

## StorNext 7.2.2 Xcellis Workflow Director Gen 3 Release Notes

---

<b>Original Product/Software Release Date</b>	August, 2024
---	--------------

---

### Contents

About the Xcellis Workflow Director Gen 3 System .....	2
Training and Documentation Resources .....	2
Learn more about StorNext 7.2.2 .....	2
What StorNext Releases are compatible with my system? .....	3
Known Issues .....	3
Quantum Appliance Licenses .....	24
Contacting Quantum Support .....	24

---

© 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. ActiveScale, DXi, DXi Accent, FlexSync, FlexTier, iLayer, Lattus, Myriad, Quantum, the Quantum logo, QXS, Scalar, StorNext, SuperLoader, Unified Surveillance Platform, USP, 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 Xcellis Workflow Director Gen 3 System

Quantum Xcellis® is the control center for fast, reliable StorNext® environments worldwide, driving and managing every aspect of the file system and shared storage environment, connected users, and attached storage systems. Xcellis can be deployed as a highavailability pair of two servers, then can add more servers at any time to offer more client connectivity performance or additional network share performance.

The newest generation of Xcellis offers higher peak and sustained performance, and up to 1.6x performance improvement for file systems operations, ensuring that even the largest shared storage environments are highly responsive for all users. Xcellis is engineered to get the highest throughput from the widest range of shared storage appliances like Quantum QXS™-5 Series HDD and SSD storage systems, and offers the fastest performance ever for StorNext on Quantum F-Series NVMe flash storage systems for the most demanding streaming file needs.

Xcellis is also designed to make the most of the fastest available connection technologies, offering a choice of 100 Gigabit Ethernet with RDMA support for NVMe-oF or 32Gb Fibre Channel, and is ready for 200 Gigabit Ethernet or 64Gb Fibre Channel as your connectivity environment evolves.

---

## Training and Documentation Resources

[Xcellis Workflow Director Gen 3 training and documentation](#)

---

## Learn more about StorNext 7.2.2

### 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.

What StorNext Releases are compatible with my system?

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

---

## Known Issues

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

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

### StorNext File System Known Issues

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

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

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	67363	n/a	<p>StorNext 5.4.0.x incorrectly allowed the <b>Unix ID Mapping</b> type to be set to <b>none</b> when the <b>Security Model</b> is set to <b>acl</b>. As a result, file systems fail to start when the <b>Unix ID Mapping</b> type is set to <b>none</b> when the <b>Security Model</b> is set to <b>acl</b>.</p> <p>Beginning with StorNext 6, the FSM does not start when this invalid combination of settings is used.</p> <p><b>Workaround:</b></p> <p>To prevent this issue, set the <b>Unix ID Mapping</b> to either <b>winbind</b> or <b>algorithmic</b> for any file system where the <b>Security Model</b> is set to <b>acl</b>. You can make the adjustment before or after upgrading.</p>
All	75633	n/a	<p>A StorNext NAS client cannot rename a file if the file has the read-only attribute set. This problem only affects StorNext NAS clients.</p> <p><b>Workaround:</b></p> <p>A StorNext NAS client must remove the read-only attribute before it can rename the file.</p>
All	78082	575600	<p>If you set the parameter <b>audit=true</b> in the file system configuration file, the result causes all I/O activity, including reads, to be recorded in the <b>mdarchive</b>. This can greatly increase the number of updates applied to the <b>mdarchive</b> which, in turn, increases the amount of <b>mdarchive</b> compaction activity by the FSM.</p> <p>There is one instance where this increased compaction activity caused the <b>mdarchive</b> to grow beyond the capacity of the HA shared file system.</p> <p><b>Workaround:</b></p> <p>To work around this issue, Quantum recommends you increase the parameter <b>metadataArchiveCache</b> from the default 2 GB to at least 4 GB, but preferably larger when setting audit to <b>true</b> in the configuration file.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
macOS	66948	322824, 336945	<p>If you access StorNext file systems from Apple Xsan clients, then you might encounter I/O error messages in the system log that do not contain details about real I/O errors detected on the Xsan client.</p> <p><b>Workaround:</b></p> <p>If you encounter the errors on an Xsan client, contact Apple.</p>
macOS	75819	n/a	<p>An Xsan client cannot mount a StorNext File System volume when a cluster number is included in the local <b>fsnameservers</b> file. An unexpected EOF reading reply error is displayed.</p> <pre># xsanctl mount snfs3 xsanctl: unexpected EOF reading reply</pre> <p>If you encounter the EOF error on your Xsan client, do the following workaround to prevent the issue.</p> <p><b>Workaround:</b></p> <p>If your MDC <b>fsnameservers</b> file includes a cluster number, remove the cluster number (<b>@_cluster_xx</b>) from the <b>mysan.configprofile</b> file <b>before</b> you copy the file to the Xsan client.</p> <p>For example, change:</p> <pre>10.65.181.158@_cluster0</pre> <p>to</p> <pre>10.65.181.158</pre> <p>See <a href="#">Mount the StorNext File System on Xsan 5.0 (or later)</a> for additional information.</p>

## StorNext Storage Manager Known Issues

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

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	43320	1581004	<p>File retrieves from media to disk can be suboptimal for fast tape drives like the Oracle STK T10K drives. This scenario can occur when the retrieve event is initiated on a host that is different from the host running the mover process, which requires the use of synchronous direct I/O.</p> <p><b>Workaround:</b></p> <p>To work around this issue and achieve optimal performance for both file stores and retrieves with the T10K drives, increase the default I/O size used by the mover process and make the mover process use asynchronous buffered I/O when the use of synchronous direct I/O is not required, using the following steps:</p> <ul style="list-style-type: none"> <li><b>i Note:</b> This workaround might also help improve the performance of the faster LTO drives by updating the <b>FS_LTO_BLOCK_FACTOR</b> system parameter.</li> <li><b>i Note:</b> Changes to <b>FS_XXX_BLOCK_FACTOR</b> only affects tapes formatted after the change.</li> </ul> <ol style="list-style-type: none"> <li>1. Change the <b>FS_T10K_BLOCK_FACTOR</b> system parameter from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <pre data-bbox="834 1136 1455 1199" style="background-color: #f0f0f0; padding: 5px;">FS_T10K_BLOCK_FACTOR=32;</pre> <ul style="list-style-type: none"> <li><b>i Note:</b> The T10K default I/O block size is 512 KB or 8 * 64 KB. With the block factor changed to 32, the new T10K I/O block size will be 2 MB or 32 * 64 KB. Presently, the <b>FS_T10K_BLOCK_FACTOR</b> system parameter must not be set to a value that exceeds 32.</li> </ul> </li> <li>2. Restart Storage Manager to ensure the change in <b>Step 1</b> goes into effect: <pre data-bbox="818 1583 1455 1694" style="background-color: #f0f0f0; padding: 5px;"># tsmstop # tsmstart</pre> </li> </ol>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>3. Verify the <b>FS_T10K_BLOCK_FACTOR</b> system parameter contains the new value:</p> <pre data-bbox="818 436 1458 541"># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> <p>4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients.</p> <p>5. Modify <code>/etc/fstab</code> on the MDCs and the DDM clients to use the <code>auto_dma_write_length</code> and <code>auto_dma_read_length</code> mount options as follows:</p> <pre data-bbox="818 814 1458 953">snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_ dma_read_length=16m 0 0</pre> <p>6. Unmount and re-mount your file systems.</p> <p>7. Use new T10K media to store a copy of the file from the disk.</p> <p><b>i Note: Step 7</b> is very important; when the new copy is made to the new tapes, the new tapes are labeled with a 2 MB block size, which is used for subsequent writes or reads to and from the media. Tapes on which <b>fsformat</b> was run before the change will use the block factor in use at that time. This change will not impact those tapes.</p>



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

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

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

## StorNext GUI Known Issues

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

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

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69958	373823	<p>There is a known issue where stripe group expansion using the StorNext GUI can fail and requires manual intervention to restore file system operation.</p> <p>Stripe group expansion allows an additional disk (LUN) to be added to an existing stripe group, growing the file system user data space.</p> <p>The other way to add space to a file system is to create a new stripe group and add this to the file system.</p> <p>When the GUI executes stripe group expansion, it stops the file system, modifies the configuration file and runs <b>cvupdatefs</b>. If <b>cvupdatefs</b> takes more than 5 minutes to complete, the GUI kills the <b>cvupdatefs</b> utility and reports an error. At this point the file system does not start because the configuration file does not match the current state of the metadata.</p> <p><b>Workaround:</b></p> <p>To address this issue, you can do one of two things:</p> <ul style="list-style-type: none"> <li>• Optimally, you would not attempt the stripe group expansion at all. Instead add a new stripe group to the file system.</li> <li>• If stripe group expansion is deemed necessary, use the <b>cvupdatefs</b> CLI directly instead of through the GUI. This method does not have any time limitations.</li> </ul> <p>In the case that the StorNext GUI attempt was made and hit the 5 minute timeout, file system operations can be resumed by performing the following:</p> <ol style="list-style-type: none"> <li>1. Restore the previous version of the configuration file. This can be found in the following directory: <div data-bbox="820 1470 1458 1539" data-label="Code-Block"> <pre>/usr/cvfs/data/&lt;fs&gt;/config_history</pre> </div> </li> <li>2. Run the <b>cvfsck</b> utility to verify and potentially correct the metadata.</li> <li>3. Start the file system.</li> </ol>

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

All	79549	n/a	You cannot use the StorNext GUI to configure or modify a policy class with the <b>GOOGLES3</b> media type. The process fails with an error similar to the following:
-----	-------	-----	--

```
Failed to create storage policy or some of its attributes, e.g. associated directories and schedules.
```

```
Failed to update steering parameters for Copy [1] for Media Type [GOOGLES3]. Verify that the Media Type and Media Format is correct for each copy defined on the Steering tab.
```

**Workaround:**

To workaround the issue, Quantum recommends you use the **GOOGLE** media type rather than the **GOOGLES3** media type, not only because the **GOOGLES3** media type provides very limited functionality compared to the **GOOGLE** media type, but also because you cannot use the StorNext GUI to configure or modify a policy class with the **GOOGLES3** media type. However, if you have to use the **GOOGLES3** media type, then use the CLI **fsaddclass** command to configure a policy class and the **fsmodclass** command to modify an existing policy class.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	80068	n/a	<p>The status and capacity monitors for file systems, libraries, storage disks, and tape drives does not display data on the StorNext GUI <a href="#">home page</a>. This issue is caused as a result of missing font RPM packages on your system.</p> <p><b>Workaround:</b></p> <p>To workaround the issue, Quantum recommends you do the following:</p> <ol style="list-style-type: none"> <li>1. Install the <b>java-1.8.0-openjdk</b> package on your system:</li> </ol> <pre>yum install java-1.8.0-openjdk</pre> <ol style="list-style-type: none"> <li>2. Restart the Apache Tomcat service:</li> </ol> <pre>systemctl restart stornext_web</pre>
All	80600	n/a	<p>If you have an open StorNext user interface browser session, then the browser page might flash due to an ICEfaces error. This occurs when the browser page displays partial data rows, empty data, or all the data.</p> <p><b>Workaround:</b></p> <p>To workaround the issue, Quantum recommends you do the following:</p> <ol style="list-style-type: none"> <li>1. Log out of your StorNext user interface browser session(s).</li> <li>2. Close all browser tabs and windows.</li> <li>3. Open a new browser window and access the StorNext user interface (see <a href="#">Access the StorNext User Interface (new UI introduced in StorNext 7.0.1)</a> or <a href="#">Access the StorNext Software GUI</a>).</li> </ol>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	47954	n/a	<p>The Safari browser becomes unresponsive when you attempt to configure an Email server using the StorNext GUI.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, perform the following procedure:</p> <ol style="list-style-type: none"> <li>1. Shut down the Safari browser window(s).</li> <li>2. Restart the Safari browser, and then retry the operation.</li> <li>3. Uncheck the <b>Verify SMTP Server Connectivity</b> box, and then retry the operation.</li> <li>4. Set <b>Authentication</b> to <b>NONE</b>, and then retry the operation.</li> <li>5. Disable the Safari <b>User names and passwords</b> AutoFill under <b>Safari &gt; Preferences &gt; AutoFill</b>, and then retry operation.</li> </ol>
All	SNXT-1329	n/a	<p>When you use the StorNext GUI to configure the <b>Timeserver</b> or <b>Pool</b> (located within the <b>Network Time Protocol (NTP)</b> section on the <b>Configuration &gt; System &gt; Date &amp; Time</b> page), and enter multiple comma separated IP addresses in the <b>Timeserver Pool</b> field, and then click <b>Test NTP</b>, the NTP test hangs and prevents you from accessing the GUI.</p> <p><b>Example of Timeserver Pool Entry</b></p> <div style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;"> <p><b>10.65.162.1,10.65.162.2</b></p> </div> <p><b>Workaround:</b></p> <p>Do the following to workaround this issue:</p> <ol style="list-style-type: none"> <li>1. Restart the StorNext GUI.</li> <li>2. Do <b>not</b> enter multiple IP addresses in the <b>Timeserver Pool</b> field. You can only enter one IP address.</li> </ol>

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

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



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

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

All	SNXT-389	n/a	The rpm database might corrupt if an rpm operation or a query is unexpectedly interrupted by a signal, reboot, or a system crash on an Xcellis XWD environment in particular, but could also impact a StorNext MDC in a customer-supplied environment and anywhere that the Appliance Controller software or the Unified User Interface (UI) software runs.
-----	----------	-----	---

Querying the database might result in errors and report that an installed package is not installed. For example:

```

$ rpm -q snfs-server
error: rpmdb: BDB0113 Thread/process
7492/140149665155136 failed: BDB1507 Thread died
in Berkeley DB library
error: db5 error(-30973) from dbenv->failchk:
BDB0087 DB_RUNRECOVERY: Fatal error, run database
recovery
error: cannot open Packages index using db5 - (-
30973)
error: cannot open Packages database in
/var/lib/rpm
error: rpmdb: BDB0113 Thread/process
7492/140149665155136 failed: BDB1507 Thread died
in Berkeley DB library
error: db5 error(-30973) from dbenv->failchk:
BDB0087 DB_RUNRECOVERY: Fatal error, run database
recovery
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 a

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

reboot of the system while an rpm activity occurred. StorNext, Appliance Controller, and 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:**

Run the following command to repair the rpm database:

```
$ rpmdb --rebuilddb
```

After you rebuild the database, run a query to verify the database is repaired:

```
$ rpm -q snfs-server
snfs-server-7.1.1-91E.RedHat7.x86_64
```

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	SNXT-547	n/a	<p>You might see the following error messages when the Rocky 8 kernel first boots:</p> <pre>DMAR: [Firmware Bug]: No firmware reserved region can cover this RMRR [0x000000006f760000-0x000000006f762fff], contact BIOS vendor for fixes DMAR: [Firmware Bug]: Your BIOS is broken; bad RMRR [0x000000006f760000-0x000000006f762fff] DMAR: [Firmware Bug]: RMRR entry for device 1a:00.0 is broken - applying workaround</pre> <p><b>Workaround:</b></p> <p>You can ignore the error messages.</p> <p><b>Note:</b> Your system might require an updated BIOS from the manufacturer to fix and eliminate the warnings.</p>
All	SNXT-1537	n/a	<p>If your system is running StorNext version 7.2.0 (or later) and Appliance Controller version 4.1.0, you might receive the following operation failure error when you run the <b>qtmcontroller iscsiadm</b> command, and your card link goes down:</p> <pre>Operation failure: Initiator created, but p3p1.100:isca interface link is down, p3p2.100:iscb interface link is down (E-2003)</pre> <p><b>Workaround:</b></p> <p>To workaround this issue, run the following Appliance Controller command:</p> <pre>qtmcontroller -c iscsiadm network start</pre>

## Quantum UUI Known Issues

The table below lists known issues specific to the Quantum UUI.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	CON-11364	n/a	<p>The <b>CLIENTS</b> page of the UUI displays duplicate clients.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, do the following.</p> <ol style="list-style-type: none"> <li>1. On the <b>APIGW</b> page of the UUI, click <b>Deregister Endpoint</b> for the endpoint of the duplicate client.</li> <li>2. Restart the <b>Unified Connector</b> on the client that contains the duplicate entries.</li> </ol>
All	CON-11487	n/a	<p>When you navigate to the <b>Metrics &gt; Performance and Metrics &gt; Utilization</b> UUI page, you might experience a long delay for the page to display properly.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, refresh your browser page.</p>
All	CON-11504, CON-11505, CON-11506	n/a	<p>If you navigate to the <b>Metrics &gt; Utilization</b> page, and attempt to download a <b>Historical Usage</b> graph, the image file for the graph might not display properly.</p> <p><b>Workaround:</b></p> <p>Do the following to workaround this issue:</p> <ol style="list-style-type: none"> <li>1. Click the <b>Print chart</b> icon in the UUI.</li> <li>2. Save the document as a PDF.</li> </ol>
All	CON-11508	n/a	<p>When you navigate to the <b>Metrics &gt; Performance and Metrics &gt; Utilization</b> UUI page, you might receive an error when you view a widget in full screen mode and attempt to print or export a data chart.</p> <p><b>Example</b></p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>Trouble loading Capacity data. Request aborted.</p> </div> <p><b>Workaround:</b></p> <p>To workaround this issue, close your browser window and then reopen a new browser window.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	CON-11541, CON-11542	n/a	<p>If you navigate to the <b>Metrics &gt; Performance</b> page, and attempt to save any data graph as a PNG file or a JPEG file, the image for the graph contains a black background with black text, which results in an illegible data graph.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, download the data graph in a different format, such as a PDF document, or an SVG vector image.</p>
All	CON-11587	n/a	<p>During a fresh installation of the StorNext 7.2 software on the secondary node of an Xcellis system, the UUI Unified Connector client registration is missing after StorNext is converted to an HA pair. On the top navigation menu of the UUI, click <b>CLIENTS</b> to verify the issue; you might notice a client registration for both the primary node and the secondary node.</p> <p><b>Workaround:</b></p> <p>If you are missing the client registration for the secondary node, Quantum recommends you perform either of the following options to workaround this issue.</p> <p><b>Workaround Option 1</b></p> <p>As the <b>root</b> user, SSH to the secondary node and execute the following command:</p> <pre data-bbox="764 1220 1458 1291">systemctl restart unified-connector</pre> <p><b>Workaround Option 2</b></p> <p>Reboot the secondary node.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	CON-11592	n/a	<p>In the UII (<b>Clients &gt; Software Downloads</b>), the <b>Software Downloads</b> dialog displays a DDM client binary file that is not supported for Rocky 9, Alma 9, and Oracle 9 clients in this release of StorNext.</p> <p><b>Example</b></p> <pre data-bbox="764 562 1458 663">snfs_client_plus_mover_RedHat9_x86_64.bin</pre> <p><b>Workaround:</b></p> <p>To workaround this issue, ignore the DDM client binary file that appears in the UII and do not attempt to download the binary file. The DDM client is not supported for Rocky 9, Alma 9, and Oracle 9 clients in this release of StorNext.</p>
All	CON-11597	n/a	<p>In the UII (<b>Metrics &gt; Dashboards</b>), the error below might appear in your widget when you perform the following steps:</p> <ol data-bbox="776 1003 1422 1339" style="list-style-type: none"> <li>1. On the top navigation menu, click <b>Metrics</b>, and then click <b>Dashboards</b> on the left navigation menu.</li> <li>2. Create a new dashboard.</li> <li>3. Add a widget for a file system to the dashboard you created in <b>Step 2</b>.</li> <li>4. Stop and start the file system.</li> <li>5. On the top navigation menu, click <b>Metrics</b>, and then click <b>Dashboards</b> on the left navigation menu.</li> </ol> <p><b>Example of error</b></p> <pre data-bbox="764 1440 1458 1541">File System &lt;name of file system&gt; is no longer exist.</pre> <p><b>Workaround:</b></p> <p>To workaround this issue, remove the existing widget from dashboard and add a new widget.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	CON-11618	n/a	<p>If you install the UUI software on a customer-supplied system, the following error message appears in the <code>./uui-install.sh</code> script output when the installation completes.</p> <p><b>Example of error</b></p> <pre>Error: failed to remove gateway config: gateway config empty</pre> <p><b>Workaround:</b> You can ignore the error message.</p>
All	CON-11753	n/a	<p>The Quantum UUI <b>Clients &gt; Software Downloads</b> page displays StorNext version 7.2.0 for the clients.</p> <p><b>Workaround:</b> The client download files are actually for StorNext version 7.2.2.</p>

---

## Quantum Appliance Licenses

See [StorNext Licenses](#).

---

## Contacting Quantum Support

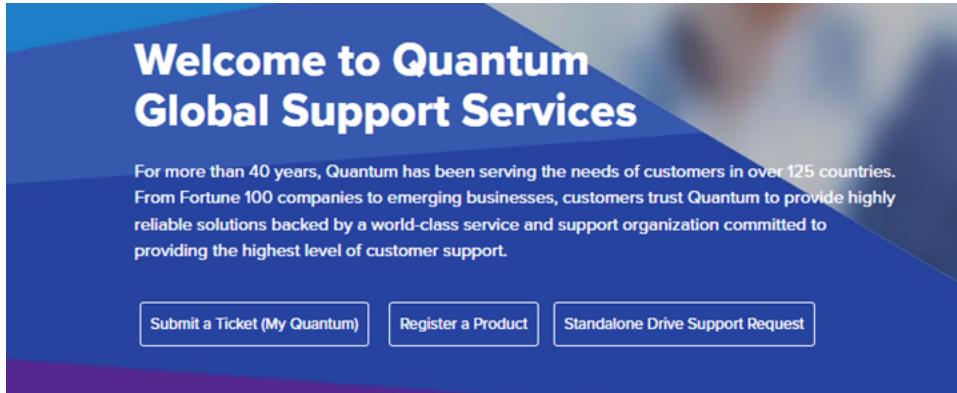
Below is information related to contacting Quantum Support as well as steps to improve your Quantum customer journey.

- [Submit a Ticket \(Service Request\) on the next page](#)
- [Use MyQuantum Service Delivery Platform on the next page](#)
- [Use Cloud Based Analytics \(CBA\) on page 26](#)
- [Escalate a Case on page 26](#)
- [Contact Quantum Sales on page 26](#)

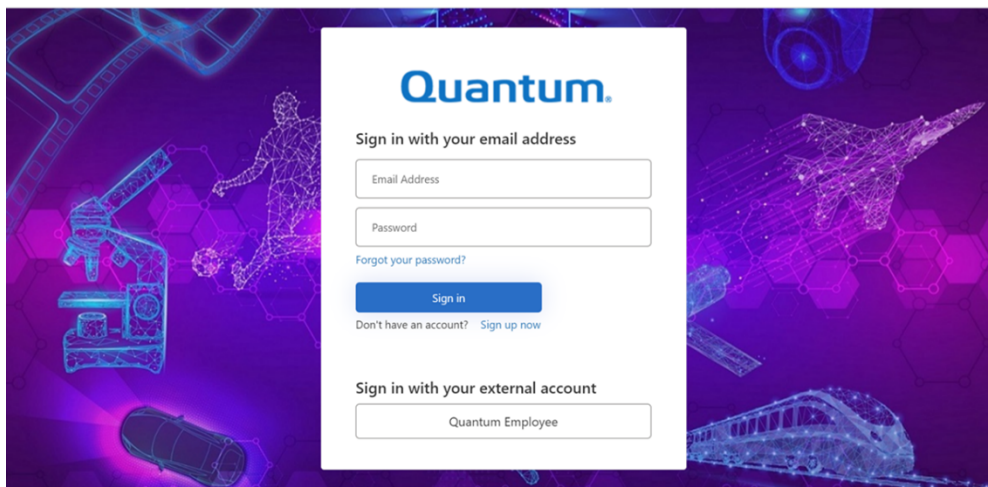


## Submit a Ticket (Service Request)

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/>