



StorNext 4.7.0 M660 Metadata Appliance Release Notes

Contents

Purpose of this Release	2
StorNext M660 Metadata Appliance Requirements	2
Fixed Issues	5
Known Issues	6
Upgrading Firmware	7
Converting to HA	11
Contacting Quantum	14

Product of USA.

Quantum Corporation provides this publication “as is” without warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. Quantum Corporation may revise this publication from time to time without notice.

COPYRIGHT STATEMENT

© 2014 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.

TRADEMARK STATEMENT

Quantum, the Quantum Logo, Backup. Recovery. Archive. It's What We Do., Be Certain, Be Quantum Certain, DLT, the DLT Logo, DLTSage, DLTtape, the DLTtape Logo, DXi, DXi Accent, Dynamic Powerdown, FastSense, FlexLink, GoProtect, GoVault, iLayer, Lattus, MediaShield, Optyon, Pocket-sized., Well-armed., Preserving the World's Most Important Data. Yours., Q-Cloud, Quantum Certain, Quantum Certainty, Quantum vmPRO, Scalar, SDLT, SiteCare, SmartVerify, StorageCare, StorNext, Super DLTtape, SuperLoader, and Vision are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners.

Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners.

Quantum specifications are subject to change.



The StorNext M660 Metadata Appliance combines industry-proven Quantum hardware and StorNext software into one convenient, out-of-the-box solution. Your M660 system has been pre-installed and is ready to operate with a minimum of additional configuration.

This document contains additional information related to your M660 system.

For the complete list of documentation for the M660 Metadata Appliance, click on the following web page:

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

Purpose of this Release

StorNext 4.7.0 is a maintenance release of StorNext software and includes important bug fixes.

The complete list of documentation for StorNext 4.7.0 can be found here (click the "Select a StorNext Version" menu to view the documents for a given release of StorNext):

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

StorNext M660 Metadata Appliance Requirements

This section contains important things you should know about your M660 Metadata Appliance.

M660 Metadata Appliance

The StorNext M660 metadata controller (MDC) nodes are designed to run the MDC file system and, optionally, with the StorNext Storage Manager feature installed. The standby node of the M660 may, optionally, be configured for use as a Distributed Data Mover host. The M660 may be configured as a Gateway server in non-Lattus environments. The M660 may also, optionally, offload processing activities to another StorNext SAN client configured as a Distributed Data Mover host.

Note: Do not enable gateway functionality for M662s connected to Lattus systems. While there are no functional issues in enabling the LAN gateway functionality within a Lattus environment, observed throughput when moving data to Lattus through the 10GbE ports will be slower than expected if this feature is enabled.

Data Replication and Deduplication

The M660 supports the use of an optional StorNext Replication license.

Note: M662s used in the Lattus system need both Eth8 and Eth9 10GbE ports available for access to the Lattus “private” network. If Replication is desired, Eth5 through Eth7 1GbE ports are available.

The M660 is not designed for Deduplication, so the standard StorNext Deduplication license is not supported with the M660.

SNSM Wide Area Storage License

The SNSM (StorNext Storage Manager) Wide Area Storage license is installed and can be enabled on the M662, providing access to Lattus object storage as a storage destination. This access is limited to the M662 configured for use in a Lattus system.

StorNext Software Features Supported

With the exceptions noted above, all other optional StorNext software is supported on the M660, and is purchased separately.

M660 Metadata Appliance File System Restrictions

The StorNext M660 does not support running NFS, or CIFS/SAMBA directly on the Metadata Appliance.

M660 Hardware Expansion

The M660 comes with unfilled expansion slots and drive bays. The M660 offers optional hardware upgrade kits for networking and disk capacity. Hardware upgrade kits require professional installation; other hardware upgrades are not supported.

M660 Memory Requirements

The M660 memory allocation settings that were tested for use with MySQL in StorNext 4.3.x and later are presented in [Table 1](#). Other settings may be reasonable, depending on specific conditions and workloads. However, these settings have been qualified in release testing.

Table 1 Memory Allocation Settings

	FSM BufferCacheSize	FSM InodeCacheSize	MySQL innodb_buffer_pool_size	Reserved for General Usage	Total
M660	8 GB per file system (e.g. 8 x 8 GB = 64 GB)	512 K inodes per file system (e.g. 8 x ~1 GB = 8 GB)	40 GB	32 GB	144 GB

Linux Device Mapper Multipath Support

The StorNext M660 Metadata Appliance supports the standard Linux Dynamic Multipath Mapping driver (DMMP) for the metadata array and additional storage. For most configurations, the upgrade process will automatically

configure the DMMP settings. If a `multipath.conf` file is already configured on the system, a RAS ticket will be generated and the upgrade process will be halted. Please contact Quantum Customer Support for resolution assistance. In addition:

- Ensure that the `/etc/multipath.conf` file is identical on both the primary and secondary nodes of the MDC HA pair.

Configuring Clients for the M660

In order to prevent a split-brain condition between the HA pair of MDCs on the M660, at least one additional StorNext client must mount the HA file system. This will allow the additional client to "vote" in the event of a split-brain condition.

Because the shared file system on the M660 is on the internal RAID and not visible, you must mount the client using the "diskless=yes" option.

On Linux systems, put into the `/etc/fstab` an entry similar to this:

```
shared-02637 /stornext/shared cvfs diskless=yes 0 0
```

(The name "shared-SV1249CKD29435" used in the example will vary. The format is "shared-NNNNNNNNNNNNNN")

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

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

Note: You need to do this on only one client machine.

Target Reset and Fiber Channel Tape Support on Qlogic HBAs

The Enable SCSI Bus Target Reset parameter is enabled by default on all Fiber 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 on the M660, the following commands can be run on each node.

- 1 Connect to each node via ssh and login using the "stornext" user ID.
- 2 Change to root user permissions by running "sudo rootsh"
- 3 Disable Target Resets on the tape SAN port by issuing "`/usr/local/bin/scli -n 1 TR 0`"

4 Confirm that the setting is correct by issuing `"/usr/local/bin/scli -c"` and comparing the output for Port 2. It should look like:

```
[root@Acadia1-1 scripts]# scli -c
-----
-----
HBA Instance 0: QLE2562 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: QLE2562 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         : Disabled
LUNs Per Target             : 128
Enable Out Of Order Frame Assembly: Disabled
```

Fixed Issues

[Table 2](#) lists the fixed issues for all StorNext Metadata Appliances in this StorNext Release.

Table 2 Fixed Issues

Appliance Affected	CR Number	SR Number(s)	Description
All	31644	N/A	The StorNext Online Help has been updated to include information about the Gateways and Firmware in tabs located on the Help > About page.
	33418	N/A	Fixed an issue preventing login to the user interface after a change to the timezone.
	33557	N/A	Applied consistent formatting to online help.
	33740	N/A	Added an <code>/etc/multipath.conf.quantum</code> file as example for various array configurations.
G300	33230	N/A	Updated RAS tickets and alerts to more clearly specify certain hardware failures.
	33855	N/A	Fixed an issue where gateway metric may stop updating.
	28077	N/A	Fixed an issue where time was not updated properly via NTP server pools.
M330 only	33778	N/A	Updated operating system to CentOS 5.9.
M660, M440 and M330	31894	N/A	Fixed an issue where licensing totals and limits were not always correctly displayed in the user interface.
	32999	N/A	Metadata array firmware updated to 10.84.
	33441	N/A	Updated RAS tickets to identify MDC node for power supply failures.

Known Issues

All Metadata Appliances

[Table 3](#) lists the known issues for all StorNext Metadata Appliances for StorNext 4.7.0.

Table 3 Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	SN 38291	n/a	In an HA failover, an Admin Alert is issued if the new primary MDC attempts to initiate an fs_fmover process on the new standby MDC while the standby MDC is still being rebooted.	Once the impacted standby MDC finishes rebooting and becomes functional again, use fsddmconfig (or the GUI) from the master MDC to re-enable DDM for the standby MDC, as follows: <pre># fsddmconfig -u -s e standby_mdc_hostname</pre>
Linux	SN 37538/ SN 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 snfs_config(5) man page or the MAN Pages Reference Guide for details.
Linux	SN 38128	1395540	Using the GUI while a large Media import is kicked off via the command line can cause the GUI to timeout or crash.	Wait until a bulk load from tape is finished prior to opening the StorNext GUI.
Linux	29098/ SN 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.

Upgrading Firmware

The Firmware Upgrade option allows you to perform a firmware upgrade on StorNext M660 Metadata Appliances. Upgrading the firmware also upgrades StorNext, if applicable.

Note: Use the StorNext M660 Metadata Appliance GUI to perform all firmware upgrades and HA conversions.

Upgrade Considerations

Before you begin the upgrade you should note the following considerations so you can plan accordingly:

Consider the following prior to upgrading the M660 Metadata Appliance:

- 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 on the StorNext M660.

For information about supported upgrade paths for StorNext, consult the *StorNext Compatibility Guide*. If your system is running a StorNext release prior to the supported upgrade releases for a given StorNext release, consult an earlier version of the *StorNext Compatibility Guide* that applies to your specific upgrade, and the dependencies for StorNext Clients in the environment.

- Firmware upgrade installation files must first be acquired from Quantum. To obtain the firmware files (both are required) you wish to install for the Metadata Appliance:

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

- a Go to the CSWeb site and log in.
 - b Navigate to the StorNext Products page for your appliance (on the lefthand side of the CSWeb site, look for the appropriate link under the StorNext Products section).

The first section of the page contains Downloads for the given appliance.
 - c Scroll down to the Current Software section, and download both firmware image files.
- Firmware upgrade installation files, which contain the .fw suffix, must be uploaded to the system prior to beginning the upgrade process. 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 Metadata Appliance.
 - When using the firmware upgrade process from the StorNext GUI, the license for the system will be automatically applied to the StorNext M660.
 - Each StorNext firmware upgrade requires a reboot of both nodes. In some cases, multiple reboots of the nodes are required. The StorNext GUI is not available during reboots which can each take 30 minutes or longer per node. The GUI for the secondary node should not be started during this process.
 - Each time the firmware upgrade is done, the secondary node of the Metadata Appliance is left out of the HA configuration. As a result, you will need to convert the StorNext M660 to an HA system after each upgrade in order to regain failover operations. After the HA conversion, both nodes will reboot, which can take 30 minutes or longer per node.

- The Storage Manager components will need to be restarted after the HA conversion is complete by clicking the **Start** button in the Storage Manager panel of the **Tools > System Control** page.


Upgrade Procedure

To upgrade the StorNext M660:

- 1 Download the required firmware file(s) from Quantum CSWeb for the StorNext release needed.
- 2 Log into the StorNext GUI.
- 3 Choose **Tools > Firmware Upgrade**.

The **Firmware Upgrade** page appears.

Figure 1 Firmware Upgrade Page



- 4 Do one of the following:
 - a Select **Auto Upload** to upload the file immediately after you select it.
 - b Do not select **Auto Upload**.
- 5 Click **Browse...**, and then navigate to the directory where the file resides. Firmware files are identifiable by the **.fw** extension.

Note: There are two **.fw** files required for updating firmware. The filenames are similar to QTM-DXiSNA-upd-4.7.0.OS6-10313-0.1of2.fw and QTM-DXiSNA-upd-4.7.0.OS6-10313-0.2of2.fw. Since it is a two-part upgrade, upload both files to the GUI. To begin activation of the upgrade, select either of the uploaded files and then click **Apply**. Both parts are applied to the system.

If you selected the **Auto Upload** option, the file is immediately uploaded. Proceed to [Step 7](#).

- 6 If you did not select the **Auto Upload** option and want to validate the file before uploading, click **Validate**. After a message informs you that the file is valid, click the **Upload** button located to the right of the **Browse...** button.

Note: Files are automatically validated after you click **Apply** ([Step 7](#)), but you won't receive a message telling you the file is valid.

7 Click **Apply** to begin the upgrade.

The green status indicator at the top of the page indicates upload progress, not the upgrade progress. To monitor upgrade progress, check the logs available under the Reports menu.

Note: After the upgrade to the primary MDC node completes, metadata operations will be interrupted for 30 minutes or more, and both MDC nodes will reboot, which could take an additional 30 minutes to complete. It could be a long time before you are able to log back in, so plan upgrade times accordingly.

8 Convert the system to HA, according to [Converting to HA](#) on page 11. (Not necessary when upgrading to StorNext 4.7.0).

GUI Feedback During Upgrades

There are some indications within the GUI that the system is being upgraded. Here are some notes about this visual feedback:

- On the primary MDC node, the GUI will display different status messages throughout the installation, including messages that the system will reboot, and red icons indicating that the primary MDC node, secondary MDC node, File System, and Storage Manager are also disabled.
- The current user account will eventually time out and the GUI will stop functioning when the power to Metadata Appliance is removed during reboot. Status updates will cease and the GUI will not be fully-functional again until the system completes rebooting both MDC nodes.

Post-Upgrade Failover

If you desire to failover your system after the upgrade, see [Initiating a Graceful System Failover](#) on page 12.

Deleting Uploaded Files

Follow these steps to delete uploaded files you no longer need:

- 1 Log into the StorNext GUI.
- 2 Choose **Tools > Firmware Upgrade**.
- 3 The **Firmware Upgrade** page appears
- 4 Select from the list the file you want to delete, and then click **Delete**. (If you want to delete multiple files, you must delete them one at a time.)
- 5 When a confirmation appears, click **Yes** to proceed or **No** to abort.

When a message informs you that the file was successfully deleted, click **OK**.

Converting to HA

This section describes the configuration steps necessary to convert two StorNext MDC nodes into a High Availability MDC pair connected to a shared file system. Converting to HA consists of selecting the dedicated unmanaged StorNext file system for use as the controlling shared file system, and then instructing StorNext to convert each MDC node to operate as one MDC node of the HA pair.

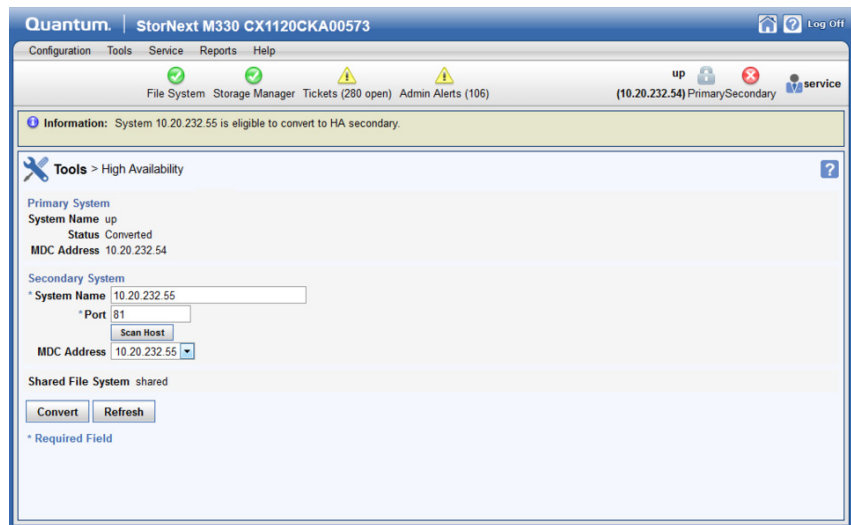
- The UIDs for the **quantumdb** and **tdlm** users and the **adic** group **must** be identical on both MDC nodes of an HA pair. If the UIDs are not identical, the MySQL database will not start (due to file permission errors), which in turn prevents storage manager from fully starting up.

HA Conversion Procedure

Follow these steps to configure HA:

- 1 Choose **Tools > High Availability > Convert**.

Figure 2 Tools > HA > Convert



Note: The primary MDC node **Status** displays “Converted”.

- 2 The IP address of the secondary MDC node of the StorNext M660 will appear in the System Name field along with a port number. Click **Scan Host**. The system should resolve the secondary MDC node - the MDC Address will auto-fill with a value.
- 3 Click **Convert** to convert the secondary MDC node.

Note: Both MDC nodes will reboot, which can take 30 minutes or more per node to complete.

- 4 Storage Manager may need to be started following the HA conversion if the system was in config mode at the time that HA conversion was initiated. To

restart the Storage Manager components, click the **Start** button in the Storage Manager panel of the **Tools > System Control** page.

GUI Feedback During HA Conversion

There are some indications within the GUI that the system is being upgraded. Here are some notes about this visual feedback:

- After the StorNext upgrade has completed, and the HA configuration has been done, the GUI for the secondary MDC node provides a message stating it is not the primary MDC node and a link to launch the primary MDC node.

Caution: Do not login to the GUI of the secondary MDC node at any point during the upgrade/HA conversion process. System configuration and licensing for the system could be compromised.

- When you are able to log into the primary /node, after accepting the EULA, the system will automatically display the **Tools > System Control** page. Click the **Start** button to restart the Storage Manager components.
- Wait until the system icons for both MDC nodes of the system as well as File System and Storage Manager are green, which indicates normal operation.

Post-Conversion Steps

If you are using the DDM feature, do the following:

- If you use the secondary MDC node as a DDM mover, make sure the file systems are mounted.

Initiating a Graceful System Failover

To initiate the failover of a Metadata Appliance after converting to HA or any time a failover is desired:

- 1 Open an SSH connection to the MDC node operating as the **primary**.
- 2 Login to the command line of the **primary** MDC node.
 - User: **stornext**
 - Password: (the customer should have the password for the **stornext** account)
- 3 Type **sudo rootsh** to gain root user access.
- 4 Enter the password a second time.
- 5 Confirm that the MDC node is operating as the **primary** by entering the following at the command line:

```
snhamgr -m status
```

- 6 Verify the output is (bold used for clarification):

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

- 7 On the MDC node operating as the **primary**, initiate an HA failover to the MDC node operating as the **secondary**.

```
service cvfs stop
```

- 8 Wait until the **secondary** MDC node becomes the **primary**, and leave your SSH connection to this node open. (Time may vary.)
- 9 Open an SSH connection to the MDC node now operating as the **primary**.
- 10 Login to the command line of the **primary** MDC node.
 - User: **stornext**
 - Password: (the customer should have the password for the **stornext** account)
- 11 Type **sudo rootsh** to gain root user access.
- 12 Enter the password a second time.
- 13 Confirm that the MDC node is operating as the **primary** by entering the following at the command line:

```
snhamgr -m status
```

- 14 Verify the output is:

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

- 15 From the SSH connection to the MDC node now operating as the **secondary**, enter the following:

```
service cvfs start
```

- 16 Confirm that the MDC node is operating as the **secondary** by entering the following at the command line:

```
snhamgr -m status
```

- 17 Verify the output is:

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

- 18 Repeat if desired to fail over to the original system operating as the **primary**.
- 19 Verify that all clients have full access.

20 Test access to all file systems.

Contacting Quantum

More information about this product is available on the Quantum Service and Support website at www.quantum.com/ServiceandSupport. The Quantum Service and Support website contains a collection of information, including answers to frequently asked questions (FAQs). You can also access software, firmware, and drivers through this site.

For further assistance, or if training is desired, contact Quantum Technical Assistance Center.

North America	1 800-284-5101 (toll free) 1-720-249-5700
EMEA	+00800 7826 8888 (toll free) +49 6131 3241 1164
APAC	1-800-7826-8887 (toll free) +603-7953-3010
Service and Support Web Site	www.quantum.com/ServiceandSupport
Online Service Request System	www.quantum.com/OSR

(Local numbers for specific countries are listed on the Quantum Service and Support Website.) [Converting to HA](#) on page 11