

QXS Gx222xxxx Release Notes

Product	<p>The Gx222R050 firmware includes two versions:</p> <ul style="list-style-type: none"> • GL222R050 for: <ul style="list-style-type: none"> • QXS-312 and QXS-324 • QXS-412, QXS-424, QXS-448, and QXS-456 • StorNext QXS-1200, QXS-2400, and QXS-5600 • DH-3534, DH-3524, DH-3834, and DH-3824 • DH-4534, DH-4524, DH-4544, DH-4554, DH-4834, DH-4824, DH-4844, and DH-4854 • GF222R050 for: <ul style="list-style-type: none"> • QXS-648 and QXS-656 • DH-6544, DH-6554, DH-6844, and DH-6854
<p>Note: DH models require the current firmware to be at Dot Hill branded GL210 before proceeding with an upgrade to Quantum's Gx222R050.</p>	
Operating Systems	<p>Microsoft Windows Server 2008 Microsoft Windows Server 2008 R2 Microsoft Windows Server 2008 Hyper-V Microsoft Windows Server 2012 Microsoft Windows Server 2012 R2 Microsoft Windows Server 2012 Hyper-V Red Hat Enterprise Linux 5, 6, and 7 SuSE Linux Enterprise Server 11 and 12 VMware ESXi 5.0/5.1/5.5 Apple Mac OS X Mavericks 10.9.0 and 10.9.2 (Fibre Channel only)</p>
Date	August 2017

© 2017 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. Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, DXiV-Series, FlexTier, Lattus, the Q logo, The Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the SuperLoader, Symform, the Symform logo (and design), 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. 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.



Contents

GL222xxxx Binary Wrapper	2
Description	5
Recent Enhancements	6
Recent Fixes.	8
Supported/Unsupported Configurations	13
Firmware Notes.	14
Installation Instructions.	16
Known Issues	21
Calculating Maximum IQN Name Length	27
Contacting Quantum	29

GL222xxxx Binary Wrapper

Note: The GL222xxxx binary wrapper is only used and/or seen when installing the GL222xxxx firmware. Customers will see the GL222xxxx binary wrapper when installing the GL222xxxx firmware.

An update to GL222xxxx firmware must **not** be performed if the READ-AHEAD-CACHE settings are other than the default (Adaptive) setting.

Caution: A disruption in service is possible with the unit if the READ-AHEAD-CACHE setting is not changed back to the default (Adaptive) setting before the firmware update is started.

The setting can be restored once the firmware update is complete.

The G222 binaries have been wrapped and a warning message displays to help convey the importance of making this setting change prior to the firmware update.

This section provides the following:

- [Binary Wrapper Contents](#)
- [Security Warning Pop-Up](#)
- [Warning Message](#)

Binary Wrapper Contents

The binary wrapper contains the firmware binaries and requires users to acknowledge the warning message before firmware binaries are available.

- The wrapper executable runs on Windows only, outside of the firmware.

Note: The binary wrapper is only needed for GL222xxxx firmware.

The GF222xxxx firmware is **not** to be wrapped; it is not affected.

- The binary wrapper includes the followings files:
 - GL222xxxx-QTM.bin
 - GL222xxxx-QTM.md5
 - README-FIRST.txt
 - README-GL222xxxx.txt
- Once binaries are unwrapped, user proceeds with the normal firmware update process.
- Refer to [Firmware Notes](#) on page 14 for the firmware update process.

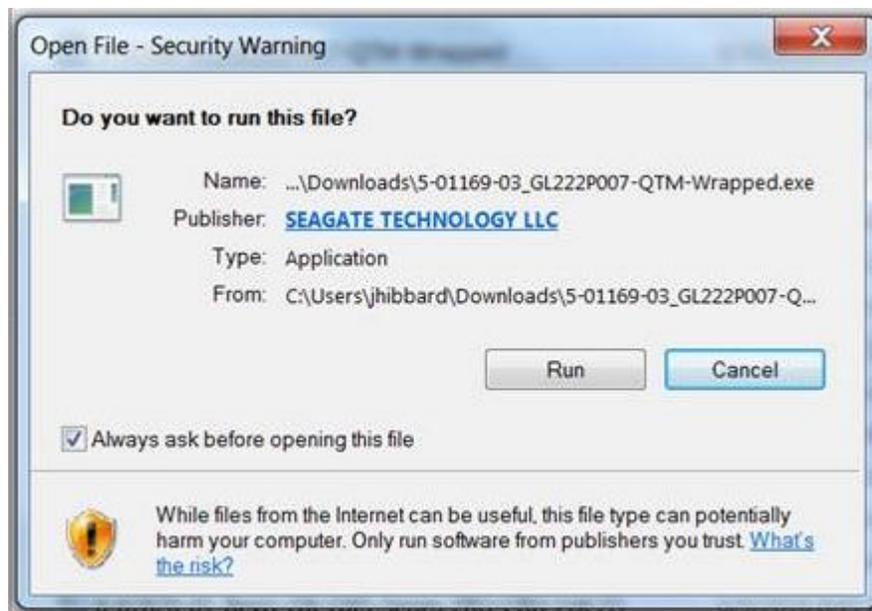
Security Warning Pop-Up

Note: Depending on security settings and version of Windows, a security warning pop-up might occur before the GL222xxxx-QTM-Wrapped.exe file is executed.

Click the "Run" button to execute the GL222xxxx-QTM-Wrapped.exe file.

Refer to [Figure 2](#) for a representative example of the security warning pop-up.

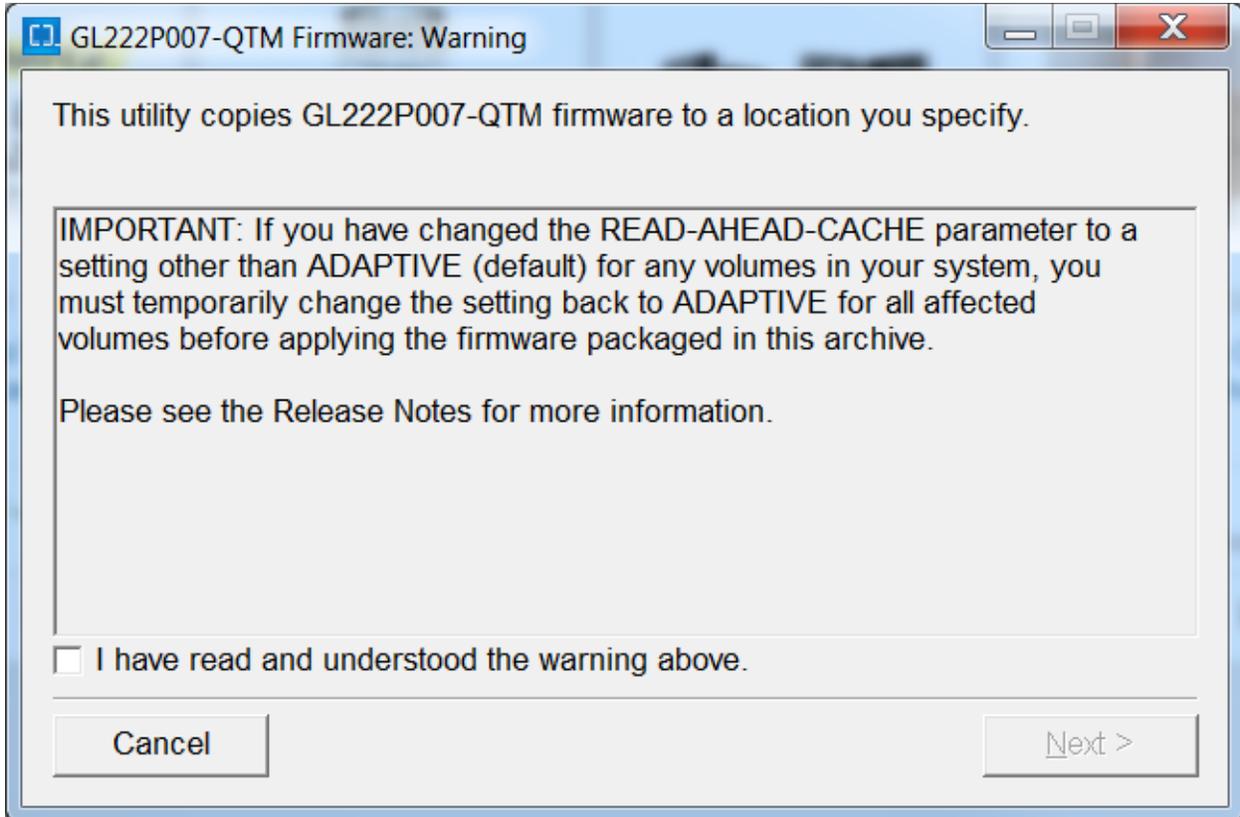
Figure 1 Security Warning
Pop-Up



Warning Message

Refer to [Figure 2](#) for a representative example of the GL222xxxx-QTM Firmware: Warning message.

Figure 2 Warning Message



Note: Quantum highly recommends upgrading all systems running earlier firmware versions to current firmware Gx222xxxx version. See [Firmware Notes](#) on page 14.

Note: A different version of Gx222xxxx is available for Dot Hill branded systems.

Before upgrading to Gx222xxxx, READ-AHEAD-CACHE must be set to "Adaptive". See [Known Issues](#) on page 21.

Description

The Gx222R050 firmware release includes two versions: GL222R050 and GF222R050. Firmware versions and systems covered are as follows:

- GL222R050 for:
 - QXS-312 and QXS-324
 - QXS-412, QXS-424, QXS-448, and QXS-456
 - StorNext QXS-1200, QXS-2400, and QXS-5600
 - DH-3534, DH-3524, DH-3834, and DH-3824
 - DH-4534, DH-4524, DH-4544, DH-4554, DH-4834, DG-4824, DH-4844, and DH-4854
- GF222R050 for:
 - QXS-648 and QXS-656
 - DH-6544, DH-6554, DH-6844, and DH-6854

Releases

This section provides the following release information:

- [Gx222R050 Firmware Release](#)
- [Gx222P013 Firmware Release](#)
- [Gx222P007 Firmware Release](#)

Gx222R050 Firmware Release

The Gx222R050 firmware release is a maintenance release for the Quantum Hybrid QXS-3/4/6 Series disk solutions. A variety of bug fixes are included in this firmware release.

Gx222P013 Firmware Release

The Gx222P013 firmware release is a maintenance release for the Quantum Hybrid QXS-3/4/6 Series disk solutions. A variety of bug fixes are included in this firmware release.

Gx222P007 Firmware Release

The Gx222P007 firmware release was the foundation release for the Quantum Hybrid QXS-3/4/6 Series disk solutions. A variety of bug fixes were included as well as the new features that allows for flexible choices on how to use the QXS hybrid systems.

Features

This section provides the following features information:

- [Gx222R050 Firmware Release Features](#)
- [Gx222P013 Firmware Release Features](#)
- [Gx222P007 Firmware Release Features](#)

Gx222R050 Firmware Release Features

This firmware release has no new features.

Gx222P013 Firmware Release Features

This firmware release has no new features.

Gx222P007 Firmware Release Features

This firmware release includes the following features:

- Full support of the virtualization features
- One-button-configuration profiles filtered depending on the max drive count of the RAID
 - Ability to replace the one-button-configuration profile set at time of install with a different set based on vertical applications
 - One-button-configuration warn user if global spare requirements not met
- Ability for one-button-configuration profiles to NOT write a StorNext label
- Improved system alerts infrastructure reports serial number in alerts through RAS-like infrastructure
- Alerts system pre-configured to send alerts to Quantum
- V3 GUI is the default GUI interface for all new installs and updates

Recent Enhancements

This section contains the following recent enhancements:

- [Gx222R050 Enhancements](#)
- [Gx222P013 Enhancements](#)
- [Gx222P007 Enhancements](#)

Gx222R050 Enhancements

This firmware release has no new enhancements.

Gx222P013 Enhancements

This firmware release has no new enhancements.

Gx222P007 Enhancements

The enhancements for Gx222P007 include:

- Added the following virtual storage features:
 - Asynchronous Replication
 - Volume Tier affinity
 - Enhanced snap pool management
 - Enhanced volume copy
- Added scheduler for virtual storage based snapshots.
- Added persistent SNMP alerts for critical events until the condition is cleared.
- Added capability for larger virtual pools.
- Server certificate is signed using SHA-256 algorithm.
- Updated OpenSSH to 7.1 to mitigate CVE-2015-5600, CVE-2015-5352, CVE-2014-2653, CVE-2014-2532, CVE-2014-1692, and CVE-2013-4548.
- Added additional information to FRU events.
- Added empty page reclamation information.
- Added hover tip to QXS-312 and QXS-324 with Fan Out configuration to include Recommended Actions.
- Added new poolHealthReason to enhance the understanding of virtual pool health.
- Added the ability to de-quarantine linear disk-group from QTCR state when drives were removed/inserted during and expand operation.
- Deprecate non-optimal Disk Group chunk sizes.
- Disable RC4 cipher suite to mitigate CVE-2015-2808.
- Disabled 96-bit and MD5-based HMAC and CBC mode ciphers.
- Disabled SSL compression to mitigate CVE-2012-4929 and CVE-2012-4930.
- Disabled SSLv2.
- Disabled Telnet protocol by default. Disabled weak SSL ciphers.
- Enhanced accuracy of size shown in disk management utility (V3 GUI) for drives larger than 1 TB.
- Enhanced managed logs emails.
- Enhanced the severity for Event 571 in the case of allocated pages crossing the high threshold.
- Enhanced the wording for Event 573.
- Improved CLI error messaging.
- Improved SMI-S indication for events.

- Improved SMI-S indication support for FRUs.

Recent Fixes

This section contains the following recent fixes:

- [Gx222R050 Fixes](#)
- [Gx222P013 Fixes](#)
- [Gx222P007 Fixes](#)

Gx222R050 Fixes

The recent fixes for Gx222R050 include:

Note: This list is the key customer-reported fix list. Other fixes, not described, are included in this release.

Issue	Root Cause	Work-around
LUNs are inaccessible following a scheduled system reboot.	Invalid logic in the shutdown routine	Manual intervention required by support personnel to get LUNs online
LUNS inaccessible	Logic error in bad block recovery (BBR) routine	Power cycle the array (temporary recovery)
Customer had over commit enabled on the array. <ul style="list-style-type: none"> • Alerts were configured for pool space thresholds; however, the customer wasn't monitoring the emails. • The application experienced failed writes when it attempted to write beyond the available pool space. 	Lack of a persistent health state in the user interfaces when approaching capacity limit	Add additional disk groups or delete unnecessary/extraneous files to reclaim space.
Pool goes offline following a system reboot.	Logic error	Manual intervention required by support personnel to get LUNs online

Gx222P013 Fixes

The recent fixes for Gx222P013 include:

- Firmware upgrade/downgrade fail between GL210 and GL222 when volume read-ahead-size parameter changed from default

- Data corruption observed during repetitive fail-unfail with replications running.
- Manage Logs e-mail message is mis-formed
- Scheduler: Unable to delete the snapshot
- Logs from managed-log feature are incomplete - ONLY contain the SC Boot Log and nothing else
- SES broadcast from third-party SAN causing disk channel events, aborts and eventually leads to controller crash
- Add support for SCSI Send Diagnostic command when in-band SES is disabled (so command is not rejected for the default LUN).
- VMware: SC crash during XCOPY or WRITESAME operations, OSMEnterDebugger(), File: fxc/xor_util.cpp, Line: 434 (VMware certification fails on GF, but not GL)
- DMS - CAPI HANG and host IO failures
- A corner-case corruption detected by precommit_linear test on Rackbert
- Communication Down fault type remains on a controller, while running init_tests of continuous reboot
- Dual controller crash in AFA with multi host I/O with high queue depth setting on HBAs - OSMEnterDebugger(), File: drivers/dskclass.cpp, Line: 439 Exceeded MAX_SYNCIO_TIMEOUT - AFA-
- While deleting 100's of snaphosts started a format on two A pool volumes from windows disk management, then A controller crashed.
- Occasionally a Controller Killed due to Heartbeat loss (Fault Type: NMI CThr: PSP192)
- After an upgrade to GL22P008 a controller crashed
- Occasionally after a controller reboots the system reports a degraded health state for a period of minutes because a disk is using the alternate path
- While running the repetitive failover/failback test script Cntrlr-B killed Cntrlr-A due to Heartbeat lost and then subsequently Unkilled it automatically resulting in a data corruption
- Controller killed for Recovery fail-over stall
- Controller Page Fault during FO/FB testing
- AR:A replication was suspended internally by the system. The operation cannot be completed because communication cannot be established with the peer.
- Osiris - South Africa - large arrays with 66 SSDs (along with HDDs in hybrid config) hitting what appears to be queue depth issues
- MC crashes when deleting a schedule.
- AFA: Backend SSD Queue Depth set to 64 instead of 125.

Gx222P007 Fixes

The recent fixes for Gx222P007 include:

- Controller crashes with page fault and snap pool deleted after power cycle.
- Prevent possible incorrectly shared pages caused by controller failure during reset/rollback and controller I/O.
- Decrease the PS unmap size to 1 page.
- Corruption with simultaneous new writes and controller crash.
- Fix page faults seen when snapshots are used and virtual LBA space exceeds ~157TB.
- Don't update the flash copy of the enclosure table if another thread owns the flash or config semaphores.
- A bug where ATS/XCOPY delays failover/failback.
- Allow child MLIOs through to local cache when failback is active.
- An issue with MLIO IOs delaying failback.
- The Install License menu item, and other features such as Create Snapshot, might not appear in the disk management utility interface.
- In the disk management utility (V2 GUI) interface, the turn on LEDs option is disabled in the Vdisk Overview Disks panel.
- After the firmware is updated on the array, the disk management utility interface might freeze or become nonresponsive.
- In locales other than English, the Date, Time panel might not reflect the correct NTP settings.
- In Internet Explorer, when the online help button is clicked in the disk management utility interface, it can take a couple of minutes for the online help window to be populated with content.
- The Create Volume/Volume-set page highlights Access field in red (only once per volume).
- In disk management utility interface popups, buttons aren't enabled until clicking outside of the data entry fields or selecting Tab.
- Partner Firmware Update can fail if there is a health issue in the system. If the code load health checker determines that firmware should not be loaded, no event will be logged.
- The disk management utility (V3 GUI) Create Virtual Volume panel uses Copies terminology instead of Number Of Volumes.
- Related Snapshots shown in the Volumes Pane in the disk management utility (V3 GUI) does not list all descendant snapshots of the volume.
- A message incorrectly states "The specified volume was not found" when creating a task to reset a virtual volume snapshot.
- The CLI displays invalid pool details for virtual disk groups temporarily.
- The filter of Severity is truncated if there are two events in the event list.
- The disk management utility (V3 GUI) is not rendering several elements properly in IE 10 and IE 11.

- Thin provisioning information reported in the disk management utility (V3 GUI) and the CLI may not be in sync.
- The disk management utility interface online help does not show internal IP addresses in use by the array.
- The profile for an initiator is not shown in the disk management utility (V3 GUI).
- Initiator profiles cannot be set in the disk management utility (V3 GUI).
- Unable to select SNMPv3 Privacy Type in locales other than English in the disk management utility (V3 GUI).
- Text strings are truncated in some cases in the disk management utility (V3 GUI) when locales other than English are used.
- In non-English locales the disk management utility (V3 GUI) does not reflect the NTP enabled state when it has been enabled in the CLI.
- Snapshot schedules didn't begin until an MC restart, or schedules that had no associated tasks were deleted.
- The disk management utility interface might incorrectly display volume creation sizes with base 10 instead of base 2.
- In the disk management utility (V2 GUI) interface, saving disk performance statistics might fail.
- Creating and scheduling a replication set command fails if a volume name contains a space.
- The global spares count may be incorrect on the Home page in Korean locale.
- SCVMM R2 may not accurately display available virtual pool capacity.
- In the disk management utility (V3 GUI), SSD spares may not be displayed in the global spares information panel.
- Filter options are missing in the Event log viewer when a single event is listed in the viewer.
- Changes to email settings might show Include Logs as being enabled.
- Event logs show status of USB Drive that is not being used.
- Power Supply status shows OK when Fan is bad.
- Unable to set the partner controller IP address due to communication issue.
- A single controller configuration always indicates the chassis right ear fault LED is on.
- Power Supply SMI-S 314 Indication has WARNING severity.
- Controller crashed with NMI: NOCP Errors.
- An HTTPS web browser might not function correctly without a valid certificate.
- Certificate status is unknown.
- Controller crashed due to Critical Error OSMEnterDebugger CThr: miOL.
- Virtual disk-group drain stuck and will not complete.

- Invalid vdisks allowed for selection when creating a replication set in the disk management utility (V2 GUI) interface.
- The Create snapshot command fails immediately after the execution of delete snap-shot with cleanup.
- Disk failure can cause data unavailability until a power cycle.
- Data is not draining from CRITICAL performance tier.
- Controller crash due to resources being exhausted.
- Host Port information popup does not show entire iSCSI initiator IQN.
- Virtual pool is read-only due to multiple failures and rare race condition.
- Rare management controller crash.
- MTU packet size not correctly documented.
- Controller crash due to PCI Link Recovery Stuck in a loop.
- The array was not working correctly in single controller mode.
- IE 11 browser can cause replication to fail when changing the scheduled replication.
- Cannot highlight or get hover tip for RAID-50 in the disk management utility (V3 GUI).
- Text displayed for Hosts information is not clear in the disk management utility (V3 GUI) Home page.
- User session timeout is not displaying in the disk management utility (V3 GUI).
- Event 8 Unknown Reason Drive failure does not create 314 event.
- Not sending an event and SNMP trap when available size is 0 MB.
- Drives that have failed previously due to SMART event should not be allowed to be added to a Vdisk or Disk Group.
- The disk management utility (V2 GUI) interface volume cache settings still contains Write Optimization.
- Creation of a performance tier disk group without a performance license generates Disk group offline event.
- Hard (non-media) errors reported during disk scrub due to illegal request command to disk.
- CompactFlash product ID does not include part number in notifications.
- Problems when using an HTML tag in the search text box on the Event Log Viewer.
- The Delete Volumes option is disabled for a snap of a snap.
- In the disk management utility interface's Full Disk Encryption Settings panel, the action buttons might not become active immediately after you enter required information on the page.

Supported/Unsupported Configurations

This section provides the supported and unsupported configurations as follows:

- [Hybrid QXS Supported Configurations](#)
- [Hybrid QXS Unsupported Configurations](#)
- [StorNext Supported Configurations](#)
- [StorNext Unsupported Configurations](#)

Hybrid QXS Supported Configurations

[Table 1](#) identifies the supported hybrid QXS supported configurations:

Table 1 Hybrid QXS Supported Configurations

Product	Chassis	Max Chassis	Max Drives
QXS-312	2U12	4	48
QXS-324	2U24	4	96
QXS-412	2U12	8	96
QXS-424	2U24	8	192
QXS-448	2U48	4	192
QXS-456	4U56 (4) & 2U12 (2)	6	248
QXS-648	2U48	4	192
QXS-656	4U56 (4) & 2U12 (2)	8	248

Note 1: Always consider the max drive number when mixing different chassis combinations.

Note 2: It is impossible to identify all possible chassis combinations within the hybrid QXS supported configurations.

Hybrid QXS Unsupported Configurations

Currently, with the Gx222R050 FW release, there are no unsupported configurations.

StorNext Supported Configurations

The supported configurations are:

- QXS-1200/QXS-2400 base supports QXS-1200/QXS-2400 expansions.

- QXS-5600 base supports QXS-5600 expansions.

StorNext Unsupported Configurations

The unsupported configurations are:

- QXS-1200/QXS-2400 base does not support QXS-5600 expansions.
- QXS-5600 base does not support QXS-1200/QXS-2400 expansions.

Note: If you have supported configuration questions, refer to the QX and QXS Setup Guide and/or see [Contacting Quantum](#) on page 29 to contact Services and support.

Firmware Notes

Caution: Reverting to a previous firmware version is not recommended. Notify Quantum support for additional information.

Always update controller firmware to the latest when:

- Installing a new system
- Adding expansion chassis
- Replacing a controller I/O module(s) or expansion I/O module(s)

Note: Updating controller firmware with expansion I/O modules active ensures that the controller firmware and expansion I/O module(s) firmware are at a compatible level.

Before Installing Firmware

Complete the following before installing firmware:

- Create a backup of system data (a full backup is strongly recommended).
 - Schedule an appropriate time to install the firmware:
 - For single domain systems, I/O operations must be halted.
 - For dual domain systems, because the online firmware upgrade is performed while host I/Os are being processed, I/O load can impact the upgrade process.
 - Select a period of low I/O activity to ensure the upgrade completes as quickly as possible and avoid disruptions to hosts and applications due to timeouts.
- Allocate sufficient time for the update:
- It takes approximately 45 minutes for the firmware to load and for the automatic restart to complete on the first controller module.

- When dual modules are installed, the full process time is approximately 90 minutes.
- If cascaded drive enclosures are also being updated, the total process time may be as long as 180 minutes.
- Set the Partner Firmware Update option so that, in dual-controller systems, both controllers are updated.
 - When the Partner Firmware Update option is enabled, after the installation process completes and restarts the first controller, the system automatically installs the firmware and restarts the second controller.
 - If Partner Firmware Update is disabled, after updating software on one controller, you must manually update the second controller.

During Installation Process

Monitor the system display during firmware installation to:

- Determine update status
- See when the update is complete

After Installation is Complete

After the installation process is complete and all systems have automatically restarted:

- Verify system status in the disk management utility, and confirm that the new firmware version is listed as installed.
- Review system event logs.
- Updating array controller firmware may result in new event messages that are not described in earlier versions of documentation.
- For comprehensive event message documentation, see the most current version of the Event Descriptions Reference Guide.

Additional FW Notes

The firmware notes include the following:

- Windows Server 2012 management integration:
 - Quantum recommends that you update the Windows cache by using the cmdlet **Update-StorageProviderCache -DiscoveryLevel Full -Name <storageProviderName>** manually, after attempting any storage provision operations that use the QXS management interface (Disk Storage Management Utility or the Command Line Interface).
 - Quantum recommends that you use the QXS management interface to modify volume mappings, delete volumes, or modify volume names. Manually update Windows cache by using the cmdlet **Update-StorageProviderCache -DiscoveryLevel Full -Name <storageProviderName>**.
- System Center VMM integration:

- Running operations concurrently is supported, except for Windows 2008 R2, up to the limit of four concurrent operations. This includes creating objects (e.g., LUNs, clones, snapshots) and registering objects to hosts or four node clusters.
- Windows Server 2012 management and System Center VMM integration:
 - Quantum recommends that you disable Windows Indication subscription, if SCVMM or Windows Server 2012 manages only QXS systems and not any other arrays. To disable the Indication subscription, modify the registry key
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\Current Version\Storage Management\EnableIndications value from 1 to 0 and restart Windows Standards Based Storage Management Service.
 - If you want to enable the Indication subscription at Windows Server 2012, Quantum recommends that you configure the Indication based on the instructions provided at <http://blogs.technet.com/b/filecab/archive/2013/05/22/using-indications-with-the-windows-standards-based-storage-management-service-smi-s.aspx>.
- The QXS contains an embedded SMI-S provider for use by SMI-S client applications. The embedded provider is designed to support QXS configurations with up to 24 hard drives and up to 250 mapping paths. A mapping path is defined as an QXS volume presented through an QXS target port to a host initiator.
- When using Windows Dynamic Disk (software RAID) on top of a hardware RAID, cautions must be considered. For more information, see the section “Real World: Dynamic versus Basic Disks” on the topic at <http://technet.microsoft.com/en-us/library/dd163558.aspx>.
- Failover and failback times are affected by the number of system volumes; the more volumes there are on the system, the more time is required for failover and failback to complete.

Installation Instructions

Caution: Upgrades from a system running Gx105 firmware to the latest firmware requires two steps.

1. Upgrade from Gx105 to Gx210.
2. Upgrade from Gx210 to latest release.

This section describes how to install the firmware.

WARNING: Do not cycle power or restart devices during a firmware update. If the update is interrupted or there is a power failure, the module could become inoperative. If this occurs, contact technical support. The module may need to be returned to the factory for re-programming.

Caution: Before upgrading firmware, ensure that the system is stable and is not being reconfigured or changed in any way. If changes are in progress, monitor them and wait until they are completed before proceeding with the upgrade.

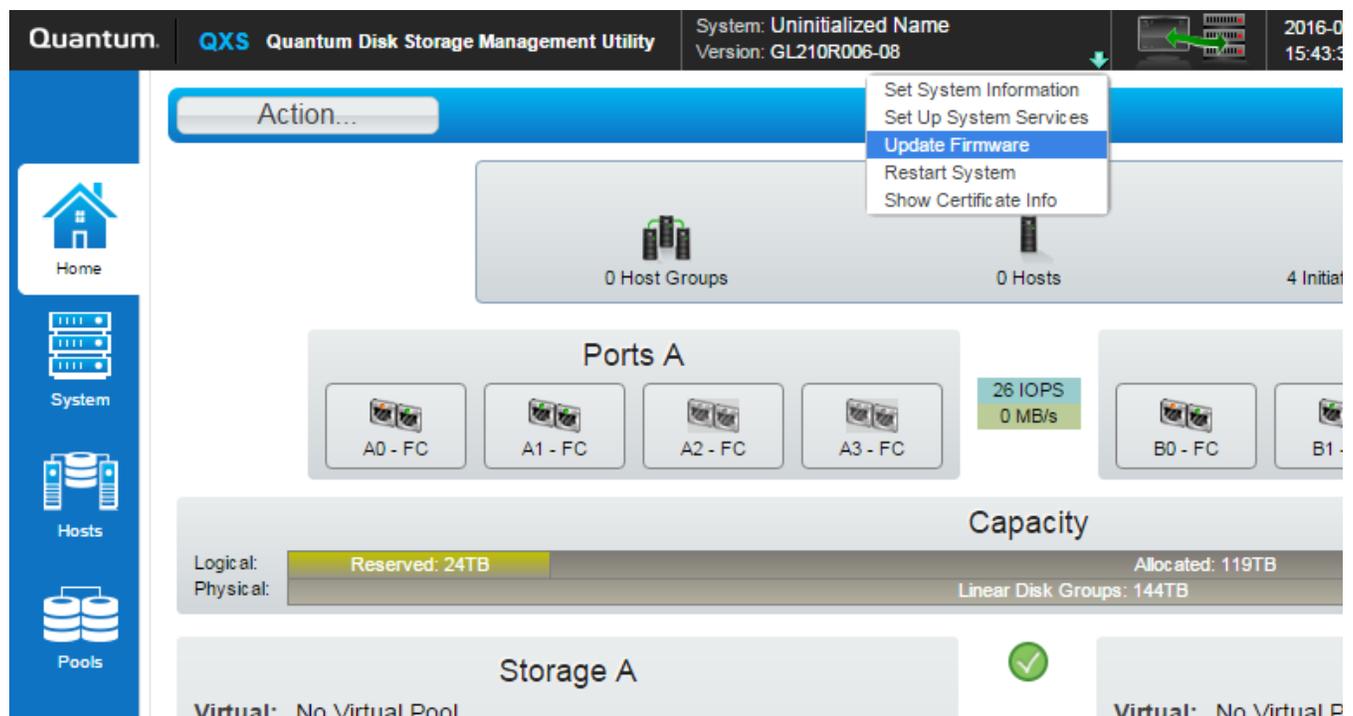
IMPORTANT: In dual-module enclosures, both controllers or both I/O modules must have the same firmware version installed. Running different firmware versions on installed modules may cause unexpected results.

Install Firmware Using Disk Management Utility V3

Follow these steps to install the firmware package using Disk Management Utility:

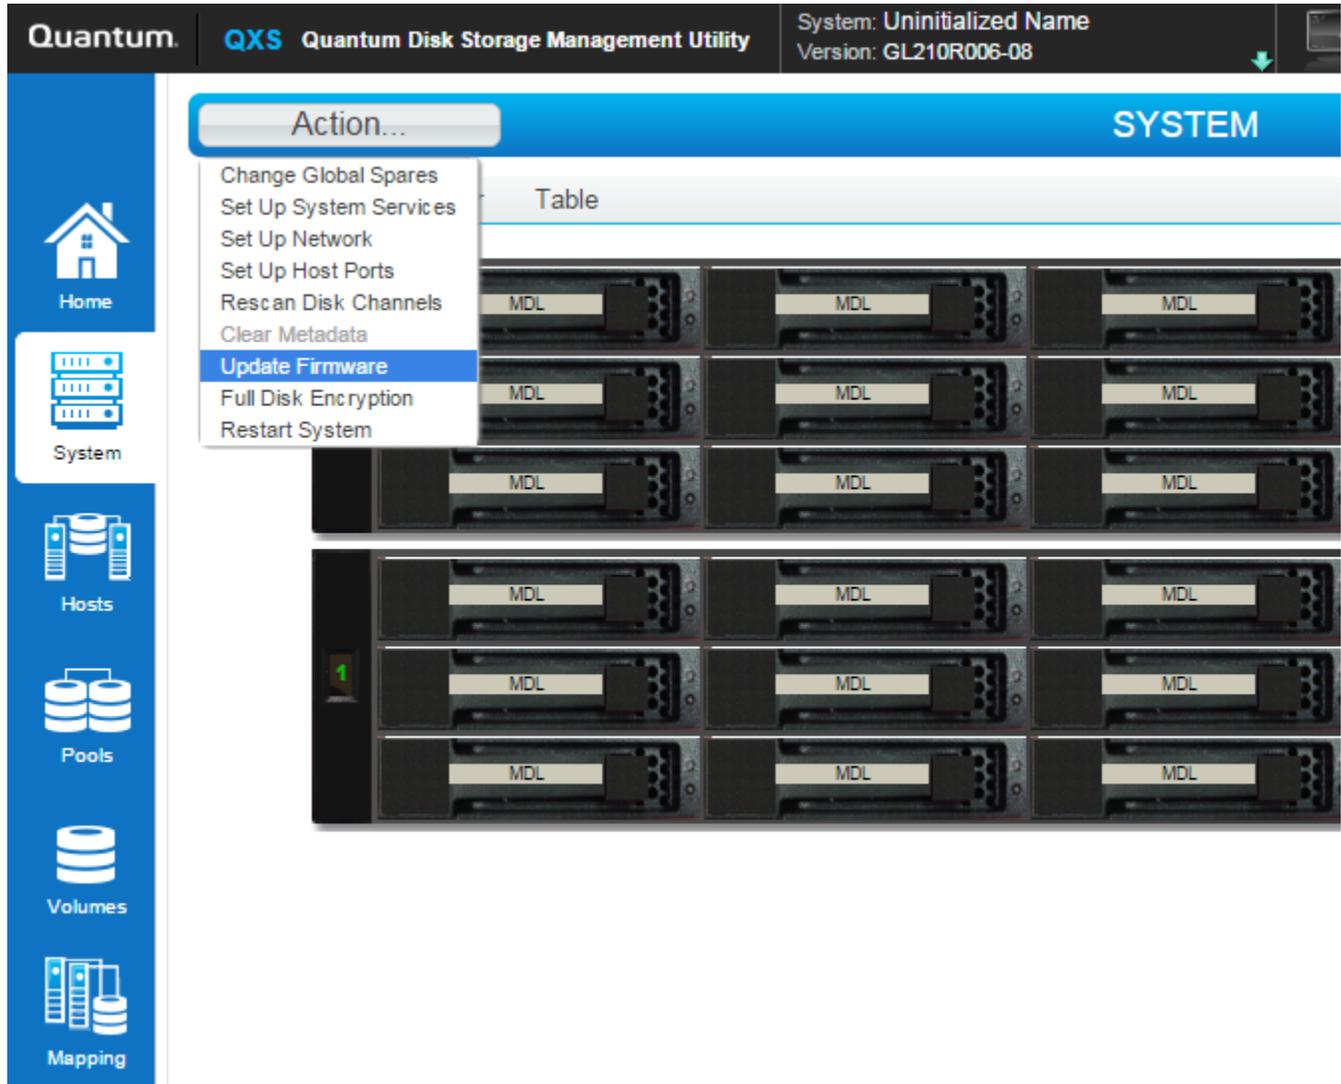
- 1 Obtain the firmware package, and save it to a temporary directory. This firmware release can be obtained by contacting Quantum Technical Support.
- 2 Locate the firmware file in the downloaded/extracted folder. The firmware filename is in the following format: GxxxRyyy-zz.bin.
- 3 In single-domain environments, stop all I/O to vdisks in the enclosure before starting the firmware update.
- 4 Log in to Disk Storage Management Utility.
- 5 Go to the Update Firmware option. There are two ways to do this:
 - a On the top of the home screen (Figure 3), click the first arrow on the upper right where the System name and Version are shown. Click Update Firmware.

Figure 3 Sample Home Screen



- b Click the System tab on the left navigation panel (Figure 4), click the Action button at the top of the system panel, and click Update Firmware.

Figure 4 Sample System Screen



Note: With either of these methods, the panel displayed contains a table that shows the currently installed firmware versions.

- 6 Click **Browse**, and select the firmware file to install.
- 7 Click **OK** to install the Controller-Module Firmware File.
- 8 Wait for the installation to complete. During installation, each updated module automatically restarts.
- 9 In the Disk Storage Management Utility display, verify that the expected firmware version is installed on each module.

Install Firmware Using Disk Management Utility V2

Follow these steps to install the firmware package using Disk Management Utility:

- 1 Obtain the firmware package, and save it to a temporary directory. This firmware release can be obtained by contacting Quantum Technical Support.
- 2 Locate the firmware file in the downloaded/extracted folder. The firmware filename is in the following format: GxxxRyyy-zz.bin.
- 3 In single-domain environments, stop all I/O to vdisks in the enclosure before starting the firmware update.
- 4 Log in to Disk Storage Management Utility and, in the Configuration View panel, right-click the system, and then select **Tools > Update Firmware**.
A table is displayed that shows currently installed firmware versions.
- 5 Click **Browse**, and select the firmware file to install.
- 6 Click **Install Controller-Module Firmware File**.
- 7 Wait for the installation to complete. During installation, each updated module automatically restarts.
- 8 In the Disk Storage Management Utility display, verify that the expected firmware version is installed on each module.

Installing Firmware Using FTP

Follow these steps to install the firmware package using FTP:

- 1 Obtain the firmware package, and save it to a temporary directory. This firmware release can be obtained by contacting Quantum Technical Support.
- 2 Locate the firmware file in the downloaded/extracted folder. The firmware filename is in the following format: GxxxRyyy-zz.bin.
- 3 Using Disk Storage Management Utility, prepare to use the FTP:
 - a Determine the network-port IP addresses of system controllers.
 - b Verify that the system FTP service is enabled.
 - c Verify that the user login has permission to use the FTP interface and has manage access rights.
- 4 In single-domain environments, stop I/O to vdisks in the enclosure before starting the firmware update.
- 5 Open a command prompt (Windows) or a terminal window (UNIX), and navigate to the directory containing the firmware file to load.
 - a Enter a command with the following syntax:
ftp <controller-network-address>. (For example: **ftp 10.1.0.9**)
 - b Log in as an FTP user (user = **ftp**, password = **!ftp**).
 - c Enter a command by using the following syntax:
put <firmware-file> flash
where **<firmware-file>** represents the binary firmware filename

- 6 Wait for the installation to complete. During installation, each updated module automatically restarts.
- 7 If needed, repeat these steps to load the firmware on additional modules.
- 8 Quit the FTP session.
- 9 Verify that the expected firmware version is installed on each module.
 - In Disk Storage Management Utility, right-click the system in the **Configuration View** panel, and then select **Tools > Update Firmware**.
 - In the Command Line Interface (CLI), execute the **show version** or the **show enclosures** command.

Installation Troubleshooting

If you experience issues during the installation process, do the following:

- 1 When viewing system version information In the Disk Storage Management Utility System Overview panel, if significantly more than an hour has elapsed and the components do not show that they were updated to the new firmware version, refresh the browser. If version information is still incorrect, proceed to the next troubleshooting step.
- 2 If version information does not show that the new firmware has been installed, even after refreshing the browser, restart all system controllers. For example, in the CLI, enter the **restart mc both** command. After the controllers have restarted, one of three things happens:
 - Updated system version information is displayed, and the new firmware version shows that it was installed.
 - The Partner Firmware Update process automatically begins and installs the firmware on the second controller. When complete, the versions should be correct.
 - System version information is still incorrect. If system version information is still incorrect, proceed to the next troubleshooting step.
- 3 Verify that all system controllers are operating properly. For example, in the CLI, enter the **show disks** command, and read the display to confirm that the displayed information is correct.
 - If the **show disks** command fails to display the disks correctly, communications within the controller have failed. To re-establish communication, cycle power on the system, and repeat the show disks command. (Do not restart the controllers; cycle power on the controller enclosure.)
 - If the **show disks** command from all controllers is successful, perform the Firmware Update Process again.

Known Issues

This section contains the following known issues and workarounds.

- [Gx222R050 Known Issues](#)
- [Gx222P013 Known Issues](#)
- [Gx222P007 Known Issues](#)

Gx222R050 Known Issues

This firmware release has no known issues.

Gx222P013 Known Issues

This firmware release has no known issues.

Gx222P007 Known Issues

The known issues and workarounds for firmware Gx222P007 are as follows:

Issue	Workaround
<p>During Gx222P007 firmware installation, the code installation activity can halt during the upgrade of the second controller in the update process.</p> <p>Note: This interruption in code installation only occurs when the READ-AHEAD-CACHE settings have been altered from the original default "Adaptive" settings.</p>	<p>The process to avoid this issue is as follows:</p> <ol style="list-style-type: none"> 1. Before upgrading to or from Gx222P007 in any RAID chassis, ensure that the READ-AHEAD-CACHE setting is at the default "Adaptive". 2. Perform a shutdown-restart of the array from the WBI. <p>Note: If this activity has occurred at any point in time after the READ-AHEAD-CACHE setting have been changed from "Adaptive" then there is no need to perform this step again.</p> <ol style="list-style-type: none"> 3. Execute the bundle installation. 4. Once the upgrade is complete the READ-AHEAD-CACHE setting can then be changed to the desired value.
<p>A firmware upgrade from GL105 to GL205 could cause a power supply to lose its FRUID information and display a "PSU FRUID data was not found for this FRU" error message.</p>	<p>Power off the chassis and then power the chassis back on to clear this message.</p>
<p>Sending commands to retrieve the associations of supercap via SMIS may cause an MC crash.</p>	<p>MC will reboot. Retrieve information on the supercap through other interfaces.</p>
<p>In rare cases during a PFU, the EC might not update.</p>	<p>Reloading code on the affected controller will update the EC to the appropriate version.</p>

Issue	Workaround
When swapping controllers between a high-density and low-density chassis, the appropriate EC might not be loaded automatically.	Loading code on the affected controller will update the EC to the appropriate version.
In a low-density chassis, the CLI show expander-status command may incorrectly report the expander is in the “left drawer” although this chassis has no drawers.	This is a display issue only and can be ignored. No workaround is necessary.
A controller crash might occur if the host operating system issues very large unmap commands to the controller.	Reboot the crashed controller. Some operating system settings might be available to limit unmap commands to under 1GB.
In some cases, downgrading to an earlier firmware version might fail to downgrade the MC properly.	Reload the desired firmware bundle.
Unable to create a peer connection after deleting a previous peer connection when using Asynchronous Replication.	Reboot both SCs at the same time using the CLI restart SC both command. Or, from the WBI, choose Restart System from the pull-down menu under the version number at the top. Select Restart with Storage and select both controllers.
In some cases, when moving controllers from a high-density chassis to a low-density chassis the proper EC component might not be loaded.	Loading code on the affected controller will update the EC to the appropriate version.
In a very rare case, after a controller failover, replication schedules might be removed.	Manually recreate the replication schedule.
An aborted replication might show status as “Unsynchronized” rather than “Aborted.”	This is a display issue only and can be ignored. No workaround is necessary.
In the online help for the “Scheduling replications” task, some details are missing, including the permitted number of characters and which characters are invalid.	The limitations on replication schedule names are the same as for replication set names. Refer to the disk management utility guide for limitations on replication names replication set names.
When trying to delete a pool with a replication set defined, this error message might display: Error: Command is not supported. - Failed to delete the virtual pool.	Remove the replication set and then delete the pool.
Modifying the password for an SNMP v3 user might default the authentication type and privacy type to none.	When modifying the password for an SNMP v3 user, also modify the authentication type and privacy type to the desired values.
Email notification for “Event 175: The network-port Ethernet link is up for a controller” might not be sent.	Use the WBI or CLI to check the status of Ethernet links.
If a replication set has entered a suspended state, the events may be inconsistent between the peer systems with respect to the reason for suspension.	Review the events on both peer systems.
CLI help for de-quarantine shows example of virtual disk groups, but virtual disk groups cannot be de-quarantined without support assistance.	None

Issue	Workaround
When swapping controllers between a high-density and low-density chassis, the appropriate EC might not be loaded automatically.	Loading code on the affected controller will update the EC to the appropriate version.
In a low-density chassis, the CLI show expander-status command may incorrectly report the expander is in the “left drawer” although this chassis has no drawers.	This is a display issue only and can be ignored. No workaround is necessary.
A controller crash might occur if the host operating system issues very large unmap commands to the controller.	Reboot the crashed controller. Some operating system settings might be available to limit unmap commands to under 1GB.
In some cases, downgrading to an earlier firmware version might fail to downgrade the MC properly.	Reload the desired firmware bundle.
Unable to create a peer connection after deleting a previous peer connection when using Asynchronous Replication.	Reboot both SCs at the same time using the CLI restart SC both command. Or, from the WBI, choose Restart System from the pull-down menu under the version number at the top. Select Restart with Storage and select both controllers.
In some cases, when moving controllers from a high-density chassis to a low-density chassis the proper EC component might not be loaded.	Loading code on the affected controller will update the EC to the appropriate version.
In a very rare case, after a controller failover, replication schedules might be removed.	Manually recreate the replication schedule.
An aborted replication might show status as “Unsynchronized” rather than “Aborted.”	This is a display issue only and can be ignored. No workaround is necessary.
In the online help for the “Scheduling replications” task, some details are missing, including the permitted number of characters and which characters are invalid.	The limitations on replication schedule names are the same as for replication set names. Refer to the disk management utility guide for limitations on replication names replication set names.
When trying to delete a pool with a replication set defined, this error message might display: Error: Command is not supported. - Failed to delete the virtual pool.	Remove the replication set and then delete the pool.
Modifying the password for an SNMP v3 user might default the authentication type and privacy type to none.	When modifying the password for an SNMP v3 user, also modify the authentication type and privacy type to the desired values.
Email notification for “Event 175: The network-port Ethernet link is up for a controller” might not be sent.	Use the WBI or CLI to check the status of Ethernet links.
If a replication set has entered a suspended state, the events may be inconsistent between the peer systems with respect to the reason for suspension.	Review the events on both peer systems.
CLI help for de-quarantine shows example of virtual disk groups, but virtual disk groups cannot be de-quarantined without support assistance.	None

Issue	Workaround
If a replication is suspended by the system, the array with the secondary volume may incorrectly report an event that the replication was suspended by the user.	None. On the first subsequent attempt to resume the replication, if the error condition still exists, the user will get an appropriate error message.
A very long user name will partially hide the session timer in the disk management utility (V3 GUI).	None
The modify peer-connection section in the disk management utility (V3 GUI) online help incorrectly states that a new name can have 80 bytes. The correct maximum number of bytes is 32 bytes for peer names. Anything over 32 bytes will be rejected.	None
The confirmation message received in the disk management utility (V3 GUI) Configuration Wizard does not clearly indicate which services are disabled when notifying the user that disabling an interface will disable the interface.	None. The selected services are disabled. Verify services in the disk management utility (V3 GUI).
After installing a license file in the disk management utility (V2 GUI) interface, the Browse button for installing a second license file may not allow changing locations.	Refresh the browser session.
Downgrading firmware may fail due to enabled features that are only available in the latest firmware being unavailable in previous firmware versions.	None
After performing a trust command in the CLI, hosts may not be able to write to volumes.	Perform a rescan from the CLI.
Operations are allowed in the primary system on a replication set when the same replication set is not available in secondary system.	Review Events to identify if operations are successful.
The maximum number of characters that can be used in modify schedule is not included in the disk management utility (V3 GUI) online help.	The maximum number of characters is limited to 26.
When trying to modify a peer with a remote address of a system for which another peer connection already exists, you get the error Error: The replication set was not found.	Verify the remote system is part of the peer connection you are attempting to modify.
The amount of space reported as used in a snap-pool may be larger than the sum of all snapshot data contained in the snap-pool.	None. The space calculation difference is due to allocation granularity differences in snapshots compared to the snap-pool.
If NTP is enabled, using set controller-date in the CLI will not perform any validation of the date format.	When setting the controller date with NTP enabled, ensure that you use a valid date format.

Issue	Workaround
The CLI command show network-parameters shows 'link-speed' as '10mbps' and 'duplex-mode' as 'half' when an Ethernet port is unplugged or a controller is uninstalled.	No workaround is required; ignore invalid output.
The CLI command set volume is allowing the use of the invalid double-quote and comma characters.	Do not use the double-quote or comma characters.
In the disk management utility (V3 GUI) the SNMP Read and Write community strings allow the use of the invalid single-quote character.	Do not use the single-quote character.
The version of a failing FRU may not be in the event logs.	Look at the label on the FRU for the version number.
In the save logs operation, if a single-quote character is used in any field, the retrieved logs will be empty.	Do not include single quotes when using the save logs operation.
A very long user name will be partially hidden under the Sign Out button.	None.
In the disk management utility (V3 GUI) or disk management utility (V2 GUI) interface, the forward slash character may be replaced with a space character.	Fields that require the use of a forward slash should be entered via the CLI.
The use of the angled brackets when creating vdisks names is erroneously allowed in the CLI.	Do not use angled brackets when creating vdisks names in the CLI.
An erroneous event is generated when telnet is disabled or enabled in the disk management utility (V3 GUI) or disk management utility (V2 GUI) interface.	Ignore the erroneous event.
Minimal discovery information is retrieved from SNMP v1 when SNMP v3 is enabled.	None.
The creation of CHAP records erroneously accepts names not in IQN format.	Use only valid IQN formatted names. See Calculating Maximum IQN Name Length on page 27 for additional information.
Event Code 522 does not display the name of the volume when the show events command is run.	None.
When using the disk management utility (V2 GUI) interface, during the add replication set operation, an error message can be shown twice.	Ignore the duplicated error message.
Cannot perform actions on multiple replication sets in the disk management utility (V3 GUI) or disk management utility (V2 GUI) interface at the same time.	Perform actions on replication sets one at a time.
Disk scrub status is not available when displaying disk information.	Review event logs for disk scrub status.

Issue	Workaround
Drives in LEFTOVER state may not illuminate the amber fault LED.	Validate drive status in the user interfaces.
SMI-S will report the device ID of IOMs in disk enclosures as "0" for the B controller and "1" for the A controller.	Use the ElementName to determine whether the A or B controller is in use.
The CLI show fans command displays fan status as up for a removed power supply.	Use the CLI show power-supplies command to see the state of the power supplies.
The disk management utility (V2 GUI) interface does not enable the option to set a volume as a primary volume for local replications.	Use the CLI.
A controller crash may occur if backend cabling is mis-configured.	Follow the documented cable configuration guidelines.
FDE Lock Key IDs may display inaccurately if reconfiguration of FDE features is attempted after clearing the Lock Key IDs and before power cycling the system.	The system must be power cycled immediately after clearing the FDE Lock Keys.
Some 314 events are not being reported by SNMP trap for the partner controller.	Register for events for both controllers.
The disk management utility (V2 GUI) interface returns an error stating a value must be an unsigned integer when creating a volume set and the number of volumes is blank or not a number.	Enter a valid number in the number of volumes field.
In step 4 of the disk management utility (V2 GUI) interface's provisioning wizard, the size field may not immediately update after entering the number of volumes.	Press the TAB key twice to have the field updated.
The disk management utility (V3 GUI) or disk management utility (V2 GUI) interface does not allow for all valid special characters to be used in an email notification sender name.	Use the CLI command set email-parameters .
Messages are not translated while collecting logs.	None.
In virtual storage remote replication, the secondary site could experience issues that could cause the replication to suspend.	The replication will automatically resume in 10 minutes or the user can manually resume it.
In single controller mode, resetting all statistics can result in a message that the request cannot be completed because the other controller is not up. - Failed to reset controller statistics.	Ignore the error message; all statistics that can be reset are reset.
Firmware update status messages in the disk management utility (V3 GUI) or disk management utility (V2 GUI) interface are not translated.	None.

Issue	Workaround
The Add Initiator to Host panel in the disk management utility (V3 GUI) can display “no host” as an option.	Ignore the erroneous option.
The CLI command clear disk-metadata shows a successful command stating that the command completed successfully, but the metadata was not cleared.	Run the command again to clear the disk metadata.
In disk management utility (V3 GUI) or disk management utility (V2 GUI) popups, buttons aren't enabled until clicking outside of the data entry fields or pressing the Tab key.	Select the Tab key or click outside of the entry fields in order to enable the buttons.
Partner Firmware Update can fail if there is a health issue in the system.	Determine array upgrade health by running the CLI check firmware-upgrade-health command. Address the health issue and reload the code on the controller.

Calculating Maximum IQN Name Length

When using explicit LUN mapping and you must shorten the IQN name on the nodes, use this procedure.

The following formula is used to calculate the maximum IQN name length, based on the number of hosts being explicitly mapped to a LUN on the array:

$$\text{Maximum IQN Character Length} = (4083 / n) - 7$$

Where n = number of hosts nodes explicitly mapped to a LUN on the array.)

Note: By specification, 223 is the maximum IQN length allowed.

[Table 2](#) on page 28 provides the calculated values based on the number of hosts being explicitly mapped to a LUN on the array:

Note: The table headings identify the following information:

of Hosts = Number of hosts explicitly mapped to a LUN
Max IQN = Maximum IQN Character Length

Table 2 Calculated Values for Hosts

# of Hosts	Max IQN	# of Hosts	Max IQN	# of Hosts	Max IQN
1-7	223	33	116	49	76
18	219	34	113	50	74
19	207	35	109	51	73
20	197	36	106	52	71
21	187	37	103	53	70
22	178	38	100	54	68
23	170	39	97	55	67
24	163	40	95	56	65
25	156	41	92	57	64
26	150	42	90	58	63
27	144	43	87	59	62
28	138	44	85	60	61
29	133	45	83	61	59
30	129	46	81	62	58
31	124	47	79	63	57
32	120	48	78	64	56

Contacting Quantum

More information about this product is available on the Service and Support website at <http://www.quantum.com/ServiceandSupport/Index.aspx>. The 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 the Quantum Customer Support Center:

United States	1-800-284-5101 (toll free) +1-720-249-5700
EMEA	+800-7826-8888 (toll free) +49-6131-3241-1164
APAC	+800-7826-8887 (toll free) +603-7953-3010

For worldwide support:

<http://www.quantum.com/ServiceandSupport/Index.aspx>
