

# StorNext 7.2.0 Artico Archive Gateway

## Release Notes

---

Product/Software Release Date	July, 2016
-------------------------------	------------

---

Contents

About the Artico Archive Gateway .....	2
Training and Documentation Resources .....	2
Learn more about StorNext 7.2.0 .....	2
What StorNext Releases are compatible with my system? .....	3
Upgrade StorNext Software and System Firmware .....	3
General Notes .....	3
Enhancements, Fixed Issues, and Notes .....	3
Known Issues .....	8
Quantum Appliance Licenses .....	44
Contacting Quantum Support .....	44

---

© 2024 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. Myriad, 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.

---

# About the Artico Archive Gateway

The Artico™ Archive Gateway provides a flexible, low-cost NAS archive appliance, with the ability to scale to store petabytes of data as demand grows. Powered by StorNext® advanced data management policies, Artico stores and retains data in a tiered storage system that intelligently places the data on the right technology at the right time, both on-premise or in the cloud. Combined with StorNext-enabled policies and the data migration tools of Rocket Arkivio Autostor software, Artico enables users to move files seamlessly from primary storage to less expensive storage while maintaining full access to all files, providing the lowest TCO model for storing, accessing and protecting data.

---

## Training and Documentation Resources

[Artico \(R630\) 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.

---

## Learn more about StorNext 7.2.0

### 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.

### Learn More on the Web...

Learn more about StorNext 7 and Scale-Out Storage and take a look at the [StorNext 7 Documentation Center](#), which includes training and documentation resources for StorNext 7, including StorNext software

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

**i Note:** Firmware files are large – over 12 GB (or larger) for releases where both files are required (typical), so plan time to download the file(s) required for the upgrade. The **/scratch** file system requires at least 20 GB of free space available before you perform the upgrade.

Before upgrading your system, refer to [Known Issues on page 8](#). 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.

---

## Enhancements, Fixed Issues, and Notes

This section lists the fixed issues, enhancements and notes for StorNext 7.2.0.

## Fixed Issues

ID	Description
SNXT-122	fslocate -k file_key fails on some files in mdarchive query
SNXT-139	Improve fsmodclass & fsexpcopy man pages to include info about rebuild candidate list
SNXT-179	Persistent LOG_INFO message in query_ldap() routine produces a large number of entries in /var/log/messages
SNXT-193	sgadd: when --template used for stripe group with exclusive affinity, sgadd issues an error message and exits: *Error*: You must specify data, metadata, or both
SNXT-194	cvadmin -H remote_host -e \"disks refresh\" do not trigger on remote_host but on local_host
SNXT-195	cvadmin -H remote_host -e \"disks\" reports wrongly \"Disks known to FSMPPM on host \"local_host\" instead of \"remote_host\"
SNXT-200	CVFS Client hits \"VOP error on cookie <> cvfs error \"IO error\" (0x3) device <> after sgoffload
SNXT-210	TSM recreate tape: misleading error messages , e.g. Some files could not be recreated from alternate media due to segment size mismatches with source media
SNXT-317	Dynamic Library Pooling ( DLP ) doesnt cope well when one library of the DLP Pool runs out of available media
SNXT-335	Storage Manager Restore Procedures: missing crucial tips and warning for pitfall
SNXT-377	fsretrieve can error when object are pulled back from archive tier to cool tier
SNXT-390	A tiering job or policy should default to only affect files
SNXT-446	Fsmppm has sigabort due to accept failure in libevent code
SNXT-457	zombie fsprobe processes pile up on MDCs due to file descriptor leak in parse_json_cli()
SNXT-475	Bug 78307 - modified files that are removed before being re-stored cannot be recovered
SNXT-493	fsmecopy recreate fails indicating there are no files to copy when there are.
SNXT-536	snprobe -a can stall in environments with large numbers of DLC connections
SNXT-581	sncompare causes false positive with SDISK and Tape Copy will suggest to remove ALL_COPIES made on All Entries
SNXT-594	fsmecopy reconciliation to help with - Deleting file during store causes database to show endtime as active.
SNXT-602	cvupdatefs() should not run 'sn_dmap -c' on LUN's that have been upgraded from \"Short Labelled\"

ID	Description
SNXT-620	XWE Gateway Agent is sending invalid file system metrics for MDC file system when affinity/tiering is enabled
SNXT-643	Fsm Assert Panic - ASSERT failed "!set    (rptr->iref_flags & IREF_RETURN_ATTRS) == 0
SNXT-661	CCTL logs flooding the nssdbg
SNXT-670	snaccess test functions incorrectly states subdirectory access is allowed
SNXT-678	Bug 73290 - File specific xrep error should not fail entire mover request seen on a XWD running SN7.0.3
SNXT-684	(CVE-2023-1945) - firefox-91.13.0-1.el7.centos - Unexpected data returned from the Safe Browsing API could have led to memory corruption and a potentially exploitable crash.
SNXT-685	add mechanism to disable hardcoded mover obj store retry logic
SNXT-688	(CVE-2022-37434) zlib: heap-based buffer over-read and overflow in inflate() in inflate.c via a large gzip header extra field
SNXT-693	(CVE-2023-0767) NSS Tools nss: Arbitrary memory write via PKCS
SNXT-694	escaped chars can lead to >255 name file, silently preventing store, failing policy involving subsequent candidates
SNXT-705	7.1.1 upgrade rendered Mellanox Infiniband devices undetectable.
SNXT-714	Compatibility Guide - Support for IBM TS4300 + IBM LTO 9 HH Fibre channel drives
SNXT-731	snstatd appears to leak memory
SNXT-734	/opt/DXi/scripts/collect-local: line 557: /usr/bin/pg_dump: No such file or directory
SNXT-740	transcribe operation to LTFS tape varies media to suspect if the destination media reach EOT/EOM
SNXT-766	cvfsck -E "FS" fails and coredump if "FS" has vacant SG
SNXT-768	fsexport -b cannot copy contents from LTO to LTOW (WORM) Media
SNXT-774	Without "exec" and "suid" XML tags in nss_cctl.xml, the values are set to false by default
SNXT-775	Man page update cvfsck -A only showing results in read/write mode
SNXT-819	ConvertAllRsvdExtents() triggers: PANIC: "BUG: unable to handle kernel NULL pointer dereference at 0000000000000004"
SNXT-878	Bug 76850 - FSM selecting a Remote NDIS Compatible Device causing connection issues (Clients interlocking because registered with same internal IP)

ID	Description
SNXT-881	Schema update failure when upgrading from SN pre-5.2.0 to post 6.4.0 in 7.1 code
SNXT-891	Documentation: Update the text for "Online" library action to reflect that it updates the device path and the mailbox slot count
SNXT-912	Documentation: Need to document when dbupgrade is run and what changes are being made to avoid unexpected downtimes or delays during Firmware upgrades
SNXT-937	sntier - Make sizes human readable and clarify what is meant by size in sntier command outputs. It is incredibly difficult to keep translating numbers like 47398723478732348 into terabytes or how many "seconds old" something is
SNXT-946	tiering - not enough info provided when a job failed due to estimate bigger than free space in pool
SNXT-956	Tracking Bug: Radar 119052712 Xsan clients fail to mount at boot
SNXT-1005	CVFS Client is leaking kmalloc-96 objects
SNXT-1019	Segmentation Fault in sn_dirent_key_init dir_btree.c:79
SNXT-1062	CVFS PANIC: _cvRelLock: unlocking a non-locked lock
SNXT-1116	/usr/adic/PSE/bin/query_tape_libraries is missing suspect count in output
CON-10602	MDC node with tiering recently disabled still shows the 'Pools' feature, which is confusing to customers
CON-10604	Unified Connector - Alma 8.7 and Rocky 8.7 are not supported by UC - minimum is 9
CON-10619	StorNext 7.1 + UUI/USB/E/API GW 7.1.1: 500 error on Alerts page when with any action
CON-10689	MacOS Unified Connector Package installer is blocked because it is not from an identified developer
CON-10701	Metrics page is not displaying graphs and statistics due to exceeding qustats data length
CON-10737	Adding second separate MDC node or HA pair pointing to same FS name server can lead to unhealthy MDC nodes status.
CON-10791	USUI – Portswigger scan flags /about/api RESPONSE for sensitive version and OS info
CON-10804	USUI - Nodes - Health status is changed incorrectly from times to times for metadata controller nodes
CON-10902	Linking an existing mystornext portal account to the UUI during user-register fails due to CORS policy

ID	Description
CON-10988	USUI - Support of large integer value(Petabytes) in license report generated and downloaded from the UUI
CON-11015	email test fails if "Use Authentication" is enabled
CON-11051	Unified Connector Installation Challenges on Windows
CON-11082	USUI-UC On newer MacOS can fail to get the hostname if TCPv4/v6 dual stack is active
CON-11168	Enhancement request - UUI - display details about the required password complexity
CON-11172	Health status of H-Series nodes become 'indeterminate' multiple times each day.

---

# Known Issues

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

CR Number	Customer SR Number	Description
SNXT-118	N/A	<p><b>confignetdef.sh</b> incorrectly creates a bond1 profile in the <b>ifcfg</b> file with an incorrect DEVICE entry.</p> <p>You may see an error like the following:</p> <pre>2023-01-18-09:13:46.244095 leo kernel: CVFS: Cannot get IP address of NIC with MAC 24-6e-96-a2-a2-59 error "Not found" (2</pre> <p><b>Workaround:</b></p> <p>This issue has been fixed in StorNext 7.1.1. See <a href="#">SNXT-116</a> in the 7.1.1 Fixed Issues.</p> <p>To fix this issue on StorNext prior to 7.1.1, open the <b>/etc/sysconfig/network-scripts/ifcfg-bond1</b> file. In the in first line of the file, change the following entry:</p> <pre>DEVICE=bond0:2</pre> <p>to:</p> <pre>DEVICE=bond1:2</pre>

---



CR Number	Customer SR Number	Description
SNXT-389	N/A	<p>The StorNext rpm database can become corrupted if an rpm operation or a query is unexpectedly interrupted by a signal, reboot, or a system crash on StorNext NAS-clustered systems or the Unified User Interface (UUI) software runs.</p> <p>Querying the database might result in errors and report that an installed package is not installed.</p>

**Example:**

```
$ rpm -q snfs-server
error: rpmdb: BDB0113
Thread/process7492/140149665155136 failed: BDB1507
Thread died in Berkeley DB library
error: db5 error(-30973) from dbenv->failchk:BDB0087
DB_RUNRECOVERY: Fatal error, run databaserecovery
error: cannot open Packages index using db5 - (-30973)
error: cannot open Packages database in /var/lib/rpm
error: rpmdb: BDB0113
Thread/process7492/140149665155136 failed: BDB1507
Thread died in Berkeley DB library
error: db5 error(-30973) from dbenv->failchk:BDB0087
DB_RUNRECOVERY: Fatal error, run databaserecovery
error: cannot open Packages database in /var/lib/rpm
package snfs-server is not installed
```

These query errors could result in the inability to start StorNext or other Quantum services.

On rare occasions, you might see corruption of the rpm database after are boot of the system while an rpm activity occurred. StorNext, StorNext NAS, and the UUI invoke periodic rpm queries that could be running when a reboot occurs that could lead to corruption. The rpm database corruption is more likely to be seen in the event of an ungraceful reboot resulting from a kernel panic, power outage, or SMITH reset, or upon sending the kill signal to a running rpm command which can happen via systemd as part of a normal shutdown/reboot sequence.

**Workaround:**

1. Log in to the command line of a server node.
2. Run the following command to repair the rpm database:



```
rpmdb --rebuilddb
```

CR Number	Customer SR Number	Description
		<p>3. After you rebuild the database, run the following query:</p> <pre>rpm -q snfs-server</pre> <p>The server version will be displayed, which verifies the database has been repaired:</p> <pre>snfs-server-7.1.1-91E.RedHat7.x86_64</pre>
SNXT-951	N/A	<p>Beginning with StorNext 7.2.0, if your Xcellis (R630) system is configured with a Mellanox CX-3 card, then you might not be able to upgrade the driver for the card.</p> <p><b>Workaround:</b></p> <p>There is no workaround for this issue; Mellanox has ended support for the card driver.</p>
SNXT-1184	N/A	<p>There is an issue in the iDRAC firmware where the iDRAC stops responding to some of the <b>racadm</b> commands.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, use the <b>Reboot IPMI</b> option in the <b>Service Menu</b> to reset the iDRAC controller.</p> <ol style="list-style-type: none"> <li>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.</li> <li>2. Launch the <b>Service Menu</b>. Enter: <pre>sh /opt/DXi/scripts/service.sh</pre> <p>The <b>Service Menu</b> displays.</p> </li> <li>3. From the <b>Service Menu</b>, select <b>Hardware Configuration</b>.</li> <li>4. Select <b>Setup IPMI</b>.</li> <li>5. Select <b>Reboot IPMI</b>.</li> </ol>

CR Number	Customer SR Number	Description
80650	N/A	<p>If Bond 1 (typically configured with port em2) is not configured during initial system configuration the system can apply settings to Bond 1 which cause the cvfs file system to not mount. If you check, the following error will be included in the <code>/var/log/messages</code> file:</p> <pre>2023-01-18-09:13:46.244095 leo kernel: CVFS: Cannot get IP address of NIC with MAC 24-6e-96-a2-a2-59 error 'Not found' (2)</pre> <p><b>Workaround:</b></p> <p>To fix this issue, see <a href="#">80650 Workaround: Fix the Bond 1 configuration</a></p>
80640	N/A	<p>Applies to Xcellis Workflow Director (Gen2/R630), Xcellis Foundation, Artico (R630), and Xcellis Workflow Extender (Gen2/R630) systems.</p> <p>On systems reinstalled from ISO or if a Mellanox Connect CX3 card is installed after initial system configuration, DNS configuration settings are removed during the ISO rebuild or after the card is added, causing DNS resolution to fail.</p> <p><b>Workaround:</b></p> <p>If DNS resolution fails on your system, you will need to reconfigure StorNext DNS server settings. See <a href="#">80640 Workaround: Configure missing DNS settings</a>.</p>
80637	N/A	<p>In some rare occasions upgrades to the Appliance Controller might report an upgrade failure even if the upgrade was successful. To validate if the upgrade failure is a false positive look for this message in the controller log:</p> <pre>server.py:142 Appliance controller Failed: local variable 'pid_list' referenced before assignment</pre> <p><b>Workaround:</b></p> <p>To resolve this issue, see <a href="#">80637 Workaround: Appliance Controller upgraded but indicates it did not</a>.</p>

CR Number	Customer SR Number	Description
80565	N/A	<p>In StorNext 7.1, when a StorNext VIP is set in the StorNextUI during HA configuration steps, the HA conversion on Node 2 prevents the StorNext GUI from working.</p> <p><b>Workaround:</b></p> <p>See <a href="#">80565 Workaround: Restart "docker services" on Node 2</a>.</p>
80541		<p>The appliance StorNext upgrade can report errors and BIOS and iDRAC mismatch errors during the upgrade.</p> <p><b>Workaround:</b></p> <p>If this issue is seen, reboot the iDRAC. See <a href="#">80541 Workaround: Reboot iDRAC on page 22</a>. Then attempt to upgrade the system again.</p>
80472	678713	<p>If your system's root file system uses a lot of storage space, attempting to import a large OVA file into DAE can fill up /var, which will trigger RAS events and prevents import of the virtual machine in DAE.</p> <p><b>Workaround:</b></p> <p>If you experience this issue, contact Quantum support. When you contact Quantum, reference <b>Bug 80453</b>.</p>
80393	676195	<p>A RAS ticket may be displayed in the StorNext GUI that indicates an omcliproxy segfault error happened while an IPv6 is bounced on server node, especially when workload on the system is high.</p> <p><b>Workaround:</b></p> <p>A workaround for this issue is being investigated. For now, since the omcliproxy utility error message is benign, it can be ignored, and the core file can be deleted.</p>
80334	673385	<p>When attempting to access HPE 3par configured LUNs on Xcellis Workflow Director host systems using QLE2694L Fibre Channel HBAs, the HBAs can initiate a <b>qla2xxx 10.02.01.00.a14-k1</b> crash, and produce firmware dump files. When this happens, it may cause the server to reboot.</p> <p><b>Workaround:</b></p> <p>See <a href="#">80334 Workaround: HBA crash and firmware dump files</a></p>

CR Number	Customer SR Number	Description
80181	N/A	<p>When there is a disk or disk partition failure in software RAID partitions, a RAS ticket will be generated in the StorNext GUI. However, currently under the <b>Event Details</b> on the ticket the "View Recommended Actions" is missing the workaround for this failure.</p> <p><b>Workaround:</b></p> <p>There is no workaround for the missing information from the "View Recommended Actions" link in the ticket. We are providing the information here. See <a href="#">80181: RAS ticket recommended actions for software RAID failures</a>.</p>
80121	N/A	<p>The StorNext upgrade firmware/bios fails silently even though the upgrade progress shows that the firmware update had completed.</p> <p><b>Workaround:</b></p> <p>See <a href="#">80121 Workaround: Firmware/BIOS fails silently during StorNext upgrade on page 26</a>.</p>
80099	N/A	<p>After the StorNext upgrade completes, StorNext Connect may be left in a partially-operational or inoperable state.</p> <p>If this happens stop and start StorNext Connect manually to attempt to resolve the situation.</p> <p><b>Workaround:</b></p> <p>See <a href="#">80099 Workaround: Potential StorNext Connect issues after you upgrade StorNext on page 29</a>.</p>
79545	N/A	<p>Boot errors may be displayed when connected to F-Series RAID storage.</p> <p><b>Example:</b></p> <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <pre>2021-11-10-09:31:29.643969 rhel-dmraid-activation: ERROR: unsupported sector size 4096 on /dev/sdd.</pre> </div> <p>This error is shown because the <b>dmraid</b> software RAID device running on the appliance only supports disk partitions that have 512 byte sector block size, and F-Series storage is configured with disk partitions configured with a 4096 byte sector block size.</p> <p><b>Workaround:</b></p> <p>There is no workaround for this error. However, these errors do not cause any operational errors, and may be safely ignored.</p>

CR Number	Customer SR Number	Description
79052	612947, 612900	<p>The following applies to and Artico (R630) systems <b>ONLY</b>. It does not apply to other server model names/generations.</p> <p>Upgrades to StorNext 7.0.1 may take many more hours longer than previous upgrades. It can take several hours to shut down the MySQL database. If the upgrade is taking an unusually long time, this database issue is most likely affecting the delay.</p> <hr/> <p> <b>Caution: DO NOT</b> attempt to shut down or reboot the server nodes if the upgrade is taking a long time.</p> <hr/> <p> <b>Note:</b> This was fixed in StorNext 7.0.2, so any upgrades directly to 7.0.2 without first upgrading to 7.0.1 will not experience this issue.</p> <p><b>Workaround:</b></p> <p>There is no workaround for this for StorNext 7.0.1. To speed up the MySQL database conversion, upgrade to StorNext 7.0.2 or later and skip the upgrade to 7.0.1, if possible.</p>
77317	N/A	<p>Adding a DAE virtual machine to a file system created on a QXS or H-Series LUN can create a RAS alert stating:</p> <pre>Excessive fragmentation detected in file "/DAE_Boot_ Disk_location/pk-DAE-vm.qcow2", inode 0x4e800000ef0b19 (decimal inum 15665945)'</pre> <p><b>Workaround:</b></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 <b>extentCountThreshold</b> parameter in the file system configuration file to "0". For information about file system configuration files, see:</p> <p><a href="https://qsupport.quantum.com/kb/flare/Content/stornext/SN6_DocSite/Default.htm#Guide_Tuning/Topics/Example_FSM_Configuration.htm">https://qsupport.quantum.com/kb/flare/Content/stornext/SN6_DocSite/Default.htm#Guide_Tuning/Topics/Example_FSM_Configuration.htm</a></p>
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 <a href="#">74202 – StorNext services can fail to stop completely when NAS is Configured and StorNext services are stopped on page 30</a> for an in-depth explanation of the issue and several alternate workarounds.</p>

CR Number	Customer SR Number	Description
73688	N/A	<p>When the <b>/var</b> directory on the appliance gets full, the following serious errors can occur:</p> <ul style="list-style-type: none"> <li>• Network communication errors on the server node</li> <li>• iDRAC is not accessible</li> <li>• The server node gets stuck during reboot and require a physical power cycle</li> </ul> <p>See also <a href="#">Leave Space on Appliance File Systems</a> on the <i>Appliance InfoHub</i>.</p> <p><b>Workaround:</b></p> <p>Keep additional space available on <b>/var</b>.</p>
70067	N/A	<p><b>Issue:</b></p> <p>CentOS 7 systems display log files with the <b>journalctl</b> command. Any commands executed as the <b>rootsh</b> user are logged to the system log file <b>. A syslog filter sends all rootsh log commands to /var/log/rootshell.</b> This causes excessive information to be captured in the <b>journalctl</b> log file, and makes the log file difficult to read.</p> <p><b>Workaround:</b></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 or later:</p> <p>See the <a href="#">journalctl Filter Script</a></p> <p>For StorNext releases prior to 6.0.5, contact <a href="#">Contacting Quantum Support on page 44</a> and refer to <b>TSB 2998</b>.</p>
55384	N/A	<p><b>Issue:</b></p> <p>If <b>dmnfstthreads</b> is not set on mount, <b>nfsds</b> may be over-committed when there are many NFS processes waiting for offline files.</p> <p><b>Workaround:</b></p> <p>For managed file systems serving NFS, Quantum recommends using the "dmnfstthreads=16" mount option. This setting ensures that NFS remains responsive when Storage Manager is retrieving data from an archive tier.</p>
55318	N/A	<p><b>Issue:</b></p> <p>Strange UID on ACL when file created on non ads client.</p> <p><b>Workaround:</b></p> <p>All systems accessing the StorNext SAN or LAN clients, or the NAS clients, must be part of the same identity domain. Accessing StorNext from different identity domains can result in inconsistent file ownership attributes, as well as potential access problems.</p>

CR Number	Customer SR Number	Description
45702	N/A	<p><b>Issue:</b></p> <p>If you replace an HDD drive with an SSD or vice versa, the StorNext GUI may display a "Missing" status and an equivalent RAS ticket, instead of displaying an "Incompatible" status.</p> <p><b>Workaround:</b></p> <p>Replacement drives must be the identical type of drive removed. HDDs can only use HDD spares, and SSDs can only use SSD spares. Replacement drives must also be the same size or larger than the failed drive. The array controller will generate errors if an incompatible drive is used as a replacement.</p>
38128	1395540	<p><b>Issue:</b></p> <p>Using the GUI while a large Media import is kicked off via the command line can cause the StorNext GUI to timeout or crash.</p> <p><b>Workaround:</b></p> <p>Wait until a bulk load from tape is finished prior to opening the StorNext GUI.</p>

## 80650 Workaround: Fix the Bond 1 configuration

1. Log in to the CLI:
  - a. 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.
 

**Service Port IP addresses:**

    - Node 1 : 10.17.21.1 (dual-server node systems)
    - Node 2 : 10.17.21.2 (dual-server node, or supported "HA-ready" single-server node StorNext appliances)
  - b. Initiate an ssh session to the system using Terminal or PuTTY:



**To ssh into the system**

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- b. Enter **sudo rootsh** to gain root user access.
- c. Enter the password for the **stornext** user account again.

2. Manually change the first line of the **/etc/sysconfig/network-scripts/ifcfg-bond1** file from **DEVICE=bond0:2** to **DEVICE=bond1:2**.

The **vi** editor is included in CentOS installed on your system. The following Google search link provides links for editing with **vi**:

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

---

**i Note:** When finished making changes with the **vi** editor, save and exit the file using **:wq!**.

## 80640 Workaround: Configure missing DNS settings

Configure your DNS settings:

1. Log in to the CLI:
  - a. 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.

**Service Port IP addresses:**

- Node 1 : 10.17.21.1 (dual-server node systems)
  - Node 2 : 10.17.21.2 (dual-server node, or supported "HA-ready" single-server node StorNext appliances)
- b. Initiate an ssh session to the system using Terminal or PuTTY:

### To ssh into the system

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- b. Enter **sudo rootsh** to gain root user access.
- c. Enter the password for the **stornext** user account again.

2. Open the **Service Menu**:

- a. Launch the **Service Menu**. Enter:

```
sh /opt/DXi/scripts/service.sh -admin
```

The **Service Menu** displays.

3. Navigate to **More Tools > Engineering Tools > Base OS Network Setup > Configure host settings**.
4. Click **<ENTER>** for **Hostname**.
5. Click **<ENTER>** for **Default Gateway IP**.
6. Enter the **Domain Name**, default will show the node name, so replace **node-1** or **node-2** (the system defaults) with a valid domain name to use.
7. Enter one or more (comma-separated) **DNS IP** addresses for the system.

Once you press **<ENTER>**, you should see the settings apply similar to:

```
INFO: *** script started with arguments: addhost --hostname myhostname1 --
domain mydomain.local --dns 10.20.219.63 --defaultgateway
10.20.216.1,10.20.216.2 ***
INFO: *** script completed successfully ***
```

The **Network Menu** displays again.

8. Select the **Activate Network Changes** option to apply the DNS changes.
9. Close the SSH session for this server node.

For single-server node systems, the DNS settings change is now complete.

For dual-server node systems, continue to the next step.

10. Log in to the CLI for the other server node:

- a. 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.

**Service Port IP addresses:**

- Node 1 : 10.17.21.1 (dual-server node systems)
- Node 2 : 10.17.21.2 (dual-server node, or supported "HA-ready" single-server node StorNext appliances)

- b. Initiate an ssh session to the system using Terminal or PuTTY:

**To ssh into the system**

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- b. Enter **sudo rootsh** to gain root user access.
- c. Enter the password for the **stornext** user account again.

11. Open the **Service Menu**:

- a. Launch the **Service Menu**. Enter:

```
sh /opt/DXi/scripts/service.sh -admin
```

The **Service Menu** displays.

12. Navigate to **More Tools > Engineering Tools > Base OS Network Setup > Configure host settings**.
13. Click **<ENTER>** for **Hostname**.
14. Click **<ENTER>** for **Default Gateway IP**.

15. Enter the **Domain Name**, default will show the node name, so replace **node-1** or **node-2** (the system defaults) with a valid domain name to use.
  16. Enter the identical **DNS IP** address/addresses entered for the other server node.
  17. Once you press <ENTER>, you should see the settings apply similar to:
  18.
 

```
INFO: *** script started with arguments: addhost --hostname myhostname1 --
domain mydomain.local --dns 10.20.219.63 --defaultgateway
10.20.216.1,10.20.216.2 ***
INFO: *** script completed successfully ***
```
  19. The **Network Menu** displays again.
  20. Select the **Activate Network Changes** option to apply the DNS changes.
  21. Close the SSH session for this server node.
- For dual-server node systems, the DNS settings change is now complete.

## 80637 Workaround: Appliance Controller upgraded but indicates it did not

1. To determine if the upgrade is incomplete, locate the following error message in the /var/log/snnas\_ controller log file:

```
server.py:142 Appliance controller Failed: local variable 'pid_list'
referenced before assignment
```

2. Log in to the server node:
  - a. 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.

**Service Port IP addresses:**

  - Node 1 : 10.17.21.1 (dual-server node systems)
  - Node 2 : 10.17.21.2 (dual-server node, or supported "HA-ready" single-server node StorNext appliances)
  - b. Initiate an ssh session to the system using Terminal or PuTTY:

**To ssh into the system**

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- b. Enter **sudo rootsh** to gain root user access.
- c. Enter the password for the **stornext** user account again.

3. Verify the Appliance Controller was upgraded and is running the latest/expected release on the current server node:

```
qtmcontroller -c 'system show version'
```

4. If the release is installed as expected. Enter the following on the server node to remove the touch file placed on the system to indicate an upgrade issue on the server node:

```
rm /var/DXi/upgrade_notify
```

**80565 Workaround: Restart "docker services" on Node 2**

1. Log into the CLI of the server node, if not already logged in.

**How to log into the server CLI:**

- a. Open an SSH connection to the appropriate server and use the IP address assigned to the node on the Management or LAN Client network.
- b. Initiate an ssh session to the system using Terminal or PuTTY:

**To ssh into the system**

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- b. Enter **sudo rootsh** to gain root user access.
- c. Enter the password for the **stornext** user account again.

2. Enter the docker restart command:

```
systemctl restart docker
```

3. Exit the CLI.
4. Launch the StorNext User Interface again.

## 80541 Workaround: Reboot iDRAC

1. Log in to the server node:
  - a. 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.

**Service Port IP addresses:**

- Node 1 : 10.17.21.1 (dual-server node systems)
- Node 2 : 10.17.21.2 (dual-server node, or supported "HA-ready" single-server node StorNext appliances)

- b. Initiate an ssh session to the system using Terminal or PuTTY:

### To ssh into the system

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- b. Enter **sudo rootsh** to gain root user access.
- c. Enter the password for the **stornext** user account again.

2. Open the **Service Menu**:

- a. Launch the **Service Menu**. Enter:

```
sh /opt/DXi/scripts/service.sh -admin
```

The **Service Menu** displays.

3. Navigate to **Hardware Configuration > Setup IPMI > Reboot IPMI**.
4. At the prompt to reboot the IPMI controller, press **y** (yes).
5. The iDRAC/IPMI will reboot.

For single-server node systems, the procedure is complete.

6. For dual-server node systems, reboot the iDRAC on the other server node using the steps you just completed on this server node.

## 80334 Workaround: HBA crash and firmware dump files

When you experience this issue, you need to set the `ql2xenabledif_tgt` value to 0 and the `qlport_down_retry` value to 14:

1. Create the `/etc/modprobe.d/hpe.conf` file. Or, if it exists, create another unique but identifiable file name for this setting)
2. Log in to the server node:
  - a. Open an SSH connection to the appropriate server and use the IP address assigned to the node on the Management or LAN Client network.
  - b. Initiate an ssh session to the system using Terminal or PuTTY:

### To ssh into the system

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- b. Enter **sudo rootsh** to gain root user access.
- c. Enter the password for the **stornext** user account again.

3. Display the contents of the `/etc/modprobe.d/hpe.conf` file:

```
cat /etc/modprobe.d/hpe.conf
```

It should contain the following entries:

```
options qla2xxx qlport_down_retry=14
options qla2xxx ql2xenabledif_tgt=0
```

If not, add `retry` and `tgt` options to the file.

4. If you made changes to the file, save the file.
5. Recreate the **initramfs** with this new information.

---

**i Note:** Do not reboot the system before completing the following 3 steps.

6. Move the `img`:

```
mv /boot/initramfs-$(uname -r).img /boot/initramfs-$(uname -r).img.prehpe
```

7. Create the initial image used by the kernel for preloading the block device modules:

---

**i Note:** There are two dashes before **kver** in the following command.

```
dracut --kver $(uname -r)
```



8. Enter the following to set the reboot options that ensure the new **initramfs** file is generated:

```
ls /boot/initramfs-$(uname -r).img
```

9. Reboot the node. Enter:

```
reboot
```

10. After the reboot, verify the **qlport\_down\_retry** value is now set to **14** as described here:  
**cat /sys/module/qla2xxx/parameters/qlport\_down\_retry**

## 80181: RAS ticket recommended actions for software RAID failures

1. Log in to the CLI:
  - a. 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.

### Service Port IP addresses:

- Node 1 : 10.17.21.1 (dual-server node systems)
- Node 2 : 10.17.21.2 (dual-server node, or supported "HA-ready" single-server node StorNext appliances)

- b. Initiate an ssh session to the system using Terminal or PuTTY:

### To ssh into the system

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- b. Enter **sudo rootsh** to gain root user access.
      - c. Enter the password for the **stornext** user account again.

2. Open the **Service Menu**:

- a. Launch the **Service Menu**. Enter:

```
sh /opt/DXi/scripts/service.sh -admin
```

The **Service Menu** displays.

3. Navigate to **More Tools > SW RAID Tools**.
4. To repair a software RAID configuration, select the **SW RAID Repair** option.

## 80121 Workaround: Firmware/BIOS fails silently during StorNext upgrade

To check if there was a failure of applying Dell f/w updates, do the following. First, log in to the CLI of the server node that is currently not acting as primary, which prevents a failover if a manual CLI re-install of the StorNext software/reboot is required after the check. We will repeat the steps for the other server node during this process, and upgrade that firmware manually from the CLI, if needed.

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.

### Service Port IP addresses:

- Node 1 : 10.17.21.1 (dual-server node systems)
- Node 2 : 10.17.21.2 (dual-server node, or supported "HA-ready" single-server node StorNext appliances)

2. Initiate an ssh session to the system using Terminal or PuTTY:

### To ssh into the system

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- b. Enter **sudo rootsh** to gain root user access.
- c. Enter the password for the **stornext** user account again.

3. From the CLI, change to the firmware directory. Enter:

```
cd /opt/DXi/firmware
```

4. From the CLI, run the following command. Enter:

```
./fw_regenall.sh
```

**Example output:**

```
fw_bundle: Performing firmware cleanup.
fw_bundle: Firmware cleanup completed successfully.
fw_bundle: Checking Firmware Bundle Version.
fw_bundle: Need to generate CoR files for the first time.
fw_bundle: Performing firmware CoR generation...
fw_bundle: Firmware CoR generation completed successfully.
fw_bundle: Performing firmware description report.
fw_bundle: Generating firmware description manifests.
fw_bundle: Firmware manifest generation is disabled.
fw_bundle: Generating firmware bundle manifest.
fw_bundle: Firmware manifest generation is disabled.
fw_bundle: Generation of the firmware bundle manifest completed
successfully.

fw_bundle: Performing firmware bundle report.
fw_bundle: Performing TXT firmware report.
fw_bundle: Performing GUI firmware report.
fw_bundle: GUI Firmware reporting completed successfully.
fw_bundle: TXT Firmware reporting completed successfully.
fw_bundle: Firmware reporting completed successfully.
```

At this time, you could check the firmware versions installed by checking the FirmwareReport.txt file. To view the file contents, enter:

```
cat /opt/DXi/hwdetect/FirmwareReport.txt
```

5. From the CLI, run the following command. Enter:

```
./fw_check.sh
```

There are three possible outputs that you may see when you run the `fw_check.sh` command:

- **Example Output #1** - If you see the following message, a firmware **IS NOT** required:

```
fw_bundle: Performing firmware check.
.
.
fw_bundle: Firmware check completed successfully. All firmware is
up-to-date.
```

The firmware check output above shows that the firmware installed on the system is up-to-date. There are no further actions to take. You can exit the **rootsh** session, and close the SSH terminal client.

- **Example Output #2** - If you see the following message, a firmware update **IS NOT** required:

```
fw_bundle: Performing firmware check.
fw_bundle: WARNING: Dell r630 iDRAC8/Life Cycle Controller - iDRAC
Information,
iDRAC Firmware Information, Firmware Version: Mismatch; Expected:
'2.70.70.70',
Found: '2.75.75.75'
fw_bundle: WARNING: Dell r630 BIOS - BIOS Information, BIOS Firmware
Information,
Version: Mismatch; Expected: '2.9.1', Found: '2.11.0'e.
.
.
fw_bundle: WARNING: One or more firmware packages have a mismatched
configuration.
```

The firmware check output above shows that the firmware installed on the system is newer than the firmware included in the StorNext software update. There are no further actions to take. You can exit the **rootsh** session, and close the SSH terminal client.

- **Example Output #3** - If you see the following message, a firmware update **IS** required. The firmware check shows that there is one or more firmware on the system that is older than what is included in the included in the StorNext software update:

```
fw_bundle: Performing firmware check.
fw_bundle: WARNING: Dell r630 iDRAC8/Life Cycle Controller - iDRAC
Information, iDRAC
Firmware Information, Firmware Version: Mismatch; Expected:
'2.70.70.70', Found: '2.75.75.75'
```

```
fw_bundle: WARNING: Dell r630 BIOS - BIOS Information, BIOS Firmware
Information, Version:
Mismatch; Expected: '2.11.0', Found: '2.9.1'.
```

Continue to [Step 6](#).

6. If you need to re-apply the f/w update, run the following command:

```
./fw_install.sh
```

If a server reboot is required, the system will reboot.

7. Exit the **rootsh** session.
8. Exit the SSH terminal session.
9. Repeat [Step 1](#) through [Step 8](#) for the server node acting as primary.

If this does not fix the StorNext Connect issue, contact Quantum Support for assistance.

## 80099 Workaround: Potential StorNext Connect issues after you upgrade StorNext

Do the following to stop and restart StorNext Connect. Log in to the CLI of node 2 **ONLY**, since StorNext Connect does not run on node 1:

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.

### Service Port IP addresses:

- Node 1 : 10.17.21.1 (dual-server node systems)
  - Node 2 : 10.17.21.2 (dual-server node, or supported "HA-ready" single-server node StorNext appliances)
2. Initiate an ssh session to the system using Terminal or PuTTY:

### To ssh into the system

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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

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

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

3. From the CLI of node 2, stop StorNext Connect. Enter:

```
/opt/quantum/connect/bin/connect stop
```

4. From the CLI of node 2, start StorNext Connect. Enter:

```
/opt/quantum/connect/bin/connect start
```

5. Exit the **rootsh** session.

6. Exit the SSH terminal session.

If this does not fix the StorNext Connect issue, contact Quantum Support for assistance.

## 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- 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 page 34](#).

#### 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- 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.

## journalctl Filter Script

### Step 1: Upgrade to StorNext 6.0.5

**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.

### Service Port IP addresses:

- Node 1 : 10.17.21.1 (dual-server node systems)
- Node 2 : 10.17.21.2 (dual-server node, or supported "HA-ready" single-server node StorNext appliances)

2. Initiate an ssh session to the system using Terminal or PuTTY:

### To ssh into the system

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

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

**Note:** "password" is the default password for the stornext user account. If the password has been changed, use the current password. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

- b. Enter **sudo rootsh** to gain root user access.
- c. Enter the password for the **stornext** user account again.

3. Continue to [Step 3: Run the Journal Filter Script and View the Script Help below](#).

## 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.

## Restart the Appliance Controller

On the server node hosting NAS, do the following (use node 2 on dual-server node systems):

1. Open an SSH connection to the appropriate server node and use the IP address assigned to the node on the Management or LAN Client network.
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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

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

```
chown -R ldap:ldap /var/lib/ldap
```

5. Enter the following (Commands vary by system type. If you don't know which system you are using, see the [How to Identify Your System](#) page.):

For Artico (R630) systems only.

```
systemctl restart slapd
```

**OR**

For Artico (R520) systems only

```
service slapd restart
```

6. Restart the appliance controller:

For Artico(R630) systems only.

```
systemctl start snnas_controller
```

Then...

```
systemctl stop snnas_controller
```

**OR**

For Artico (R520) systems only

```
initctl snnas_controller stop
```

Then...

```
initctl snnas_controller start
```

## How to Disable NFS v4

To disable NFSv4 on Artico systems, especially those systems running NAS 1.2.0 or earlier, perform the following on node 2:

### Edit the nfs File

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.

#### Service Port IP addresses:

- Node 1 : 10.17.21.1 (dual-server node systems)
- Node 2 : 10.17.21.2 (dual-server node, or supported "HA-ready" single-server node StorNext appliances)

2. Initiate an ssh session to the system using Terminal or PuTTY:



### To ssh into the system

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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

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

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

3. Edit **/etc/sysconfig/nfs** file. (See [Example: Use vi to Edit the nfs File on the next page](#))
4. Change the value for **RPCNFSDARGS** to "-N 4". When finished, the line must be:  
**RPCNFSDARGS="-N 4"**
5. Change the value for **RPCMOUNTDOPTS** to "-N 4". When finished the line must be:  
**RPCMOUNTDOPTS="-N 4"**
6. Save the file.

### Restart NFS (for CentOS7-based Systems)

For Artico (R630) Archive Gateway systems.

1. Enter the following to restart the NFS configuration.:

```
systemctl restart nfs-config
```

2. Enter the following to restart the NFS server.:

```
systemctl restart nfs-server
```

### Restart NFS (for CentOS6-based Systems)

For Artico(R520) NAS Archive Appliance systems.

1. Enter the following to restart the NFS configuration:

```
service nfs-config restart
```

2. Enter the following to restart the NFS server:

```
service nfs-server restart
```

### After You Restart NFS on Node 2

Repeat both the [Edit the nfs File on page 40](#) and [Known Issues on page 8](#) sections on node 1 ( for dual-node systems).

If you need to create new NFS shares to export for user access, you may do this now. See the [About the NAS App](#) section of the **StorNext Connect Documentation Center** for information about NAS configuration using StorNext Connect.

### Example: Use vi to Edit the nfs File

1. Enter the following:

```
vi /etc/sysconfig/nfs
```

2. Move the cursor to the closing quotation mark in RPCNFSDARGS.
3. Enter the following:

```
i -N 4
```

4. Write the file and quit vi as follows:

```
:wq
```

5. For this example, you would have to [Restart NFS \(for CentOS7-based Systems\) on the previous page](#) OR [Restart NFS \(for CentOS6-based Systems\) on the previous page](#), and repeat the [Edit the nfs File on page 40](#), and [Restart NFS \(for CentOS7-based Systems\) on the previous page](#) OR [Restart NFS \(for CentOS6-based Systems\) on the previous page](#) again on node 1 (for dual-node systems).
6. Close the SSH session for the server(s) ( for dual-node systems).

## Re-enable DDM on the Secondary Node After HA Failover

Once the MDC node acting as secondary finishes rebooting and becomes functional again, use `fsddmconfig` (or the GUI) from the server node currently acting as primary to re-enable DDM for the standby server node, as follows:

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.

### Service Port IP addresses:

- Node 1 : 10.17.21.1 (dual-server node systems)
- Node 2 : 10.17.21.2 (dual-server node, or supported "HA-ready" single-server node StorNext appliances)

2. Initiate an ssh session to the system using Terminal or PuTTY:

### To ssh into the system

- 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. Beginning with StorNext 7.0.2, you must change the default password for the stornext user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.


- b. Enter **sudo rootsh** to gain root user access.
- c. Enter the password for the **stornext** user account again.

3. On the command line of the system enter the following:

```
fsddmconfig -u -s e
standby_system_hostname
```

If a system running DDMs periodically displays an **Admin Alert** when it fails over, cycles a client node, or upgrades while the server node operating as secondary is rebooting but not completely down, change the timeout value to allow more time to communicate with the node. Update the TSM configuration parameter **DDM\_CLIENT\_RETRY\_TIMEOUT** from the default 30 seconds to a larger value depending on how long the node is taking to reboot. In the event the server node acting as secondary will be down for an extended period of time, the node should be taken offline. Additionally, if this parameter is set too high, and the standby server does not come back, the command that is currently running (e.g., store or retrieve) will not return until after the specified timeout.

For more information about the **DDM\_CLIENT\_RETRY\_TIMEOUT** parameter that can be placed into the **fs\_sysparm\_override** configuration file, refer to the `/usr/adic/TSM/config/fs_sysparm.README` file.

 **Caution:** Incorrect modification of the **fs\_sysparm\_override** configuration file can cause a serious, adverse effect on StorNext functionality. Before modifying this file, we recommend you contact [Contacting Quantum Support below](#).

---

## Quantum Appliance Licenses

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

---

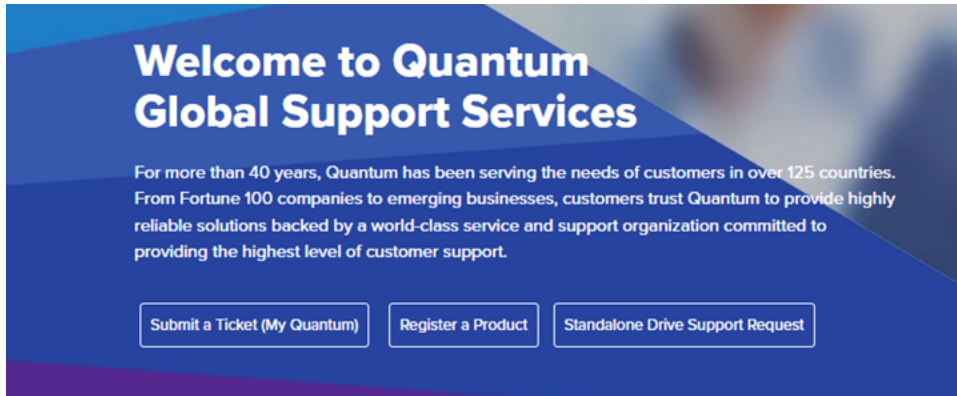
## Contacting Quantum Support

Below is information related to contacting Quantum Support as well as steps to improve your Quantum customer journey.

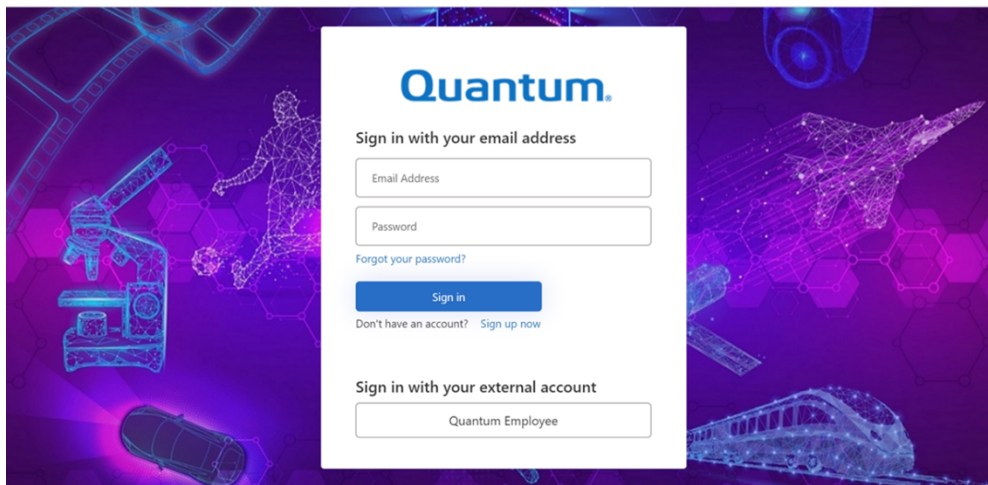
- [Submit a Ticket \(Service Request\) below](#)
- [Use MyQuantum Service Delivery Platform on the next page](#)
- [Use Cloud Based Analytics \(CBA\) on page 46](#)
- [Escalate a Case on page 46](#)
- [Contact Quantum Sales on page 46](#)

### Submit a Ticket (Service Request)

If you need to submit a ticket or speak to Quantum technical support, go to the Support page at <https://www.quantum.com/en/service-support/>




To start the process with Quantum Technical Support, click **Submit a Ticket**. From here, sign in to the MyQuantum Service Delivery Platform or create an account. For more information, refer to the [Use MyQuantum Service Delivery Platform below](#) section below.



## Use MyQuantum Service Delivery Platform

MyQuantum is a single portal for everything Quantum. You can view assets, open support cases, receive real-time updates, and search the Knowledge Base and documentation, all through a secure, online portal.

1. Create an account and log in to the [MyQuantum Service Delivery Platform](#).
2. Register a product on [MyQuantum](#).



3. Request site access to the Cloud-Based Analytics (CBA) monitoring portal and follow the instructions to set up product(s) to connect to CBA. You can use CBA to monitor Quantum products remotely, from a single dashboard, and Quantum Support can use it to help troubleshoot products more efficiently.

Refer to product documentation for product-specific information related to CBA.

## Use Cloud Based Analytics (CBA)

Quantum products are equipped with a Cloud Based Analytics (CBA) agent that can provide log files and snapshots to Quantum CBA servers that are running in the cloud.

CBA enables Quantum systems to collect data regarding system and environment performance. The collected data is bundled and uploaded to the remote CBA server for analysis. You can access Quantum system performance and health results on the CBA dashboard (at <https://insight.quantum.com>) or through the MyQuantum Service Delivery Platform.

The CBA dashboard displays the analytic results of the uploaded CBA data using flexible charting tools, along with an overall health score of each Quantum system configured for the CBA account.

## Escalate a Case

To escalate a case, follow the process documented here: <https://www.quantum.com/en/service-support/resources/escalation/>

## Contact Quantum Sales

<https://www.quantum.com/en/company/contact-us/>