



StorNext M660 Metadata Appliance

4.3.3 Release Notes

Purpose of this Release

The StorNext M660 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.

For upgrading to StorNext 4.3.3 on the M660 Metadata Appliance, refer to the “Update Firmware” section of Chapter 7 of the current version of the StorNext User’s Guide.

This document contains additional information related to your M660 system.

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

<http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/M330/Index..aspx?whattab=Third>

StorNext 4.3.3 is a maintenance release of StorNext software and includes important bug fixes.

The complete list of documentation for StorNext 4.3.3 can be found here (click the “Select a StorNext Version” menu to view the documents for that version of StorNext):

<http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx?whattab=Fourth>



About This Release

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 and the Distributed Data Mover features installed. The M660 may be configured as a Gateway server.

Data Replication and Deduplication

The M660 supports the use of an optional StorNext replication license. The M660 is not designed for deduplication, and the standard StorNext deduplication license is not supported with the M660.

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

For the StorNext M660 Metadata Appliance - StorNext 4.3.3 and later 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. (See CR [40015](#) in [Table 2](#) on page 6.)

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 1](#) 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
G300	32941	n/a	G30x System Serial Number was not shown on RAS ticket on the MDC. The G30x System Serial Number is now listed right after the name of the G300 system in the RAS ticket on the MDC.
M660	33098	1524898	HA node failover resulted in ping pong SMITH between the two nodes for M660 systems. On reboot now, ping pong does not occur any more.
All Appliances	33099	n/a	Upgrades from pre-4.3.2 were missing raid-strings and deviceparams files in <code>/usr/cvfs/config</code> . For Appliances upgraded from 4.3.2 to 4.3.3, files are now located in <code>/usr/cvfs/config</code> .
M440	33129	1529052	The M440 Service Menu had 3 options to “ Reconfigure MD LUNs ” to 2, 4 or 8 LUNs. These were M330-specific menu options. The extra menu options have been removed from the M440 Service Menu .
All Appliances	33137	n/a	Added Dot-Hill array entries to the <code>multipath.conf</code> file defaults for CentOS 5 and CentOS 6
G300	33230	n/a	RAS tickets and email attachments contained generic information on the MDC GUI when events were generated from the G300. For G300s upgraded from 4.3.2 to 4.3.3, the G300-specific information is now passed to the MDC.
M660	33309	n/a	M662 showing one of the network port (eth3) as failure when no bonding is configured. The GUI no longer reports that eth3 has failed.

Known Issues

Metadata Appliances

[Table 2](#) lists the known issues for all StorNext Metadata Appliances:

Table 3 Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	40015	n/a	The M660 <code>/etc/multipath.conf</code> file needs to be identical on both the primary and secondary nodes of the HA pair for failover to work properly.	Copy the <code>/etc/multipath.conf</code> file into the same directory on both MDC nodes of the M660 system any time changes are made to the file.
All	38419	n/a	In the GUI, the reset button on the Configuration > System -> Network tab doesn't always restore current network configuration values.	When this happens, navigate away from the page by selecting any other menu item and navigate back again to reload the current network settings.
All	38291	n/a	In an HA failover, an Admin Alert is issued if the new primary MDC attempts to initiate an <code>fs_fmover</code> 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 <code>fsddmconfig</code> (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>
All	29483/ 38267	n/a	Logical and physical Fibre Channel port numbers may not match.	There is no current workaround for this issue. This will be fixed in a future StorNext release.
Linux	37538	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 <code>snfs_config(5)</code> man page or the MAN Pages Reference Guide for details.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	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/ 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.

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.)