currently assigned, enter:

```
su sysadmin -c 'reg show nas.heartbeat.check_state_secs'
```

to display the setting of the frequency interval of the NAS heartbeat state check, or enter:

```
su sysadmin -c 'reg show stornext_service.stop_period'
```

to display the duration that StorNext services are stopped.

i Note: By default, both timers are set to 120 (in seconds=2 minutes).

Workarounds

Workaround #1

If StorNext fails to stop due to NAS resource usage, do the following:

1. Log in to the command line of the server:

Launch the Command Line With an SSH Utility

- a. Open an SSH connection to the appropriate server and use the IP address assigned to the node on the Management/LAN Client network.
- a. Log in to the command line using the following credentials:
 - User name: **stornext**
 - Password: **<stornext user account password>**

i Note: "password" is the default password for the stornext user account. If the password has been changed, use the current password.

- b. Enter **sudo rootsh** to gain root user access.
 - c. Enter the password for the **stornext** user account again.
2. Verify that the StorNext file systems are mounted:

```
grep cvfs /proc/mounts
```

If any StorNext file systems are not mounted, mount them now.

3. Mount the HA shared file system:

```
mount /usr/adic/HAM/shared
```

4. Start fsm for each StorNext file system (shown below as "<FSNAME>") mounted on the server node:

```
cvadmin -e "start <FSNAME> on localhost"
```

5. Stop StorNext services:

For CentOS7 systems, enter:

```
systemctl stop cvfs
```

For CentOS6 systems, enter:

```
service cvfs stop
```

6. To stop NAS, enter:

```
/usr/cvfs/lib/snnas_control stop
```

7. To stop the Appliance Controller, enter:

```
/usr/local/quantum/bin/sml_service_tool stop snnas_controller
```

When you are ready, restart the Appliance Controller and NAS services:

8. To restart the Appliance Controller, enter:

```
/usr/local/quantum/bin/sml_service_tool start snnas_controller
```

9. To restart NAS, enter:

```
/usr/cvfs/lib/snnas_control start
```

Workaround #2

For **Scenario #1**, Quantum recommends that you extend the time period to a value of 10 minutes (600 seconds) for both variables. Set a value that exceeds the amount of time necessary for typical StorNext shutdown.

To set new values:

1. Log in to the command line:

Launch the Command Line With an SSH Utility

- a. Open an SSH connection to the appropriate server and use the IP address assigned to the node on the Management/LAN Client network.
- a. Log in to the command line using the following credentials:
 - User name: **stornext**
 - Password: **<stornext user account password>**

i Note: "password" is the default password for the stornext user account. If the password has been changed, use the current password.

- b. Enter **sudo rootsh** to gain root user access.
 - c. Enter the password for the **stornext** user account again.
2. Change the frequency interval of the NAS heartbeat state check to 600 seconds (10 minutes). Enter:

```
su sysadmin -c 'reg set nas.heartbeat.check_state_secs 600'
```

i Note: This is run with a temporary log in to the Appliance Controller as the **sysadmin** user.

3. Change the duration that StorNext services are stopped to 600 seconds (10 minutes). Enter:

```
su sysadmin -c 'reg set stornext_service.stop_period 600'
```

4. **i Note:** This is run with a temporary log in to the Appliance Controller as the **sysadmin** user.

Workaround #3

For **Scenario #2**, to eliminate the confusion of setting a long duration for a maintenance window, and then having to change the values back again to a shorter value again, you can simply stop NAS services and the Appliance Controller, and then restart these services when you are finished.

To stop the services and restart them:

1. Log in to the command line:

Launch the Command Line With an SSH Utility

- a. Open an SSH connection to the appropriate server and use the IP address assigned to the node on the Management/LAN Client network.
- a. Log in to the command line using the following credentials:
 - User name: **stornext**
 - Password: **<stornext user account password>**

i **Note:** "password" is the default password for the stornext user account. If the password has been changed, use the current password.

- b. Enter **sudo rootsh** to gain root user access.
 - c. Enter the password for the **stornext** user account again.
2. Stop NAS by executing the following command:

```
/usr/cvfs/lib/snnas_control stop
```

3. Stop the Appliance Controller by executing the following command:

```
/usr/local/quantum/bin/sml_service_tool stop snnas_controller
```

When you are ready, restart the Appliance Controller and NAS services:

4. To restart the Appliance Controller, enter:

```
/usr/local/quantum/bin/sml_service_tool start snnas_controller
```

5. To restart NAS, enter:

```
/usr/cvfs/lib/snnas_control start
```

6. Change the settings for the NAS heartbeat state check and the duration that StorNext services are stopped according to [Workaround #2 on the previous page](#).

Workaround #4

An alternate method for maintenance windows, instead of implementing workaround #3 for **Scenario #2**, you can choose a much longer value if you need the NAS heartbeat state check interval and

StorNext services stopped for a long period of time (for example, during a planned maintenance window in a data center). While this is another option, the downside of this is that you will need to know how much downtime to expect maintenance to take, and then set these times to be less than what is needed for that maintenance. It is also difficult to extend the maintenance period reliably once those timers have started. In addition, you will most likely have to reset the values back to their original settings when the maintenance window is complete, which could be a hassle.

Here is one example which extends the time values to 2 hours (7,200 seconds):

1. Log in to the command line:

Launch the Command Line With an SSH Utility

- a. Open an SSH connection to the appropriate server and use the IP address assigned to the node on the Management/LAN Client network.
- a. Log in to the command line using the following credentials:
 - User name: **stornext**
 - Password: **<stornext user account password>**

i Note: "password" is the default password for the stornext user account. If the password has been changed, use the current password.

- b. Enter **sudo rootsh** to gain root user access.
 - c. Enter the password for the **stornext** user account again.
2. Change the frequency interval of the NAS heartbeat state check to 7,200 seconds (10 minutes). Enter:

```
su sysadmin -c 'reg set nas.heartbeat.check_state_secs 7200'
```

i Note: This is run with a temporary log in to the Appliance Controller as the **sysadmin** user.

3. Change the duration that StorNext services are stopped to 7,200 seconds (10 minutes). Enter:

```
su sysadmin -c 'reg set stornext_service.stop_period 7200'
```

4. **i Note:** This is run with a temporary log in to the Appliance Controller as the **sysadmin** user.

Change the `component.properties` Timeout Setting

1. Log in to the command line of the primary server node (typically node 1). See [Launch the Command Line \(for Xcellis Workflow DirectorGen 2\)](#) or [Launch the Command Line \(for R630 systems\)](#).
2. Open the `component.properties` file with the vi editor:

```
vi /usr/adic/gui/config/component.properties
```

Pro Tip #1: Edit With vi

The vi editor is included in CentOS. If you are unfamiliar with using vi to open, edit (copy/paste/delete), exit with or without saving changes, use the following Google search URL to find related links for editing with vi:

<https://www.google.com/search?q=edit+with+vi>

3. Scroll down to the web services timeouts section. It should look similar to (note the default firmware upgrade value is 300 seconds - highlighted in green):

```
#
# Web Services Timeouts (seconds)
#
ws.convert.to.ha.timeout=300 ws.apply.fw.upgrade.timeout=300
ws.get.system.id.timeout=10 ws.ntp.timeout=30 ws.datetime.timeout=30
ws.network.config.timeout=180 ws.network.hapeer.timeout=30
ws.logging.level.timeout=30 ws.capture.state.timeout=600
ws.job.close.timeout=30 ws.nas.close.timeout=60

#
```

4. Move the cursor to sit over the number 3, enter **<SHIFT>+*** and then **R** (for "replace") and enter the number **6** to change the ws.apply.fw.upgrade.timeout= value from 300 seconds (5 minutes) to 600 seconds (10 minutes):

```
ws.apply.fw.upgrade.timeout=600
```

5. Save the changes to the file by entering **SHIFT:wq** to write the file and quit the vi editor.

The setting change will be replicated to the other server node in 2-3 minutes after the file change is saved, so this only has to be applied to server node 1.

6. Run the **Firmware Upgrade** process again:

i Note: Even though the upgrade failed previously, the firmware upgrade files are still available for the upgrade.

- a. Select either of the uploaded files click **Apply**. In the case of a single file, select the one file uploaded and click **Apply**.
- b. Click **Yes** to confirm the installation. For two-file upgrades, both firmware files are applied to the system during the upgrade.

The upgrade should complete after the timeout setting is changed. However, if the firmware upgrade fails again, contact [Contacting Quantum on page 58](#) for assistance with manually upgrading the node that failed to upgrade.

Server Boot (QXS/NetApp)

This **does not** need to be applied to **all** systems connected to storage arrays, only impacted systems, as described here.

Systems Impacted:

- This **ONLY** affects systems upgrading from StorNext 5.x releases to 6.x (prior to 6.0.5)
 - This **ONLY** affects systems upgrading that are Fibre Channel-connected to **BOTH** NetApp and QXS arrays.
- i** **Note:** This issue was fixed in StorNext 6.0.5. See [70300 on page 17](#).
If you are upgrading to StorNext 6.0.5, you will not see this issue, and do not have to perform this workaround.

BEFORE applying the upgrade to StorNext 6.x, do the following:

Log in to the system command line

1. Open an SSH connection to the appropriate server and use the IP address assigned to the node on the Management or LAN Client network, or use the Service Port IP address, if connected to the Service Port.

Service Port IP address:

- 10.17.21.1 (dual-server node systems)

1. Log in to the command line using the following credentials:
 - User name: **stornext**
 - Password: **<stornext user account password>**

i **Note:** "password" is the default password for the stornext user account. If the password has been changed, use the current password.

2. Enter **sudo rootsh** to gain root user access.
3. Enter the password for the **stornext** user account again.

Edit the Grub Configuration File

1. Create a backup copy of the CentOS7 grub.d/40_custom configuration file. Enter the following:

```
cp /etc/grub.d/40_custom /scratch/40_custom.orig
```

2. Create a backup copy of the CentOS7 grub.cfg file. Enter the following:

```
cp /boot/grub2/grub.cfg /boot/grub2/grub.cfg.orig
```

3. Edit the /etc/grub.d/40_custom file using the vi editor, included in CentOS7.x. Enter the following:

```
vi /etc/grub.d/40_custom
```

4. To enter vi "insert" mode, enter i.
5. Look for the lines beginning with "linux16" as shown in bold, blue below, and add "<space>rdblacklist=ses rdblacklist=enclosure" to the end of those lines, as shown in green in the following example. (Ellipses "..." included below denote code prior to and after the "linux16" sections.):

```
...
```

```
linux16 /vmlinuz-3.10.0-514.6.2.el7.qtm.x86_64 root=UUID=3d9116b2-c5de-41fd-  
bd38-2042c0f8ebad ro console=tty1 vconsole.keymap=us crashkernel=768M@0M  
vconsole.font=latarcyrheb-sun16 selinux=0 enforcing=0 intel_idle.max_cstate=0  
processor.max_cstate=1 elevator=deadline transparent_hugepage=never quiet  
LANG=en_US.UTF-8 rdblacklist=ses rdblacklist=enclosure
```

```
initrd16 /initramfs-3.10.0-514.6.2.el7.qtm.x86_64.img  
}
```

```
menuentry 'Xcellis SOL' --class centos --class gnu-linux --class os --  
unrestricted {  
load_video  
set gfxpayload=keep  
insmod gzio  
insmod part_msdos  
insmod ext2  
set root='hd0,msdos1'  
if [ x = xy ]; then  
search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
```

```

efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1 --hint='hd0,msdos1' a0467663-
8cf0-48be-b967-d5ef7eb486e2
else
search --no-floppy --fs-uuid --set=root a0467663-8cf0-48be-b967-d5ef7eb486e2
fi
linux16 /vmlinuz-3.10.0-514.6.2.el7.qtm.x86_64 root=UUID=3d9116b2-c5de-41fd-
bd38-2042c0f8ebad ro console=ttyS1,115200n8 vconsole.keymap=us
crashkernel=768M@0M vconsole.font=latarcyrheb-sun16 selinux=0 enforcing=0
intel_idle.max_cstate=0 processor.max_cstate=1 elevator=deadline transparent_
hugepage=never quiet LANG=en_US.UTF-8 rdblacklist=ses rdblacklist=enclosure
initrd16 /initramfs-3.10.0-514.6.2.el7.qtm.x86_64.img
}
...

```

6. To exit "insert" mode and return to "command mode", press <Esc>.
7. To save changes and exit vi, enter :wq and press <Enter>. You will return to the command prompt. (To exit vi without saving changes, enter :q! and press <Enter>.)
8. Repeat [Log in to the system command line on page 49](#) and [Edit the Grub Configuration File on the previous page](#) for the other server node.

Edit the Blacklist Configuration File and Enable the Changes

1. Enter the following to create a backup copy of the blacklist configuration file:

```
cp /etc/modprobe.d/blacklist.conf /scratch/blacklist.conf.orig
```

2. Enter the following to edit the /etc/modprobe.d/blacklist.conf file using vi:

```
vi /etc/modprobe.d/blacklist.conf
```

3. To edit, change to the vi "insert" mode. Enter i.
4. Add a separate line for "rdblacklist ses" and another line for "rdblacklist enclosure" to the file, as shown in green in the following example.

```

blacklist rivafb
blacklist savagefb
blacklist sstfb
blacklist neofb

```

```
blacklist tridentfb
blacklist tdfxfb
blacklist virgefb
blacklist vga16fb
blacklist viafb

# ISDN - see bugs 154799, 159068
blacklist hisax
blacklist hisax_fcpcipnp

# sound drivers
blacklist snd-pcsp

# I/O dynamic configuration support for s390x (bz
#563228)
blacklist chsc_sch

blacklist sb_edac
blacklist edac_core
blacklist ses
blacklist enclosure
```

5. To exit "insert" mode and return to "command mode", press <Esc>.
6. To save changes and exit vi, enter :wq and press <Enter>. You will return to the command prompt. (To exit vi without saving changes, enter :q! and press <Enter>.)
7. To enable the changes, enter:

```
dracut -f
```

8. Continue to [Reboot the system on the next page](#).

Workaround: System Already Upgraded

If the system has been upgraded and cannot boot, do the following to recover the system:

1. Detach/disconnect the FC ports from the nodes or disable the FC zoning configuration for the QXS array(s).
2. Boot the system and enter the following to stop cvfs:

```
systemctl stop cvfs
```

3. Edit the grub configuration file for your version of CentOS. See [Edit the Grub Configuration File on page 50](#).
4. Edit the blacklist configuration file. See [Edit the Blacklist Configuration File and Enable the Changes on page 51](#).
5. Reattach or re-enable FC zoning configuration for the QXS array(s).
6. Continue to [Reboot the system below](#).

Reboot the system

1. Reboot the server you are already logged into.

```
systemctl reboot
```

Continue to [Step 3](#), for single-node systems. For dual-node systems, continue to the next step.

2. Log in to the other server node. See [Log in to the system command line on page 49](#).

```
systemctl reboot
```

3. For dual-node systems, repeat [Step 3](#) for server node 2. The process is complete.

Change Upgrade Timeout Setting

BEFORE applying the upgrade to StorNext 6.x, do the following:

Log in to the system command line

1. Open an SSH connection to the appropriate server and use the IP address assigned to the node on the Management or LAN Client network, or use the Service Port IP address, if connected to the Service Port.

Service Port IP address:

- 10.17.21.1 (dual-server node systems)

1. Log in to the command line using the following credentials:
 - User name: **stornext**
 - Password: **<stornext user account password>**

i **Note:** "password" is the default password for the stornext user account. If the password has been changed, use the current password.

2. Enter **sudo rootsh** to gain root user access.
3. Enter the password for the **stornext** user account again.
1. Change to the config directory:

```
cd /usr/adic/gui/config
```

2. Copy the component.properties file:

```
cp component.properties component.properties.orig
```

3. Edit the file using vi:

```
vi component.properties
```

4. Scroll down the file until you get to the **# Web Services Timeouts (seconds)** section, and the upgrade timeout line shown in bold text. The current value should be similar to the value shown in bold, green:

```
#  
# Web Services Timeouts (seconds)  
#  
ws.convert.to.ha.timeout=300  
ws.apply.fw.upgrade.timeout=300  
ws.get.system.id.timeout=10  
ws.ntp.timeout=30  
ws.datetime.timeout=30  
ws.network.config.timeout=180  
ws.network.hapeer.timeout=30  
ws.logging.level.timeout=30  
ws.capture.state.timeout=600  
ws.job.close.timeout=30  
ws.nas.close.timeout=60
```

5. To edit, change to the vi "insert" mode. Enter i.
6. Change the value to **1200**:

```
#  
# Web Services Timeouts (seconds)  
#  
ws.convert.to.ha.timeout=300  
ws.apply.fw.upgrade.timeout=1200  
ws.get.system.id.timeout=10  
ws.ntp.timeout=30  
ws.datetime.timeout=30  
ws.network.config.timeout=180  
ws.network.hapeer.timeout=30  
ws.logging.level.timeout=30  
ws.capture.state.timeout=600  
ws.job.close.timeout=30  
ws.nas.close.timeout=60
```

7. To exit "insert" mode and return to "command mode", press <Esc>.
8. To save changes and exit vi, enter :wq and press <Enter>. You will return to the command prompt. (To exit vi without saving changes, enter :q and press <Enter>.)
9. Continue to [Reboot the system on the next page](#).

Reboot the system

1. Reboot the server you are already logged into.

```
systemctl reboot
```

Continue to [Step 3](#), for single-node systems. For dual-node systems, continue to the next step.

2. Log in to the other server node. See [Log in to the system command line on page 53](#).
3. Reboot the other server, node 2.

```
systemctl reboot
```

4. Repeat [Step 3](#) for server node 2.

The process is complete.

journalctl Filter Script

Step 1: Upgrade to StorNext 6.0.5

i Note: StorNext 6.0.5 provides the script, so you need to upgrade to StorNext 6.0.5 to access the file.

See [Upgrade a Gateway System \(Upgrade Firmware\) Using the Service Menu](#).

Step 2: Log in to the System Command Line

1. Open an SSH connection to the appropriate server and use the IP address assigned to the node on the Management or LAN Client network, or use the Service Port IP address, if connected to the Service Port.

Service Port IP address:

- 10.17.21.1 (dual-server node systems)

1. Log in to the command line using the following credentials:

- User name: **stornext**
- Password: **<stornext user account password>**

i Note: "password" is the default password for the stornext user account. If the password has been changed, use the current password.

2. Enter **sudo rootsh** to gain root user access.
3. Enter the password for the **stornext** user account again.
4. Continue to [Step 3: Run the Journal Filter Script and View the Script Help on the next page](#).

Step 3: Run the Journal Filter Script and View the Script Help

- Enter the following to use the script to view the log output in a way that filters out all the rootsh commands:

```
/opt/DXi/scripts/journalctl_filter.sh -f
```

- Enter the following to view the script's help file:

```
/opt/DXi/scripts/journalctl_filter.sh --help
```

i Note: The `journalctl_filter.sh` output can also be saved to a file.

Create the Missing Hardware Monitor Directory

Do the following to check if the directory is missing, and re-create the directory if it is and restart the `hwmond` service:

1. Open an SSH connection to the appropriate server and use the IP address assigned to the node on the Management or LAN Client network, or use the Service Port IP address, if connected to the Service Port.

Service Port IP address:

- 10.17.21.1 (dual-server node systems)
1. Log in to the command line using the following credentials:
 - User name: `stornext`
 - Password: `<stornext user account password>`

i Note: "password" is the default password for the `stornext` user account. If the password has been changed, use the current password.

2. Enter `sudo rootsh` to gain root user access.
3. Enter the password for the `stornext` user account again.
4. Enter the following on the command line:

```
cd /usr/adic/HAM/shared/HwMonStatus
```

If you see the following message, the directory does not exist. Continue to the next step.

```
bash: cd: /usr/adic/HAM/shared/HwMonStatus: No such file or directory
```

5. Enter the following on the command line to create the directory:

```
mkdir /usr/adic/HAM/shared/HwMonStatus
```

6. Enter the following to verify the directory has been created :

```
cd /usr/adic/HAM/shared/HwMonStatus
```

7. Change the permission of the hwmond service to "read-only":

```
chmod -r 775 HwMonStatus
```

8. Restart the hwmond service.

```
systemctl restart quantum-hwmond.service
```

Contacting Quantum

For further assistance, 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>