



# StorNext 5.3.x M330 Metadata Appliance Release Notes

---

StorNext 5 Releases Supported	Release Notes Updated
StorNext 5 Release 5.3.1	March 2016
StorNext 5 Release 5.2.2	August 2015
StorNext 5 Release 5.2.1	June 2015
StorNext 5 Release 5.2.0	February 2015
StorNext 5 Release 5.1.0	July 2014
StorNext 5 Release 5.0.1	March 2014
StorNext 5 Release 5.0	March 2014

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



---

# Contents

About the StorNext M330 Metadata Appliance .....	2
System Documentation .....	2
About StorNext 5 .....	3
StorNext Release Compatibility – M330 .....	3
StorNext Licenses .....	3
General Notes .....	6
Fixed Issues and Enhancements .....	14
Known Issues .....	19
Known Issues Workarounds .....	24
Considerations - Before You Upgrade Firmware .....	26
How to Upgrade Firmware .....	30
What to do after Firmware Upgrades .....	34
Upgrade Times .....	35
Initiate a Graceful Server Fail-over .....	37
Contacting Quantum .....	43

---

## About the StorNext M330 Metadata Appliance

The StorNext M330 offers the powerful file-sharing capabilities of StorNext in an optimized appliance package. The appliance includes a pair of MDC (metadata controller) nodes in a High Availability (HA) configuration and a high-performance metadata/data array, which is available with HDDs.

---

## System Documentation

The complete list of documentation for M330 Metadata Appliances can be found here:

<http://www.quantum.com/snmdcdocs>

---

# About StorNext 5

StorNext 5 is a new generation of Quantum StorNext that performs faster, scales farther, and expands flexibility. StorNext 5 has been built from the ground up with a new architecture designed to meet the needs of today's evolving digital workflows.

The complete list of documentation for StorNext 5, including the StorNext 5 Release Notes, can be found here:

<http://www.quantum.com/sn5docs>

---

## StorNext Release Compatibility – M330

StorNext Releases																
4.3.2	4.3.3	4.7	4.7.0.1	4.7.1	4.7.2	5	5.0.1	5.1.0	5.1.1	5.2	5.2.0.1	5.2.0.2	5.2.1	5.2.2	5.3	5.3.1
✓	✓	x	✓	✓	x	✓	✓	✓	x	✓	x	x	✓	✓	x	✓

---

## StorNext Licenses

The following StorNext Licenses apply to M330systems.

### StorNext Licenses (Pre-Installed)

The following StorNext licenses are factory-installed, factory-licensed, and included in the base price of M330 systems. No further action is required to enable these factory-installed licenses.

- Note:** Do not mix auto-generated/evaluation and permanent licenses. When one or more permanent licenses are entered in StorNext, all auto-generated/evaluation licenses are deleted. If an auto-generated or evaluation license period is used to evaluate new features, be aware that any remaining time on those licenses is lost as soon as one or more permanent licenses is entered. Maintenance and LAN Client licenses can be mixed with other types of licenses. After permanent licenses are entered, do not install an evaluation license.

## SAN Client

A StorNext **SAN Client** enables a host computer to mount a StorNext File System with direct block level access to the disk arrays via Fibre Channel or iSCSI connectivity. (SAN Clients have direct block-level access to the disk arrays via SAN Fibre Channel or iSCSI connectivity.) The StorNext File System is licensed on a per-client basis. Any machine which directly mounts the file system is considered a client, including the metadata controllers (MDCs).

Includes 10 file system SAN clients available to the user for any supported OS. This count includes one SAN client required for each server node. Additional licenses may be purchased separately.

## LAN Client

StorNext **LAN Clients** use IP-based protocol to read and write data to the StorNext File System through customer-supplied gateways. You must have a LAN Client license for each LAN Client you intend to use with StorNext (in addition to any SAN Client licenses). These licenses are based on operating system type: Linux, Windows or UNIX.

---

**Note:** For Quantum appliances that can operate as StorNext LAN Gateways, LAN Clients are referred to as Gateway Clients, and are not licensed on a "per-seat" basis. See the **Gateway** license.

## Failover (HA)

A Failover HA (High Availability) license enables automated failover from the primary server node to the secondary server node, in the event of a primary server node failure. One HA license is required for each server node pair. This single HA license is applied to both server nodes. In addition to the client licenses applied to each server node, a third client is also required to ensure proper failover functionality. This client can be one that is already accessing a StorNext file system. However, this third client must be purchased separately. See the [StorNext File System Client Mount Requirement on page 11](#).

Includes one HA license. The HA license enables automated failover from the server node operating as the primary to the server node operating as secondary, in the event of a primary server node failure.

## Distributed Data Mover (DDM)

A DDM license is required if you plan to retrieve data using the secondary node of the M330 system. This license is required for each DDM host. The license must be applied to the M330 system.

---

**Note:** Each DDM host also requires one File System SAN Client license in addition to the DDM license.

Includes one DDM license, which allows the secondary server node to function as a DDM host.

## Storage Manager

A Storage Manager license provides full access to the base functionality of StorNext Storage Manager. StorNext Storage Manager is licensed based on the capacity of data stored to secondary tiers of storage (Tape or Storage Disk). Lattus Object Storage and Q-Cloud use Storage Manager technology, but are

licensed separately. See [StorNext Licenses on page 3](#) and [Vaulting below](#) for complete details.

## Maintenance

A Maintenance license is associated with the expiration date of your current service contract. A valid Maintenance License is required to enable StorNext software upgrades. This license is a time-based license key which expires at the same date as the maintenance contract. Each time the maintenance contract is extended, a new maintenance license key is generated.

- i Note:** License enforcement is based on the date of the software to which you are upgraded, regardless of the current date. For example, if your maintenance license key was valid between January 16th, 2010 and January 16th, 2011, you could upgrade to a software version that was released prior to the maintenance license (and associated underlying service contract) expiration date, but not to a version released after January 16th, 2011.

The Maintenance License comes with 12 months of Bronze-level support. At the time of installation, the Maintenance License is generally updated to correspond exactly to the duration of your service contract.

## Gateway Licenses

The following two gateway license types, StorNext LAN Gateway and NAS are supported for various appliances, as described below:

## Additional StorNext Feature Licenses

The following licenses are optional and licensed separately:

Object Storage is licensed based on the capacity of data stored to secondary tiers of Lattus Object Storage media only.

The Object Storage license must be installed and configured on in order to configure StorNext to use Lattus Object Storage as a storage destination.

## Vaulting

A Vaulting license provides the ability to move seldom-used media to a manual archive vault, freeing room for media in the managed archives.

If a vaulting license is purchased, any data that is stored in a vault does not apply to the Storage Manager capacity license.

## Data Replication

A replication license is required if you want to use the StorNext Data Replication feature. Replication is licensed on a per-MDC (or MDC pair) basis. If replication is used between multiple M330 systems, each system must have a replication license.

- Note:** The Storage Manager replication license must be purchased even if the file system being replicated is not managed.

## Other StorNext Feature Licenses

With the exception noted in the next section, all other optional StorNext features are supported, but must be purchased separately

## Unsupported StorNext Feature Licenses

### Data Deduplication

The system is not designed for Deduplication, so the StorNext Deduplication license is not supported on M330systems.

- Note:** Do not purchase or try to install the Deduplication license.

## StorNext Disk Licensing/Certification

See the StorNext Disk Licensing and Disk Certification section of the current *StorNext Licensing Guide* for information about Quantum Branded, Quantum Certified and Quantum Uncertified Storage requirements.

---

## General Notes

This section contains important information you should know about your system.

### Hardware Expansion and Upgrades

The system can come with unfilled expansion slots and drive bays. While server drive expansion is not supported, a metadata array storage upgrade is and a 10 GbE expansion card upgrade kit is available. Hardware upgrade kits require professional installation by Quantum service or an Authorized Service Partner. Additional hardware upgrades are not supported.

### File System Restrictions

The system does not support directly running NFS or CIFS/SAMBA.

## Third-Party Software Support

Quantum support for customer-installed 3rd party software is limited to core system drivers. Example: “EMC Power Path” (metadata array storage driver). Any other installed software operating on the system is not supported.

## Linux Device Mapper Multipath Support

The system supports the standard Linux Dynamic Multipath Mapping driver (DMMP) for all disk storage. For most configurations, the upgrade process will automatically configure the DMMP settings.

The `/etc/multipath.conf` file **MUST** be identical on both the server nodes.

## Multipath Fix for Metadata Array Ping-Pong Condition

If a path failure occurs on one of the MDC nodes in the HA pair while I/O is running, such as a cable disconnect, a ping-pong situation can occur, where connection to the LUN(s) of the metadata array will bounce between the failed path and the good path causing I/O traffic to be paused during these failover conditions. Performance is impacted greatly while this condition persists.

To fix this condition, manually change the settings in the Linux Dynamic Multipath Mapping driver (DMMP) `multipath.conf` file as follows:

1. Open an SSH connection to the appropriate MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.
2. At the command prompt, enter **stornext** for the username.
3. Enter the password for the **stornext** user account. The default password is “password”, but may have been changed after initial configuration.
4. At the prompt, enter **sudo rootsh** to gain root user access.
5. At the prompt, enter the password for the **stornext** user account again.
6. Press **Enter**.
7. Open another ssh terminal session for the other MDC node and repeat the previous steps to login to the command line of that node. Be sure to use the IP address for that node.
8. On the command line of both MDC nodes, type:

```
service cvfs stop
```

9. On the command line of both MDC nodes, type:

```
service multipathd stop
```

10. Modify the `/etc/multipath.conf` file on both MDC nodes for the "LSI" device section to look like:

```
device {
    vendor                "LSI|NETAPP"
    product               "INF-01-00"
    path_grouping_policy  group_by_prio
    getuid_callout        "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout          "/sbin/mpath_prio_rdac /dev/%n"
    path_checker          rdac
    path_selector         "round-robin 0"
    hardware_handler     "1 rdac"
    failback              manual
    no_path_retry         30
    rr_weight             priorities
    features              "2 pg_init_retries 50"
    rr_min_io             16
}
```

11. On the command line of both MDC nodes, type:

```
service multipathd start
```

12. On the command line of both MDC nodes, type:

```
cd /boot
```

13. On the command line of both MDC nodes, type:

```
INITRD=$(ls initrd*.img | head -1)
```



14. On the command line of both MDC nodes, type:

```
KERNEL=$(uname -r)
```

15. On the command line of both MDC nodes, type:

```
mv -f $INITRD $INITRD-org
```

16. On the command line of both MDC nodes, type:

```
/sbin/mkinitrd --omit-raid-modules --omit-lvmodules --without-dmraid --
preload=scsi_dh_rdac -- preload=scsi_mod --preload=sd_mod -- preload=dm_
multipath --preload=dm_round_robin -- preload=sg --preload=sr_mod --
with=scsi_dh_rdac -- with=scsi_transport_fc --with=mpt2sas -v $INITRD $KERNEL
```

17. Reboot both nodes, or type the following on the command line of both MDC nodes:

```
service cvfs start
```

## Change the MDC VIP Address

To change the Virtual IP address for the system, do the following:

- i Note:** Note: This procedure must be used any time the VIP address needs to be changed after initial system configuration, and may include the MAC address of one of the network ports embedded on your system's motherboard.
1. Open an SSH connection to the appropriate MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.
  2. At the command prompt, enter **stornext** for the username.
  3. Enter the password for the **stornext** user account. The default password is "password", but may have been changed after initial configuration.
  4. At the prompt, enter **sudo rootsh** to gain root user access.
  5. At the prompt, enter the password for the **stornext** user account again.
  6. Press **Enter**.
  7. From the command line, update the VIP Address using the following command (refer to the **vip\_control man page** or the *StorNext Man Pages Reference Guide* for information on changing the

address).

**i Note:** For replication VIP use, you will need the MAC addresses to be used for each MDC node.

- a. Enter the following to navigate to the correct directory:

```
cd /usr/cvfs/bin/vip_control
```

- b. Enter the vip\_control command, following the specific variables described here:

```
vip_control -u <vip_str>
```

Syntax of **<vip\_str>** = "<values for MAC 1>;<values for MAC 2>,"

The values for the **<vip\_str>** consist of the following comma-separated values: MAC address, IPV4 VIP, netmask, IPV6 VIP, and then prefix length. Where no value is set, use commas to separate from the other values. A semi-colon separates first and second VIPs and closes the string, followed by an ending straight quotation mark.

Example:

```
vip_control -u
"0030482D38F6,10.0.0.2,255.255.255.0,,;0030482d38f7,10.1.0.2,25
5.255.255.0,,;"
```

In this example, the following are included:

8. To update the firewall rules for the MDC node operating as the secondary, enter the following:

```
/opt/DXi/scripts/netcfg.sh --reset_snvip
```

9. Close the ssh connection for the MDC operating as the primary.
10. Open an ssh connection to the system using the IP for the MDC operating as the secondary.
11. Run the vip\_control on the secondaryMDC with the same values entered earlier on the primary MDC, or copy the `/usr/cvfs/config/ha_vip.txt` file from the primary MDC to the secondary.
12. To update the firewall rules for the MDC node operating as the primary, enter the following:

```
/opt/DXi/scripts/netcfg.sh --reset_snvip
```

The procedure is complete.

## Memory Requirements

The memory allocation settings presented in the table below are general recommendations. Other settings may be reasonable, depending on specific conditions and workloads.

**Table 1:** Memory Allocation Settings

Total System Memory	FSM BufferCacheSize	FSM nodeCacheSize*	MySQL innodb_buffer_pool_size	Reserved for General Use
24 GB	2 GB per file system (e.g. 4 x 2 GB = 8 GB)	256 K inodes per file system (e.g. 4 x ~ 0.5 GB = 2 GB)	5 GB (default is 2 GB)	18 GB

\*For systems not running the maximum number of file systems, it is recommended that memory be provisioned to increase the FSM BufferCacheSize setting to a value up to 8 GB and the FSM InodeCacheSize setting to a value up to 512 K. Increasing these values may improve file operation performance on systems with many concurrently active files.

For information about changing memory allocation settings, see the following sections in the *StorNext 5 File System Tuning Guide*:

- MySQL innodb\_buffer\_pool\_size
- BufferCacheSize
- InodeCacheSize

## SNAPI Security Restriction

SNAPI should only be allowed from the IP address(es) of SNAPI clients or a StorNext AEL Archive. For instructions on how to set this up, contact Quantum Support.

## StorNext File System Client Mount Requirement

To prevent a split-brain condition between the HA pair of the server nodes, at least one additional StorNext file system client must mount the HA file system. This will allow the additional client to "vote" in the event of a split-brain condition.

**Note:** You need to do this on only one client machine, and you must mount the path for the HA shared file system on the client as read-only.

## Linux

On Linux systems, add an entry to the `/etc/fstab` file on the StorNext file system client similar to the following:

```
shared-SV1249CKD2943502637 /stornext/shared cvfs ro,diskless=yes 0 0
```

The name "shared-SV1249CKD2943502637" used in the example will vary. The format is "shared-NNNNNNNNNNNNNNNN".

## Windows

On Windows clients, use the Mount Options field to add "ro,diskless=yes".

For more information about this procedure, see the HA chapter in the *StorNext User's Guide*.

## Target Reset and Fibre Channel Tape Support on Qlogic HBAs

The Enable SCSI Bus Target Reset parameter is enabled by default on all Fibre Channel HBAs. The parameter exists for disk arrays, but poses a problem for tape drives.

### Problem

When the SCSI bus target (the tape drive) is reset when a backup job is running, the backup job may abort. If the tape drive does not receive the rewind and unload commands from the backup job, it leaves the tape in the drive. This may cause the drive to be seen as not ready, and then be marked offline in the backup application when the next job tries to use the drive.

### Solution

To disable Target Resets on the tape SAN port of the system, do the following commands on both nodes.

1. Open an SSH connection to the appropriate MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.
2. At the command prompt, enter `stornext` for the username.
3. Enter the password for the `stornext` user account. The default password is "password", but may have been changed after initial configuration.
4. At the prompt, enter `sudo rootsh` to gain root user access.

5. At the prompt, enter the password for the **stornext** user account again.
6. Press **Enter**.
7. Repeat the steps above to create a separate SSH session for the other node.
8. Disable Target Resets on the tape SAN port by entering the following command:

```
/opt/QLogic_Corporation/QConvergeConsoleCLI/scli -n 1 TR 0
```

9. Confirm that the setting is correct by entering the following command:

```
/opt/QLogic_Corporation/QConvergeConsoleCLI/scli -c
```

10. Compare the output for Port 2. It should look similar to the following example.

**Note:** <Qlogic model> in the example could be QLE2564L or QLE2662.

```
[root@Acadia1-1 scripts]# scli -c
-----
-----
HBA Instance 0: <Qlogic model> Port 1 WWPN 21-00-00-1B-32-9D-4A-8D PortID 00-00-
00
-----
Connection Options                : 2 - Loop Preferred, Otherwise Point-to-Point
Data Rate                         : Auto
Frame Size                        : 2048
Hard Loop ID                      : 0
Loop Reset Delay (seconds)       : 5
Enable Host HBA BIOS             : Disabled
Enable Hard Loop ID              : Disabled
Enable FC Tape Support           : Enabled
Operation Mode                   : 0 - Interrupt for every I/O completion
Interrupt Delay Timer (100ms)    : 0
Execution Throttle                : 65535
Login Retry Count                 : 8
```

```
Port Down Retry Count           : 30
Enable LIP Full Login           : Enabled
Link Down Timeout (seconds)     : 30
Enable Target Reset             : Enabled
LUNs Per Target                 : 128
Enable Out Of Order Frame Assembly : Disabled
```

-----  
-----  
HBA Instance 1: <Qlogic model> Port 2 WWPN 21-01-00-1B-32-BD-4A-8D PortID 00-00-00

-----  
-----

```
Connection Options              : 2 - Loop Preferred, Otherwise Point-to-Point
Data Rate                       : Auto
Frame Size                      : 2048
Hard Loop ID                    : 0
Loop Reset Delay (seconds)      : 5
Enable Host HBA BIOS            : Disabled
Enable Hard Loop ID             : Disabled
Enable FC Tape Support          : Enabled
Operation Mode                  : 0 - Interrupt for every I/O completion
Interrupt Delay Timer (100ms)   : 0
Execution Throttle              : 65535
Login Retry Count               : 8
Port Down Retry Count           : 30
Enable LIP Full Login           : Enabled
Link Down Timeout (seconds)     : 30
Enable Target Reset             : Enabled
LUNs Per Target                 : 128
Enable Out Of Order Frame Assembly : Disabled
```

---

## Fixed Issues and Enhancements

This section lists the fixed issues that affect M330systems for different StorNext Releases.

## Fixed issues and enhancements in StorNext 5 Release 5.3.1

The following table lists the fixed issues/enhancements for this StorNext Release.

See the StorNext 5 Release Notes for information about StorNext software updates for this release here:

<http://www.quantum.com/sn5docs>

**Table 2:** Fixed Issues/Enhancements for StorNext 5 Release 5.3.1

CR Number	SR Number(s)	Description
61068	n/a	Fixed an issue which could have caused the firmware upgrade process to fail.
61070	n/a	Fixed minor spelling errors in the <b>Service Menu</b>
61073	3565236	Fixed an issue where hwmoud previously didn't catch a failed PSU if the return string contained the extra words "Failure detected".
61089	n/a	Enhanced the upgrade process to ensure custom <code>innodb_buffer_pool_size</code> memory settings for MySQL are preserved after firmware upgrades.
57627	n/a	After the system is booted, if NICs were present in slots that previously contained NICs with different speeds (Example: 10 GbE vs. 1 GbE), the system will result in the same Ethernet alias names defined for the network interfaces as the previously-installed NICs.  Ethernet alias names shown in the StorNext Metrics GUI page now reflect the change in network device type representing the alias.

## Fixed issues and enhancements in StorNext 5 Release 5.2.2

The following table lists the fixed issues/enhancements for this StorNext Release.

See the StorNext 5 Release Notes for information about StorNext software updates for this release here:

<http://www.quantum.com/sn5docs>

**Table 3:** Fixed Issues/Enhancements for StorNext 5 Release 5.2.2

CR Number	SR Number(s)	Description
56290	n/a	Updated metadata array firmware to 08.20.09.00.

CR Number	SR Number(s)	Description
56291	n/a	Upgrade CentOS to version 6.6 when the StorNext 5 appliance firmware upgrade is applied.
56583	n/a	Added a way to insert global variables into <code>smb.conf</code> through the command line.
56717	n/a	Added <code>/var/log/sa</code> contents to support bundle file-list for SAR reports.

## Fixed issues and enhancements in StorNext 5 Release 5.2.1

The following table lists the fixed issues/enhancements for this StorNext Release.

See the StorNext 5 Release Notes for information about StorNext software updates for this release here:

<http://www.quantum.com/sn5docs>

**Table 4:** Fixed Issues/Enhancements for StorNext 5 Release 5.2.1

CR Number	SR Number(s)	Description
55553	1628792	Updated the default <code>innodb_buffer_pool</code> size for appliances. This change improves Storage Manager performance for systems running the previous defaults. The new default value is: M330 - 2GB
55567	3511046, and 3512842	Changed the <b>sar</b> log time format to improve compatibility with <b>sar</b> graphing tools.
55568	3488100	Added a “screen” utility to the StorNext appliance code to aid troubleshooting and configuration.
55585	3515176	Added the Ubuntu client to the client download drop-down menu in the StorNext GUI.
55593	n/a	The default gateway is now set correctly when no DNS server is specified when configuring the network settings in the <b>Service Menu</b> .
55709	n/a	ADATA USB drives are now properly recognized when plugged into an appliance.

## Fixed issues and enhancements in StorNext 5 Release 5.2.0

The following table lists the fixed issues/enhancements for this StorNext Release.



See the StorNext 5 Release Notes for information about StorNext software updates for this release here:

<http://www.quantum.com/sn5docs>

**Table 5:** Fixed Issues/Enhancements for StorNext 5 Release 5.2.0

CR Number	SR Number(s)	Description
54108	n/a	Fixed Ghost vulnerability in Linux glibc (CVE-2015-0235).
51295	n/a	Fixed issue where the upgrading the G300 from the MDC would not work.
53765	n/a	Eliminated erroneous RAS tickets generated during the upgrade process.

## Fixed issues and enhancements in StorNext 5 Release 5.1.0

The following table lists the fixed issues/enhancements for this StorNext Release.

See the StorNext 5 Release Notes for information about StorNext software updates for this release here:

<http://www.quantum.com/sn5docs>

**Table 6:** Fixed Issues/Enhancements for StorNext 5 Release 5.1.0

CR Number	SR Number(s)	Description
48639	3398958, and 3342200	Optimized MDC node boot speed by reducing log level during bootup.
47343	n/a	Added a capacity check for the HA shared file system to ensure sufficient disk space is available for the upgrade.
47344	n/a	Added new checkpoints within the upgrade process to increase the robustness of the upgrade in the event of an error.
47358	n/a	Placed rootsh logs under logrotate control and compression to avoid unnecessary disk usage.
47379	n/a	Updated the firmware levels for the metadata controller nodes.
47408	n/a	Updated the multipath.conf.quantum file.
47565	n/a	Add version table support for 5.1 upgrades on all appliances.

CR Number	SR Number(s)	Description
47854	n/a	Resolved an issue during upgrade where the upgrade could hang if the secondary node was manually rebooted.
48000	n/a	Resolved an issue where ALUA support may not be correctly detected during the upgrade process.

## Fixed issues and enhancements in StorNext 5 Release 5.0.1

The following table lists the fixed issues/enhancements for this StorNext Release.

See the StorNext 5 Release Notes for information about StorNext software updates for this release here:

<http://www.quantum.com/sn5docs>

**Table 7:** Fixed Issues/Enhancements for StorNext 5 Release 5.0.1

CR Number	SR Number(s)	Description
35600	n/a	Add ALUA support on the M-Series arrays.
36891	1634032, and 3352678	Updated metadata array multipath settings for optimal performance and reliability.

## Fixed issues and enhancements in StorNext 5 Release 5.0

The following table lists the fixed issues/enhancements for the initial Release of StorNext 5.

See the StorNext 5 Release Notes for information about StorNext software updates for this release here:

<http://www.quantum.com/sn5docs>

**Table 8:** Fixed Issues/Enhancements for StorNext 5 Release 5.3.1

CR Number	SR Number(s)	Description
29296	n/a	Log message:  syncha.pl[24735]: Relocating shared: 'SRVCL0G/logs/srvcLog'  no longer repeats after a SAS failover on the metadata array.
34846	n/a	Upgrade CentOS to version 6.4 when the StorNext 5 appliance firmware upgrade is applied.

CR Number	SR Number(s)	Description
34847	n/a	Upgrade output log files are now consolidated into a single file during software upgrades.
35536	n/a	StorNext 5 appliance firmware upgrade includes the latest component-level firmware.
35602	n/a	HA failover no longer causes both nodes to be down during an upgrade. Changed the upgrade behavior to a “no client downtime” upgrade when upgrading from 4.7 (and later) to StorNext 5.
35612	n/a	Network restarts from the metadata appliance <b>Service Menu</b> no longer cause bond0 to not start up.
35617	n/a	Manually-configured VLAN settings are now preserved in StorNext 5.
35652	n/a	Custom network bonding configurations are now preserved when upgrading from CentOS6 and CentOS5 platforms.
35839	n/a	RAS messages are no longer generated when a health check is called on a system with no configured managed file systems.

---

## Known Issues

This section lists the known issues that affect M330systems.

CR Number	SR Number	Description	Workaround (if applicable)
n/a	n/a	If you upgrade the system to StorNext 5 Release 5.2.2 or later (which upgrades the metadata array firmware to the Kingston release [v. 08.20.09] on the array), LUNs created for the Expansion Unit will be 5 GB smaller than LUNs created from systems running older StorNext releases.	There currently is no workaround for this issue. LUN sizes cannot be adjusted.

CR Number	SR Number	Description	Workaround (if applicable)
62454	n/a	<p>On upgrades to StorNext 5 Release to 5.3.1 from StorNext releases prior to 5.0, the directory <b>/usr/adic/wsar_agent/tmp</b> is not created on the nodes during the upgrade process. Because of this, async web services will not function correctly.</p> <p>A log message similar to the following will be displayed in the <b>wsar_agent.log</b> file:</p> <pre>&gt; [0328 13:05:35.878] 0x7f36b1346700 ERR [wsarutils.c:107] [wsar_ run_cmd] [22] wsar_run_ cmd: req_id=22: failed to open(/usr/adic/wsar_ agent/tmp/wsar.22.out), errno 2</pre>	<p>To create the <b>/usr/adic/wsar_agent/tmp</b> directory on both nodes:</p> <ol style="list-style-type: none"> <li>1. Open an SSH connection to the appropriate MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.</li> <li>2. At the command prompt, enter <b>stornext</b> for the username.</li> <li>3. Enter the password for the <b>stornext</b> user account. The default password is "password", but may have been changed after initial configuration.</li> <li>4. At the prompt, enter <b>sudo rootsh</b> to gain root user access.</li> <li>5. At the prompt, enter the password for the <b>stornext</b> user account again.</li> <li>6. Press <b>Enter</b>.</li> <li>7. Verify the directory exists on the server node: <pre>cd /usr/adic/wsar_ agent/tmp</pre> </li> <li>8. If the directory does not exist, you will see a message similar to the following: <pre>bash: cd: /usr/adic/wsar_ agent/tmp: No such file or directory</pre> </li> <li>9. If needed, create the directory:</li> </ol>

CR Number	SR Number	Description	Workaround (if applicable)
60814	3635792	When hosted applications are run on a server node and are active inside of an SNFS file system, StorNext restarts may hang, requiring the server node to be rebooted.	<pre data-bbox="1045 296 1455 401">mkdir /usr/adic/wsar_agent/tmp</pre> <p data-bbox="984 443 1455 512">10. Repeat the previous steps for the other server node.</p>
60774/ 55220	n/a	(NAS-only issue) NFS version 4 is not supported and must be disabled.	<p data-bbox="984 800 1455 905">There currently is no workaround for NFS v4 support. See <a href="#">Workaround: Disabling NFS v4 on page 24</a>.</p> <p data-bbox="984 926 1455 1052">If you are running StorNext NAS on and export NFS shares, you must disable NFSv4 for systems running StorNext 5 Release 5.3.0 or later.</p>
60614	n/a	, and all other Connect-managed Linux StorNext SAN clients must have the latest Connector installed so that statistics can be passed to StorNext Connect. If a system is managed by StorNext Connect and you upgrade the firmware on that system to StorNext 5 Release 5.3.0 firmware <b>PRIOR</b> to upgrading the StorNext Connect Connector, the Volume Storage widget on the StorNext Connect Dashboard will display no data for those systems.	<p data-bbox="984 1079 1455 1388">For steps to take to update the Connect Connector(s) before upgrading system firmware, see <a href="#">Considerations - Before You Upgrade Firmware on page 26</a>. For steps to take if you have already upgraded firmware but did not first update the Connector(s), see <a href="#">Considerations - Before You Upgrade Firmware on page 26</a>.</p>

CR Number	SR Number	Description	Workaround (if applicable)
56135	n/a	<p>The StorNext GUI does not show the RHEL7 and SLES12 clients in the list of clients available for download.</p> <p><b>i Note:</b> This issue only affects systems running StorNext 5 Release 5.2.1.</p>	<p>To manually download the client installers for RHEL7 and SLES12:</p> <ol style="list-style-type: none"> <li>1. Open an SSH connection to one of the server nodes (either node will work) using the IP address assigned to that node on the Metadata network. Manually copy the Redhat7 or <b>SuSE12 .bin</b> file from <b>/usr/cvfs/CLIENTS</b> to an external USB thumb drive or copy over the network to the client system.</li> <li>2. Continue with the installation procedure for the client for your operating system as described in the <i>StorNext Installation Guide</i> or the StorNext online help.</li> </ol>
54451	n/a	<p>StorNext supports case-sensitive file names. For configurations with different client types, such as Windows and Mac sharing the same files, the default case type may be different.</p>	<p>There currently is no workaround for this issue. SMB is operating as expected.</p>
47041		<p>Adding new index to tierfiles tables can delay first TSM start up following system StorNext upgrades.</p>	<p>See <a href="#">TSM Indexing delay for large databases on page 29</a> in the <b>Before you Upgrade Firmware</b> section for the workaround.</p>
45702	n/a	<p>If you replace an HDD drive with an SSD or vice versa, the StorNext GUI will show a status of "Missing" and an equivalent RAS ticket instead of displaying an "Incompatible" status.</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>

CR Number	SR Number	Description	Workaround (if applicable)
38291	n/a	After an HA failover, an Admin Alert is issued if the MDC node acting as primary attempts to initiate an fs_fmover process on the MDC node acting as secondary while the secondary MDC node is rebooting. This situation applies to all StorNext 5 releases.	See <a href="#">Workaround: Re-enable DDM on the secondary node after HA failover on page 25</a> .
38128	1395540	Using the GUI while a large Media import is kicked off via the command line can cause the StorNext GUI to timeout or crash.	Wait until a bulk load from tape is finished prior to opening the StorNext GUI.
37538/ 36626	1398524	GUI is unable to down a stripe group when LUNs are unavailable	Mark stripe groups down in the GUI before taking the stripe group's disks offline. If that is not possible, set the stripe group down directly through the FSM configuration file and restart the FSM. See the <b>snfs_config(5)</b> man page or the <b>StorNext MAN Pages Reference Guide</b> for details.
37166	n/a	The sn_metrics database tables are not installed in the MySQL database when upgrading to StorNext 4.2.1.0.1 from an earlier version.	<p>After upgrading to 4.2.1.0.1 or later, start StorNext Storage Manager and run the sngateway_install_mysql_tables script as follows:</p> <pre>/usr/adic/.profile service cvfs start /usr/cvfs/install/sngatewa y_install_mysql_tables.pl</pre> <p>Optionally, you can verify that the sn_metrics tables have been installed by running mysqlshow as follows:</p> <pre>/usr/adic/mysql/bin/mysqls how sn_metrics</pre>

CR Number	SR Number	Description	Workaround (if applicable)
37916	n/a	Admin alerts are generated for network or FC ports that are disconnected but are configured in the system.	The only way to prevent these alerts from displaying is to remove the network or FC ports that are disconnected from your configuration, unless the ports will only be down temporarily.

## Known Issues Workarounds

### Workaround: Disabling NFS v4

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

1. Open an SSH connection to the appropriate MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.
2. At the command prompt, enter **stornext** for the username.
3. Enter the password for the **stornext** user account. The default password is “password”, but may have been changed after initial configuration.
4. At the prompt, enter **sudo rootsh** to gain root user access.
5. At the prompt, enter the password for the **stornext** user account again.
6. Press **Enter**.
7. Edit **/etc/sysconfig/nfs** file. (See [Example: Use vi to edit the nfs file on the next page](#))
8. Change the value for **RPCNFSDARGS** to “-N 4”. When finished, the line must be:  
**RPCNFSDARGS="-N 4"**
9. Change the value for **RPCMOUNTDOPTS** to “-N 4”. When finished the line must be:  
**RPCMOUNTDOPTS="-N 4"**
10. Save the file.
11. Enter the following to restart the NFS configuration:

```
service nfs-config restart
```



12. Enter the following to restart the NFS server:

```
service nfs-server restart
```

13. Close the SSH session for the node.
14. Repeat all of the above steps on node 1.

If you need to create new NFS shares to export for user access, you may do this now. See the "Manage NAS Clients" section of the *Xcellis User's Guide* for information about NAS configuration in StorNext Connect.

#### Example: Use vi to edit the nfs file

- a. Enter the following:

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

- b. Move the cursor to the closing quotation mark in RPCNFSDARGS.
- c. Enter the following:

```
i -N 4
```

- d. Write the file and quit vi as follows:

```
:wq
```

- e. Go back to [Step](#) of the procedure.

## Workaround: 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 MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.
2. At the command prompt, enter `stornext` for the username.
3. Enter the password for the `stornext` user account. The default password is "password", but may have


been changed after initial configuration.

4. At the prompt, enter **sudo rootsh** to gain root user access.
5. At the prompt, enter the password for the **stornext** user account again.
6. Press **Enter**.
7. 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, Quantum recommends you to contact Technical Support.

---

## Considerations - Before You Upgrade Firmware

### About specific upgrades

Please review the following if it pertains to your specific upgrade situation:

- Before you upgrade, you will need to obtain the firmware upgrade files. See [Obtain the firmware upgrade files on the next page](#).
- Not all StorNext releases may be upgraded to a given StorNext release. As a result, an upgrade to the current version of StorNext may require multiple, incremental upgrades, depending on the version of StorNext currently installed. For information about supported upgrade paths for StorNext 5, consult the *StorNext 5 Compatibility Guide*. If your system is running a StorNext release prior to the supported

upgrade releases for a given StorNext 5 Release, consult an earlier version of the *StorNext Compatibility Guide* that applies to your specific upgrade.

- [M330 upgrade dependency for StorNext 5 Releases 5.1.0 and later on the next page](#)
- Upgrades to StorNext 5 Release 5.1.1, 5.1.0.1, and 5.3.0 are not supported for the M330.
- Firmware Upgrade time considerations. See [Upgrade Times on page 35](#).
- Pre-upgrade consideration for Large Databases - affects multiple StorNext Upgrades. See - [TSM Indexing delay for large databases on page 29](#).

## Obtain the firmware upgrade files

**Note:** The two files are large - around 2 GB total, so plan time to download the files for the upgrade.

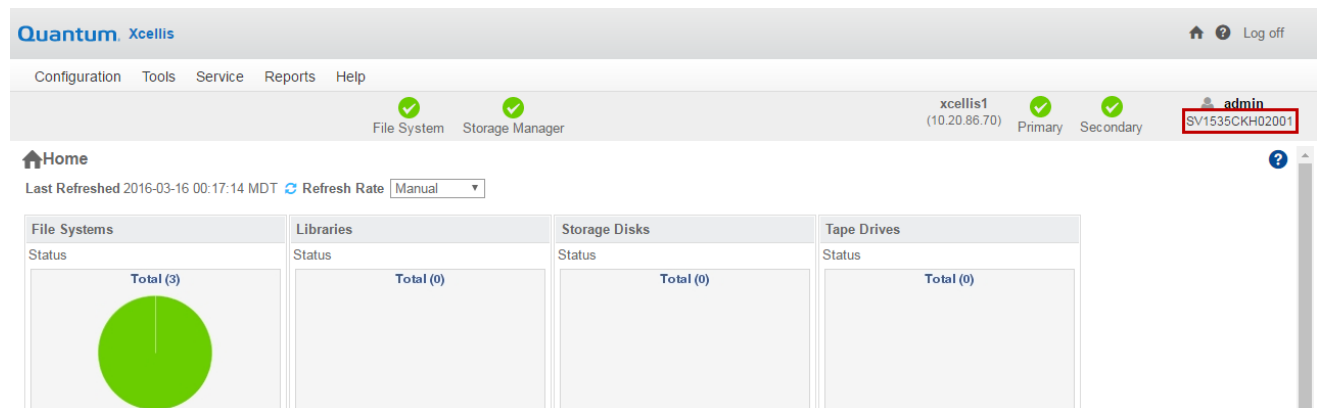
### Firmware request pre-requisites

In order to receive a license.dat file, you will need the following:

- The System Serial Number. (Use this in the "Product Information, Serial Number of the Original Media" section of the form.)

System serial numbers are alpha-numeric (example: SV1728CKH02059).

You can find the System Serial Number on the right-hand side of the StorNext GUI, as shown here:



- You could also use the StorNext Serial Number instead. (Use this in the "Product Information, Serial Number of the Original Media" section of the form.) You can find the StorNext Serial Number in the license.dat file. The license.dat file can be found on the server node at /usr/cvfs/config/license.dat. Open the license.dat file and locate the **Serial Number**.

Example:

```
# Serial Number: Q8675309
```

- The cvfsid string for each server node.

Example:

```
ECF4BCEECC0E linux 0 xcellis13
```

### How to Generate a cvfsid String

The cvfsid command is needed to generate StorNext license files. Here's how to generate the cvfsid on the command line of the MDC node:

1. Open an SSH connection to the appropriate MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.
2. At the command prompt, enter **stornext** for the username.
3. Enter the password for the **stornext** user account. The default password is "password", but may have been changed after initial configuration.
4. At the prompt, enter **sudo rootsh** to gain root user access.
5. At the prompt, enter the password for the **stornext** user account again.
6. Press **Enter**.
7. Enter the **cvfsid** command.

The cvfsid string provides the following:

```
<node MAC address> <node OS> <node name>
```

Example:

```
ECF4BCDECC0E linux 0 xcellis13
```

8. Close the SSH session for the node. Repeat for the other node.

### Firmware request

If you need firmware upgrade files to upgrade your system:

1. Enter the required information about your system at:  
<http://www.quantum.com/ServiceandSupport/License/StorNext/Index.aspx>

---

**i Note:** If you cannot access the web page or need additional help filling out the form, contact Quantum Support at <http://www.quantum.com/ServiceandSupport>.

2. After the Quantum Technical Assistance Center receives the above information, a representative will send you an ftp link to download the firmware files. Save the files to a temporary location on the system you will use to begin the firmware update.

## M330 upgrade dependency for StorNext 5 Releases 5.1.0 and later

To upgrade to StorNext releases later than StorNext 5 Release 5.1.0 on an M330, you must first upgrade to

5.1.0. For example, an upgrade from 4.3.3 to 5.3.x would require the following: **4.3.3 > 5.1.0 > 5.3.x**.

---

**i Note:** Upgrades to StorNext 5 Release 5.1.1 and 5.1.0.1 are not supported for the M330.

## TSM Indexing delay for large databases

### Issue

A database index named `classndxatimeme` will be automatically added to the `tmdb.tier000files%` and `tmdb.tier001files%` tables upon starting TSM for the first time after upgrading to specific StorNext 5 releases:

---

**! Caution:** For upgrades from 4.3.2, 4.3.3, 4.7.0, 4.7.0.1, 4.7.1, and StorNext 5 Release 5.0, **DO NOT** use this script for direct upgrades to StorNext 5 Releases 5.0.1, 5.1, 5.2, 5.2.1, 5.2.2, and 5.3.1.

This index improves the performance of certain operations such as truncation policies. However, the creation of this index can take multiple hours for very large databases. TSM will be unavailable after upgrading until the indexing has completed.

### Workaround

To minimize TSM downtime after upgrade, the `classndxatimeme` index can be created prior to performing the upgrade using the `index_tierfiles.pl` PERL script. This file is available by opening a support ticket and requesting the file. (Quantum service and service partners can obtain this file from the StorNext Metadata Appliances page on CSWeb.) The script can be run while TSM is running, although it may impact the performance of other operations while the index is being added to the database.

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

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

1. Open an SSH connection to the appropriate MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.
2. At the command prompt, enter **stornext** for the username.
3. Enter the password for the **stornext** user account. The default password is “password”, but may have been changed after initial configuration.
4. At the prompt, enter **sudo rootsh** to gain root user access.
5. At the prompt, enter the password for the **stornext** user account again.
6. Press **Enter**.
7. Copy the `index_tierfiles.pl` file to the `/tmp` directory on the primary server node.

8. Enter the following to source the profile:

```
. /usr/adic/.profile
```

9. Verify that the database is up by running:

```
mysql_control start
```

10. Execute the PERL script:

```
./  
tmp/index_tierfiles.pl
```

The procedure is complete.

---

## How to Upgrade Firmware


The **Firmware Upgrade** option in the StorNext GUI allows you to perform a firmware upgrade on Metadata Appliance systems. Depending on the version being applied, the firmware upgrade includes updates to the firmware running on the appliance servers (if applicable), the appliance OS (if applicable) and StorNext software. Firmware for the metadata array (if needed for the upgrade) is also installed.

 **Caution:** Do not use the *StorNext Upgrade Guide* or the *StorNext Installation Guide* to perform the upgrade. Only use the procedure listed here.

## Pre-upgrade considerations

Before you upgrade firmware, see [Considerations - Before You Upgrade Firmware on page 26](#).

## Upgrade Procedure

 **Note:** For information about estimating how long the upgrade might take for your given upgrade, see [Upgrade Times on page 35](#).

1. Log into the StorNext GUI.
2. Choose **Tools > Firmware Upgrade**. The Firmware Upgrade page displays.  
Select **Auto Upload** to upload the file immediately after you select it.
3. Click **Choose File** and navigate to the directory where the firmware files reside. Firmware files are identifiable by the .fw extension.

---

**i Note:** There are two .fw files required for updating firmware. The filenames are similar to QOTM-DXiSNA-upd-5.3.0.OS7-16202-16195.1of2.fw and QTM-DXiSNA-upd-5.3.0.OS7-16202-16195.2of2.fw. Since it is a two-part upgrade, upload both files to the GUI.

---

**i Note:** Uploading the firmware upgrade files in a network with low latency should only take a matter of minutes. High network latency in your environment can slow the upload of these files onto the system.

---

**i Note:** When using the firmware upgrade process from the StorNext GUI, the license for the system will be automatically applied to the system.

If you selected **Auto Upload**, the file is immediately uploaded. Proceed to the next step.

If you did not select **Auto Upload** and want to validate the file before uploading, click **Validate**. After a message informs you that the file is valid, click **Upload**.

---

**i Note:** Files are automatically validated after you click Apply in the next step, but you won't receive a message telling you the file is valid.

---

**i Note:** The green status indicator at the top of the page indicates upload progress, not the upgrade progress.

4. Select either of the uploaded files and click **Apply**. Click **Yes** to confirm the installation. Both firmware files are applied to the system during the upgrade.
5. If you see a message about clearing Tickets or Admin Alerts, click **OK** to continue the upgrade.

## Monitoring upgrade progress

While the system firmware upgrade is being applied, you will not see something that indicates upgrade progress. There are, however, a few different things that indicate the system is being upgraded.

### GUI indications that a firmware upgrade is in process

For firmware upgrades supporting this StorNext release:

- The StorNext GUI will display a message similar to the following when the upgrade begins:

```
Firmware upgrade in progress. Login session will be terminated when the upgrade
```

process has completed on the other node. Log in to the other node to continue monitoring the upgrade process. The upgrade is complete when full GUI functionality has been restored.

- The normally-green Secondary node icon changes to orange, indicating that the process is currently upgrading node 2.
- After a while, the StorNext GUI on node 1 will cease to function.
- Once the upgrade is complete on node 2, StorNext operations will fail over to node 2, and you can login to the node 2 GUI and monitor the upgrade progress.

## Command line indications that a firmware upgrade is in process

If you SSH and login to the command line of the node of a server node that is currently operating as primary while the server node currently acting as secondary is being upgraded, you will either attempt to login and the session will terminate immediately after, or will see a note similar to the following, indicating the system firmware upgrade is in process:

```
***** WARNING *****  
*** An upgrade is currently in progress on this system! ***  
*** Please do not change any settings or software until ***  
*** the upgrade completes. ***  
***** WARNING *****
```

If you attempt to SSH into the server node that is currently upgrading its firmware, the session will not connect to the system.

## Log file indications that a firmware upgrade is in process

You can monitor the upgrade process using the following log files:

Log File	Description
<code>/var/log/DXi/upgrades/NodeX.upgradeoutput.log</code>	Upgrade output log for the main part of the upgrade (not including NetApp firmware updates).
<code>/var/log/DXi/upgrades/NodeX.upgrade_progress</code>	Upgrade Progress Log for the main part of the upgrade (not including NetApp firmware updates).



Log File	Description
<code>/var/log/DXi/baseos.log</code>	Final upgrade progress at the end of the final boot of the second node (shows the NetApp firmware update has started).
<code>/var/log/DXi/UpgradeArrays.log</code>	Contains the actual NetApp firmware upgrade progress. The firmware update command will not respond for 30-40 minutes so there is no specific change in progress noted until the controllers reboot.
<code>/var/log/messages</code>	Contains other boot messages like the application of firmware bundle updates and hardware detect script execution, both of which could take several minutes.

Each of these log files will continue to log upgrade status progress detail until the StorNext GUI is once again fully-accessible at the end of the upgrade.

**i Note:** SNFS services are stopped prior to updating the controller firmware, so the StorNext GUI does not show the array(s) as available while the controller firmware is updated.

The end of the `NodeX.upgradeoutput.log` includes the following statement that indicates the array firmware update has started:

```
New NetApp FW needs to be activated. Bringing down the cluster.
```

The message below is added to the log after the array controllers reboot:

```
Successfully activated new NetApp FW on the array controllers.
```

These are the only messages provided on the upgrade progress for NetApp controllers.

The array controllers will reboot after the firmware update is complete (about 20-30 minutes after the update begins). Messages about the SCSI driver and multi-path driver devices are removed and added are displayed during the upgrade.

These occurrences are normal since the controllers drop off the SAS BUS and are added again during the update. LUN access is unavailable during the controller reboot.

**i Note:** SNFS services will start up automatically after the upgrade is completed and all LUNs will be accessible on the metadata array.

## Post-upgrade considerations

- [What to do after Firmware Upgrades below](#)
- [Post-Upgrade Browser Refresh below](#)
- [Initiate a Graceful Server Fail-over below](#)
- [What to do after Firmware Upgrades below](#)

---

# What to do after Firmware Upgrades

After the firmware upgrade is complete, do the following to complete the process:

- [Initiate a Graceful Server Fail-over below](#)
- [Post-Upgrade Browser Refresh below](#)

## Initiate a Graceful Server Fail-over

If you need to fail-over server operations from a server node operating as primary to a server node operating as secondary without losing client access to metadata. See [Initiate a Graceful Server Fail-over on page 37](#) for the procedure.

## Post-Upgrade Browser Refresh

Quantum recommends to clear the browser and cookie cache after upgrading to a StorNext release before logging into the GUI on the upgraded system.

The steps used to clear the cache depend on the browser:

### For Google Chrome

1. Click the configuration icon and the **Settings option**.
2. Click the **Show advanced settings** option.
3. In the **Privacy** section, click **Clear browsing data**.
4. Select the **the beginning of time** option in the **Obliterate the following items from** menu.
5. Select the **Clear browsing history**, **Empty the cache**, and **Delete cookies and other...** options.
6. Click **Clear browsing data**.
7. Close the **Settings** tab.

### For Internet Explorer

1. Go to **Tools > Internet Options**.
2. On the **General** tab, look in the **Browsing History** section and click **Delete**.
3. Select the **Temporary Internet Files** and **Cookies** options.
4. Click **Delete**.
5. Click **OK** to close the **Internet Options** dialog.

#### For Mozilla Firefox

1. Click **Tools > Clear Recent History**. (If this option is not visible, right-click the top of the browser window and select the **Menu Bar** option).
2. In the **Time Range to Clear** menu, select **Everything**.
3. Click **Details**, and select the **Browsing & Download History**, **Cookies**, and **Cache** options.
4. Click **Clear Now**.

---

## Upgrade Times

This section explains two important pieces of information to help you through the upgrade process:

- [Firmware upgrade process and time estimates below](#) – which provides an overview of the upgrade process.
- [Upgrade Times above](#) – which provides information on how to monitor the upgrade by viewing messages in various log files that are updated during the upgrade process.

---

**i Note:** In the StorNext GUI, **DO NOT** deactivate the "Upgrade in Progress" state or start up services manually on either server node during the upgrade.

---

**⚠ Caution:** If StorNext services are manually started while the metadata array firmware is being upgraded, the MDC nodes will likely SMITH.

## Firmware upgrade process and time estimates

There are several factors that affect the availability of the system and metadata operations during upgrades. StorNext appliance firmware upgrades (contained in the .fw files downloaded to install the upgrade) always include an update to the StorNext software in the release, and can also include updates to server node firmware, and/or metadata array controller firmware. The way in which these updates are applied, amount of time an upgrade will take and the impacts to system availability vary, depending on what is being upgraded on the system, including the StorNext release you are upgrading from, and the StorNext release you are upgrading to, and which firmware upgrade is being applied, as described in the following section.

## StorNext releases requiring downtime

**i Note:** Metadata array operations and the StorNext file system are unavailable while metadata array controller firmware is installed. We refer to this suspension of operations as a downtime upgrade, since no array I/O may take place during this time. This type of upgrade can substantially increase the total upgrade time for the appliance. See the table below to identify which upgrades include a metadata array firmware upgrade.

StorNext Releases requiring suspension of metadata operations during upgrades occur any time a newer release contains a newer version of metadata array controller firmware version than the version currently installed.

The following table contains examples of firmware upgrade versions. Upgrading StorNext Release with a newer firmware will be a downtime upgrade:

StorNext Release*	Controller Array Firmware Version
5.3.1	08.20.09
5.3.0	08.20.09
5.2.2	08.20.09
5.2.1	08.10.13
5.2.0.2	08.10.13
5.2.0.1	08.10.13
5.2	08.10.13
5.1.1	08.10.13
5.1	07.84.46
5.0.1	07.84.46
5.0	07.84.46
4.7.1	07.84.46
4.7.0.1	07.84.46
4.7.0	07.84.46
4.6.1	07.80.55
4.6	07.80.55
4.3.3	07.80.55
4.3.2	07.80.55

StorNext Release*	Controller Array Firmware Version
4.3.1	07.80.55
4.3.0	07.80.55
4.2.2.0.1	07.75.17

\* Note: Some StorNext releases are not supported upgrade paths to StorNext 5 releases. Please consult the StorNext 5 Compatibility Guides and/or older StorNext Compatibility Guides applicable for the StorNext Release you wish to upgrade to, in order to determine your particular upgrade path.

Example (based on the previous table): If you were upgrading the StorNext firmware from StorNext Release 4.7.1 to StorNext 5 Release 5.1.1, the array firmware would be upgraded from 07.84.46 to 08.10.13. This is a downtime upgrade.

## Component upgrade time estimates

Component	Upgrade Time Estimate (approx. minimum)
Server node	30 to 60 minutes per node*
Metadata Array	20 minutes**

\* Reboot times could vary widely, depending on the size of the SAN in your environment, whether or not server node firmware needs to be upgraded, and the StorNext Release installed prior to the upgrade.

\*\* Since metadata array operations, and the StorNext file system will be suspended during the array firmware upgrade, this is a downtime upgrade.

**i Note:** See sections below for information on additional reboots that may be required for specific StorNext Releases.

## Initiate a Graceful Server Fail-over

Use this procedure to fail-over server operations from a server node operating as primary to a server node operating as secondary without losing client access to metadata. This is helpful if the server has had to fail-over previously for any reason. This does not apply to standalone single node server products or if the server node is not operating as primary.

## How to initiate the fail-over

### Gracefully fail-over the server node operating as primary to the secondary

Access the command line of the node currently operating as primary, and initiate the graceful fail-over to the server node currently acting as secondary:

1. Open an SSH connection to the appropriate MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.
2. At the command prompt, enter **stornext** for the username.
3. Enter the password for the **stornext** user account. The default password is "password", but may have been changed after initial configuration.
4. At the prompt, enter **sudo rootsh** to gain root user access.
5. At the prompt, enter the password for the **stornext** user account again.
6. Press **Enter**.
7. Confirm that the server node you are connected to is operating as the primary by entering the following at the command line:

```
snhamgr -m status
```

8. Verify the output is (bold text used for clarification):

```
:default:primary:default:running:
```

This indicates the node you are connected to is running, and is set as "primary", and that the secondary node is currently "running".

9. On the node operating as the primary, initiate an HA failover to the node operating as the secondary.

```
service cvfs stop
```

Wait until the secondary node becomes the primary. (Time may vary.)

10. To verify the node has failed over, enter the following at the command line:

```
snhamgr -m status
```

11. Verify the output is (bold text used for clarification):

```
:default:stopped:default:primary:
```

This indicates that the node you are connected to has "stopped" as primary, and that the node previously operating as secondary is now operating as "primary". If this is not what is shown, wait and rerun the **snhamgr -m status** command.

You can leave the SSH connection to this server node running if you will access this server again soon.

### Verify the other server node is now operating as primary

To verify the node previously operating as secondary is now operating as primary:

1. Open an SSH connection to the appropriate MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.
2. At the command prompt, enter **stornext** for the username.
3. Enter the password for the **stornext** user account. The default password is "password", but may have been changed after initial configuration.
4. At the prompt, enter **sudo rootsh** to gain root user access.
5. At the prompt, enter the password for the **stornext** user account again.
6. Press **Enter**.
7. Confirm that the server node you are connected to is operating as the primary by entering the following at the command line:

```
snhamgr -m status
```

8. Verify the output is (bold text used for clarification):

```
:default:primary:default:stopped:
```

This indicates the node you are connected to is "primary", and that the node previously operating as primary is still "stopped".

### Restart StorNext services on the server node that is now operating as secondary

1. Open an SSH connection to the appropriate MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.
2. At the command prompt, enter **stornext** for the username.
3. Enter the password for the **stornext** user account. The default password is "password", but may have been changed after initial configuration.
4. At the prompt, enter **sudo rootsh** to gain root user access.
5. At the prompt, enter the password for the **stornext** user account again.
6. Press **Enter**.
7. Confirm that the server node you are connected to is operating as the primary by entering the following at the command line:

```
snhamgr -m status
```

8. Verify the output is (bold text used for clarification):

```
:default:stopped:default:primary:
```

This indicates the node you are connected to is operating as secondary, that its services are "stopped", and that the other node is now operating as "primary".

9. To restart StorNext services on the secondary node, enter the following:

```
service cvfs start
```

10. Confirm that the node is operating as the secondary by entering the following at the command line:

```
snhamgr -m status
```



11. Verify the output is (bold text used for clarification):

```
:default:running:default:primary:
```

### Repeat the fail-over process to set the original server node to operate as primary

If desired, use this procedure to return the original server to again operate as the primary:

1. Open an SSH connection to the appropriate MDC node and use the IP address assigned to the node on the Management or LAN Client network, or the IP assigned to the Service Port, depending on your access method.
2. At the command prompt, enter **stornext** for the username.
3. Enter the password for the **stornext** user account. The default password is "password", but may have been changed after initial configuration.
4. At the prompt, enter **sudo rootsh** to gain root user access.
5. At the prompt, enter the password for the **stornext** user account again.
6. Press **Enter**.
7. Confirm that the server node you are connected to is operating as the primary by entering the following at the command line:

```
snhamgr -m status
```

8. Verify the output is (bold text used for clarification):

```
:default:primary:default:running:
```

This indicates the node you are connected to is running, and is set as "primary", and that the secondary node is currently "running".

9. On the node operating as the primary, initiate an HA failover to the node operating as the secondary.

```
service cvfs stop
```

Wait until the secondary node becomes the primary. (Time may vary.)

10. To verify the node has failed over, enter the following at the command line:

```
snhamgr -m status
```

11. Verify the output is (bold text used for clarification):

```
:default:stopped:default:primary:
```

This indicates that the node you are connected to has "stopped" as primary, and that the node previously operating as secondary is now operating as "primary". If this is not what is shown, wait and rerun the **snhamgr -m status** command until you see the output shown above.

12. Restart StorNext services on this node:

```
service cvfs start
```

13. Verify the output is (bold text used for clarification):

```
:default:running:default:primary:
```

14. Continue to the next section, to complete the procedure.

### **Complete the procedure**

1. Verify that all clients have full access.
2. Test access to all file systems.

### **Complete the procedure**

1. Verify that all clients have full access.
2. Test access to all file systems.

# Contacting Quantum

More information about StorNext is available on the Quantum Service and Support website at <http://www.quantum.com/ServiceandSupport>. The Quantum Service and Support website contains a collection of information, including answers to frequently asked questions (FAQs).

## StorNext Appliance Upgrades

To request a StorNext software upgrade for StorNext Appliances, open a support ticket at: <https://onlineservice.quantum.com/>. For further assistance, or if training is desired, contact the Quantum Technical Assistance Center.

## Contacts

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

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

## Comments

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

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

## Getting More Information or Help

StorageCare™, Quantum's comprehensive service approach, leverages advanced data access and diagnostics technologies with cross-environment, multi-vendor expertise to resolve backup issues faster and at lower cost.

Accelerate service issue resolution with these exclusive Quantum StorageCare services:

- **Service and Support Website** - Register products, license software, browse Quantum Learning courses, check backup software and operating system support, and locate manuals, FAQs, firmware downloads, product updates and more in one convenient location. Get started at:  
<http://www.quantum.com/serviceandsupport/index.aspx>
- **eSupport** - Submit online service requests, update contact information, add attachments, and receive status updates via email. Online Service accounts are free from Quantum. That account can also be used to access Quantum's Knowledge Base, a comprehensive repository of product support information. Get

started at:

<https://onlineservice.quantum.com>

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

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

For worldwide support:

<http://www.quantum.com/serviceandsupport/index.aspx>

## Worldwide End-User Product Warranty

For more information on the Quantum Worldwide End-User Standard Limited Product Warranty:

<http://www.quantum.com/serviceandsupport/warrantyinformation/index.aspx>