

QXS 12G (Gen1) G2xxxxxx Release Notes

Contents

- Supported QXS Firmware Releases, Products, Operating Systems, and Storage Configurations 2
- Enhancements and Fixed Issues 4
- Known Issues and Workarounds 12
- Update QXS Firmware 23
- Calculating the Maximum IQN Name Length 29
- Contacting Quantum 30

© 2023 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law. ActiveScale, DXi, DXi Accent, FlexSync, FlexTier, iLayer, Lattus, Quantum, The Quantum Logo, QXS, Scalar, StorNext, SuperLoader, Vision, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Quantum specifications are subject to change.

Supported QXS Firmware Releases, Products, Operating Systems, and Storage Configurations

Supported QXS Releases and Release Notes Dates

The following releases **ONLY** are supported by these Release Notes:

QXS Firmware Releases	Release Notes Date
GT280R011-01	May 2023
GT280R010-01	November 2022
G280R009-02	September 2021
G280R006-02	October 2020
G265R009	April 2019

Supported QXS Products

The following products **ONLY** are supported by the firmware releases covered in these Release Notes:

- QXS-412 (12-drive), and QXS-424 (24-drive), and QXS-484 (84-drive) 12G (Gen 1) systems
- QXS-312 (12-drive), and QXS-324 (24-drive) 12G (Gen 1) systems

Supported QXS Operating Systems

The following Operating Systems **ONLY** are supported by the firmware releases covered in these Release Notes:

Operating System	Version
RBOD Chassis	
Microsoft Windows Server	2022, 2019
SuSE Linux Enterprise Server	15.3
Ubuntu	19.04
VMware ESXi Server	7.0 U3

Operating System	Version
Red Hat Enterprise Linux	8.5, 7.9, 7.8
JBOD Chassis	
Microsoft Windows Server	2022, 2019
Red Hat Enterprise Linux	8.5, 7.9, 7.8
SuSE Linux Enterprise Server	15.3

Supported QXS Storage Configurations

This section provides the QXS 12G (Gen1) supported configurations for RAID chassis, drives and drive form factor, and the number of expansion chassis supported using the current firmware:

Drive Expansion Chassis	QXS-312 12G (Gen1) RAID Chassis	QXS-412 12G (Gen1) RAID Chassis	QXS-324 12G (Gen1) RAID Chassis	QXS-424 12G (Gen1) RAID Chassis	QXS-484 12G (Gen1) RAID Chassis
Drive Form Factor Per Chassis	12-drives (3.5" LFF)	12-drives (3.5" LFF)	24-drives (2.5" SFF)	24-drives (2.5" SFF)	84-drives (2.5" SFF or 3.5" LFF)
2U12 Expansion Chassis	Up to 3	Up to 9	Up to 3	Up to 9	N/A
2U24 Expansion Chassis	Up to 3	Up to 9	Up to 3	Up to 9	N/A
5U84 Expansion Chassis	N/A	Up to 3	N/A		Up to 3
*Intermixed 2U Expansion Chassis	Up to 3	Up to 9	Up to 3	Up to 9	N/A

* The 2U12 and 2U24 expansion chassis can be intermixed. However, mixing 2U12/2U24 expansion chassis and 5U expansion chassis is not supported.

If you have supported and/or unsupported configuration questions, [Contact Quantum Services and Support](#).

Enhancements and Fixed Issues

This section lists the fixed issues and enhancements for the different Releases supported for your system.

- [GT280R011-01 Enhancements and Fixed Issues below](#)
- [GT280R010-01 Enhancements and Fixed Issues on the next page](#)
- [GT280R009-02 Enhancements and Fixed Issues on the next page](#)
- [GT280R006-02 Enhancements and Fixed Issues on page 8](#)
- [GT265R009 Enhancements and Fixed Issues on page 8](#)

GT280R011-01 Enhancements and Fixed Issues

Enhancement:

This update supports the latest G280 Lattice CPLD-support enhancement and the associated known issues. If an existing Intel CPLD fails, it must be replaced with a Lattice CPLD. This means the IOM or sideplane FRU that contains the CPLD must be replaced. For a successful fix, the firmware must be installed that supports the new Lattice CPLD prior to replacing the CPLD.

The following known issues exist with Lattice CPLD:

- The new Lattice CPLD is not supported by firmware version GN280R010-01 or earlier. Attempting to use that firmware with hardware using Lattice CPLD will result in unexpected behavior. The system must be updated to firmware version GN280R011-01.
- [Contact Quantum](#) if the firmware update fails with the following code load error:

```
SP: Error: Fatal
```

- In the case where a 5U84 expansion enclosure using the new CPLD is attached to a 5U84 system running GN280R010-01, and firmware update to GN280R011 is started with the PFU option enabled, the update may fail with a **Failed** error, which indicates a failure for the controller to migrate the update code to the other controller. If this occurs, restart both controllers and repeat the firmware update. This is a one-time workaround; subsequent firmware update operations should proceed normally.

Fixed Issues:

- (Reference Issue: FMW-61094) Fixed a situation where an RBOD-Sideplane cable/phy missing fault was showing from the Management Controller and displayed as "Electronic failure".
- (Reference Issue: FMW-61292) After updating to GT280R011-01 firmware, the sideplane now updates as "success".
- Fixed an issue where SNMP service stopped responding.

- Fixed an issue where SMI-S API broke after firmware upgrade to GT280R010-01.
- Added Lattice CPLD support.

GT280R010-01 Enhancements and Fixed Issues

- Fixed Overlapped command Check condition.
- Fixed an issue in the Historical Performance Statistics in the QXS **DMU** (Disk Storage Management Utility v3/GUI) and the CLI.
- Fixed a controller crash issue when global spare was assigned.
- Fixed an issue where Dell product serial numbers were not correctly getting recorded in system logs.
- Added support to SWAP bit using storage controller.
- Fixed an issue where a multicore system crashed with errors related to Macro Level I/O (MLIO).
- Fixed an issue where PCM firmware update failed with an I/O timeout error.
- Improved the event logging mechanism in storage controller logs.
- Fixed the ADAPT map corruption issue.
- Fixed an issue of swap bit of Host temp sensor in SES page 2.
- Fixed an issue where some host mappings were lost after cold booting a storage system.
- Fixed "send syslog" and SNMP notification issue when the system is configured with FQDN.
- Fixed an issue where a virtual pool's high capacity threshold was reached and controller A crashed.
- Enhanced DDR error handling in response to a dual-controller crash.
- Added drive status check in error handler to issue FDE init or retry.
- Fixed an issue where implementation of 802.1x certificates does not check for any activity in progress before activating the WPA 802.1x certificates.
- Increased sync IO timeout and re-issued original srb.
- Fixed an issue where a controller killed its peer/partner controller due to heartbeat loss.
- Fixed an issue where a data-unavailability situation occurred due to rebooting of a controller.

GT280R009-02 Enhancements and Fixed Issues

- The Management port on the chassis is now compliant with IEEE 802.1x network authentication protocol.
- Added Storage Controller capability to detect Field Accessible Reliability Metrics Specification time series log.
- Added support to invoke and monitor the Unified Data Service.
- Added Unified Data Service set protocol and show protocol commands.
- Changed the default scrub interval from 24 hours to 360 hours (15 days).

- Added support for collecting logs.
- Added support to collect drive logs for Field Accessible Reliability Metrics frames.
- Added capability to allow full resynchronization of replication sets at will.
- Developed Unified System Management (USM) Upgrade Toolkit for 5U84 and 2U12/2U24 JBODs.
- Fixed an issue where both controllers crashed after controller B was brought back up while deleting a volume.
- Fixed an issue where a dual controller crashed during the disaster reverse recovery operation, done using kill/unkill commands on the system.
- Fixed an issue where clear expander-status fails to clear expander-status for 5U84 chassis.
- Fixed an issue where a crash happened due to a timeout detected by the auto stall recovery code during the shutdown/reboot test.
- This issue happened when couple of disks become bad during the shutdown/reboot test, due to which the disk group got quarantined, and one of the controllers detected a 'shutdown stall' for the other controller.
- Fixed an issue where the wrong disk was being failed during a bad-block recovery.
- Fixed an issue of offline disk groups remaining quarantined even after the controller was live.
- Implemented a salvage mode for field recovery red-alerts.
- Fixed issue where the incorrect chassis was displayed on the **Configuration** page of the **DMU** (Disk Storage Management Utility v3/GUI), and incorrect chassis configurations were swapped, when the system is busy in SCSI rescanning.
- Fixed an issue of controller crash due to heartbeat loss.
- Fixed the wrong reporting where power supply failure reported capacitor failure.
- Fixed an improper system request for a "Performance Tier" license after upgrade to G280R002-03.
- Fixed an issue in the WBI where the Create Snapshot scheduler did not correctly convert AM/PM schedules (12-hour time format) to 24-hour time format.
- Fixed an issue where the controller crashed while deleting read-cache disk group.
- Fixed an issue where a user could successfully upload an expired certificate and install it to the array.
- Added the ability to retrieve UDS and SM2 logs from the storage system without removing the disks. The logs can be retrieved using the get logs command or the ftp command.
- Fixed an issue where, after firmware update, power supply firmware and VPD versions were not updated in the output of the CLI show versions **frus** command.
- Fixed an issue where disk scrub was being attempted on FDE locked disks, causing them to be marked as "Leftover".
- Fixed an issue where the replaced controller's IP address changed back to the controller's old IP address after chassis power cycle or shutdown restart of both controllers.

- Deprecated no-mirror cache optimization mode.
- Global spare drive now gets selected by slot affinity and gets added as disk group spare.
- Fixed the nested topology change issue.
- Fixed issues where error detected in scrub caused controller to get killed.
- The complete debug region memory buffer can now be used to extract disk logs.
- Fixed self assert in case of compact flash diagnostic failure during boot.
- QXS storage systems now implement mechanism of self-detection of multi boot failure and avoids a re-boot.
- For ADAPT disk groups, fixed an issue where disks that were not used in an expand operation were associated with the disk group during rebalance operation.
- Made interprocessor communication (IPC) path changes to handle multi-context invocation and incorrect error address reporting.
- Fixed an issue where after upgrading to GN280R002-03/GT280R002-03, system is asking for a "Performance Tier" license.
- Fixed an issue where it was not possible to unlock FDE disk after using the set fde-state repurpose command.
- Fixed an issue where while using WBI, it was not possible to set the SMTP server name if the domain suffix was more than 6 characters long.
- Fixed an issue where while receiving traps on MIB browser it was displaying IP address of a different controller rather than the address of the controller on which it was configured.
- Fixed an issue where the FDE passphrase was not masked in the GUI.
- Fixed a large memory leak in root platform.
- Fixed an error where the shared library libz.so was not loaded correctly.
- Fixed an issue of flooding of logs due to repeated SCSI enclosure services (SES) page 2 control requests.
- Fixed an issue where the virtual pool went offline due to unreadable metadata.
- Fixed the GEM CLI ddump_pwrmgr to show correct power cooling module /powersupply unit serial no.
- Fixed disk-zone distribution among stripe zones in ADAPT rebalance after disk replacement.
- Fixed a race condition that happened between task management function (TMF) sent and the TMF aborted.
- Fixed issue where a disk-group initialization failure and then a scrub abort on the same disk group caused the disk group to go offline.
- Fixed the issue of data corruption that happened during adapt target rebuild and rescan.
- Implemented security enhancements.

GT280R006-02 Enhancements and Fixed Issues

i Note: Quantum strongly recommends changing the default password for upgraded systems.

- Removed the default username named default and the default password for the default username of !manage. This applies ONLY to:
 - New QXS arrays with G280 code pre-installed
 - QXS arrays that have been re-imaged with G280 code

This has the following effects for new QXS arrays:

- a. When accessing the CLI or the **DMU** (Disk Storage Management Utility v3/GUI) for the first time, you will have to create a new user and password for the system. See the **QXS 12G Maintenance and Installation Guide**, the **QXS 12G CLI Reference Guide**, and the **QXS 12G Disk Management User Guide** for new CLI and GUI login procedures.
 - b. The “manage” and “monitor” roles will both be assigned to the first configured username created on the array.
 - c. You can no longer use default username and !manage password to access these storage arrays.
 - d. If you use the Appliance Controller qxsarray commands from the host StorNext system, you must use the username and password assigned to the system in order to use those commands, unless the system is running software where the manage username and !manage password were in place.
- Security of Connected devices
 - Security Vulnerability Scan (Appscan and Nessus) as part of code qualification
 - USGv6 cert for Data path and Management path
 - Adds Win 2019 for General Market plus latest version OS level updates
 - GT280R010-01 firmware release aligns the overall customer base to the ‘G275xxxx’ code base from the ‘G265xxxx’ code base
 - Maintains the functionality pre-existing in G265xxxx and provides unified code support
 - All G265xxxx fixes ported to G275xxxx base
 - Inherits improvements put into G275xxxx (ADAPT sequential write performance enhancement)
 - Added **OBC** (One-Button-Config) storage configuration profiles to the **DMU** (Disk Storage Management Utility v3/GUI) for the 5U84 system
 - Fixed an issue where Controller A would cycle between on and off states (Reference SR 512681)

GT265R009 Enhancements and Fixed Issues

- Added the ADAPT data protection level for linear and virtual disk groups. This feature provides:
 - Maximizes flexibility
 - Provides built-in spare capacity

- Allows for very fast rebuilds, large storage pools, and simplified expansion
- Added support for 5U84 controller and expansion chassis (a 5U chassis containing up to 84 disks).

i Note: A storage system can have up to three 5U84 expansion chassis.

- Added a JSON API output option to the CLI.
- Increased the maximum virtual pool size per controller module to 1 PiB, with the large-pools feature enabled.
- In the **DMU** (Disk Storage Management Utility v3/GUI), introduced a guided-setup interface to speed initial configuration and provisioning of the storage system.
- Re-enabled support for SMI-S for storage systems with 2U chassis only.
- Re-enabled Microsoft VSS support through use of the Seagate VSS Hardware Provider.
- Drive spin-down can be scheduled for linear disk groups and disks not in a disk group.
- For virtual storage, increased performance to 600K IOPS for 5005 systems and to 300K IOPS for 4005 systems.
- Added an option to enable/disable CHAP on the **Configure CHAP** panel of the **DMU** (Disk Storage Management Utility v3/GUI).
- Added new virtual-pool statistics to CLI API output. These statistics include:
 - Page allocations/deallocations
 - Page maps/unmaps
 - Hot/cold page moves
 - Blocked SSD promotions
- Fixed strings that were not translated in the **DMU** (Disk Storage Management Utility v3/GUI), CLI, and event messages.
- Fixed an issue where, in the **DMU** (Disk Storage Management Utility v3/GUI) when mapping a volume to a host or host group on a hybrid FC/iSCSI controller, all four host ports were selected by default despite the configured host protocol.
- Fixed alignment issues that existed when using the German locale in the **DMU** (Disk Storage Management Utility v3/GUI).
- Fixed an issue where remapping host ports caused LUN IDs to be reset.
- Fixed an issue where trying to modify a peer with the remote address of a system for which the peer already exists displayed the unintuitive error message Error: Cannot find peer connection in the remote system.
- Fixed an issue where a PFU (partner firmware update) script failed with an XML Parse Failed error.
- Fixed an issue where the **DMU** (Disk Storage Management Utility v3/GUI) did not use a string to indicate "no host" for an initiator

- Removed linear-storage references from event messages for storage systems that support only virtual storage.
- In the CLI, improved the message shown after successfully renaming a volume group.
- The confirmation message received in the **Configuration Wizard** of the **DMU** (Disk Storage Management Utility v3/GUI) did not clearly indicate which services will be disabled when notifying the user that disabling an interface will disable the interface.
- Fixed an issue where, when a user name with maximum character length was logged in, the session time information was not displayed on the **Home** page of the **DMU** (Disk Storage Management Utility v3/GUI).
- Changed help to specify that volume-groups can only include volumes that are in the same pool.
- Fixed an issue that could cause a code load failure to occur, especially with one controller module in the system.
- Fixed an issue where disks were left in an available or leftover state after a disk group was pulled from one system and inserted into a different system that is running.
- Fixed an issue where an FTP session terminated during code load.
- Fixed an issue where, after a code load from the **DMU** (Disk Storage Management Utility v3/GUI), the old version of code was still shown.
- Changed the firmware so that if an unsupported (OneStor) expansion chassis is attached to the system, the chassis status is shown as Unsupported in CLI API output.
- Fixed an issue where reversed logic for reporting virtual disk-group removal status displayed a confusing message. When the disk group was OK to remove, the message said not OK to remove, and vice versa.
- Fixed an issue where the CLI show sas-link-health command incorrectly showed the expansion port status Up even when the cable is unplugged from the controller.
- Fixed an issue where, after a test of pulling all host cables from both systems during replication, the replication state for the local system went to resume state but on the other system it went to suspended state in the **DMU** (Disk Storage Management Utility v3/GUI). Similarly, the peer connection health status showed OK on one system and Faulty on the other system.
- Fixed an issue that could cause a controller to crash while running snapshot schedules.
- Fixed an issue where, on the **System** page in the **DMU** (Disk Storage Management Utility v3/GUI), information disappeared after removing a disk.
- Fixed an issue where, while running a heavy virtual replication load, and attempting a code load, the Management Controller crashed.
- Fixed an issue where the following messages appeared while loading code:

```
Firmware Bundle - Transfer completed.
Please wait...
```

```
Please wait...
Please wait...
Please wait...
Please wait...
Please wait...
Timed out waiting for partner to connect.
No reboot required
RETURN_CODE: 10
```

- Fixed an issue where the following messages appeared while loading code:

```
The code load failed but the MC was restarted:
System health check failed.
*** Code Load Fail. ***
STATUS: *** Code Load Fail. ***
Restarting Management Controller...
RETURN_CODE: 10
```

- Removed the CLI set cli-parameters command's xml option, which was not documented.
- Fixed an issue where, on the **Replications** page in the **DMU** (Disk Storage Management Utility v3/GUI), hovering on a replication set name did not display replication progress.
- Fixed an issue where, after loading the same firmware bundle onto a system with PFU (partner firmware update) was enabled, the **DMU** (Disk Storage Management Utility v3/GUI) reported a successful code load process but then showed "Firmware upgrade failed."
- Fixed an issue where, after starting a firmware downgrade using FTP, SNMP trap event 237 reported System health is insufficient to support firmware update yet the FTP health state passed and code load continued.
- Fixed the Fan and Power Supply FRU status for Warning and Resolved events.
- Fixed an issue where code load failed with error message: Code Load Rejected. Does not appear to be a bundle at all.
- Fixed an issue where during code load the Management Controller experienced a Kernel panic: Oops - BUG: 0 [#1]
- Fixed an issue where, in the **DMU** (Disk Storage Management Utility v3/GUI), informational events that indicate a fault has been resolved were shown without a severity icon.
- Fixed an issue where auto-deletion of eligible snapshots stopped working, and new snapshots could not be created until the snapshot space limit was increased.
- Fixed an issue where initiating code load while a code load is in progress didn't return the correct message—instead it returned: Could not determine if there is an activity in progress on this system.

- Fixed an issue where code load failed while getting logs, resulting in incorrect firmware image being loaded.
- Fixed an issue where a controller crashed while loading code on a system configured to use virtual replication.
- Fixed an issue where firmware download was interrupted due to loss of power to controller or Storage Controller restart, which resulted in an incorrect firmware image being loaded due to a later timestamp on the interrupted bundle.
- Fixed an issue where, on the **Map** panel of the **DMU** (Disk Storage Management Utility v3/GUI), using filters caused the Next and Previous arrow buttons to become obscured.
- Fixed an issue where, while pulling logs via FTP or the **DMU** (Disk Storage Management Utility v3/GUI), an error message stated that logs can't be pulled because a code load is in progress.
- Attempting to log in through Telnet (if enabled) while a code load is in progress caused the Management Controller to display the message:

```
/home/appuser/appshell: source: line 69: can't open '/app/etc/app.env'.
```

You still must wait until the code load is complete before logging in, but the message has been changed to:

```
Login unavailable. Please try again later.
```

- Fixed an issue where in management interfaces, FRU data was not available for a power supply.
- Changed events 39, 40, 307, 469, 476, 477, and 524 to specify that the operating temperature range for a controller chassis is 5–35°C (41–95°F) and for an expansion chassis is 5–40°C (41–104°F).
- Fixed an issue that prevented using the **DMU** (Disk Storage Management Utility v3/GUI) to switch from static to DHCP addressing for the management port, in single controller mode.

Known Issues and Workarounds

This section lists the known issues that could potentially affect your system.

- [G280R009-02 Known Issues and Workarounds on the next page](#)
- [G265R009 Known Issues and Workarounds on page 15](#)

G280R009-02 Known Issues and Workarounds

Description

Issue:

Controller A crashed during a controller pull/push operation.

Workaround:

None. This is a rare occurrence.

Issue:

The **DMU** (Disk Storage Management Utility v3/GUI) error message (schedule was not found) is not appropriate while modifying/expanding the size of a volume.

Workaround:

Use the CLI.

Issue:

SSD disk groups can be created without a Performance Tier license.

Workaround:

None.

Issue:

Incorrect queue depth values are displayed for historical performance statistics.

Workaround:

None.

Issue:

The maximum volume size limit for an ADAPT virtual or linear disk group is 1 PiB ; for a non-ADAPT virtual disk group it is 128 TiB; and for a non-ADAPT linear disk group the maximum size is limited only by the size and capacity of member disks.

Workaround:

None.

Issue:

When the drive spin down feature is enabled, it is not possible to unlock the FDE drives.

Workaround:

Disable the drive spin down feature and wait for approximately 2-3 minutes, and perform a manual rescan. Once the drives spin up, the system can be unlocked by entering the FDE passphrase.

Description**Issue:**

The Management Controllers reboot in a loop after creating a monthly snapshot schedule.

Workaround:

While creating a scheduled snapshot, clear the Repeat every checkbox.

Issue:

One controller crashes during repetitive drive pull and insertion on RAID 10 disk group.

Workaround:

None.

Issue:

A disk detected timeout when you do a repetitive shutdown, restart with adapt reconstruction and manual rescan.

Workaround:

None.

Issue:

When using the **DMU** (Disk Storage Management Utility v3/GUI) to replace a disk, the health panel does not correctly reflect the LED status of the disks until the disk group rebuild starts. Any amber LEDs illuminated on the disks are not shown as illuminated in the health panel.

Workaround:

The **Health** panel is updated after the disk group is rebuilt.

Issue:

In the Storage Management Console and when using the CLI set chap-record command, the secret and mutual-secret parameters accept the entry of a double-quote character ("), which is an invalid character.

Workaround:

If a double-quote character was included in the secret or mutual-secret parameter, remove it and re-create those entries.

Issue:

When creating a username during initial CLI login, the following invalid characters are allowed in the username: angle brackets (< >) and single quote (').

Workaround:

Re-create the username without using angle brackets (< >) and single quote (').

G265R009 Known Issues and Workarounds

Description

Issue:

RAID write errors occur and message RAID had write error appears in the event log.

Workaround:

None, but rebooting the Storage Controllers one at a time will clear the write errors from the system.

Issue:

If Telnet is enabled, the Telnet session to a controller may drop for no reason.

Workaround:

Attempt to restart the Telnet session. If not successful, restart the Management Controller from the peer controller.

Issue:

A controller's management port stops working.

Workaround:

Restart the affected controller's Management Controller from the peer controller.

Issue:

Some strings are not translated in the **DMU** (Disk Storage Management Utility v3/GUI) and in event messages.

Workaround:

Use a translation service, such as Google Translate, to translate unknown English words into the preferred language.

Issue:

In the **DMU** (Disk Storage Management Utility v3/GUI), forward slash characters (/) in name fields are replaced with spaces.

Workaround:

Locations or other names that require the use of a forward slash (/) should be entered via the CLI.

Issue:

When replication peers are in different time zones, incorrect values are shown for replication **Start Time** and **Estimated Completion Time**.

Workaround:

Restart the Storage Controller that shows the incorrect time.

Description**Issue:**

The CLI create snapshot command does not handle listing multiple volume or snapshot names that are enclosed in quotes, if any of the names contain a space.

Workaround:

Run the command separately for a volume or snapshot whose name includes a space.

Issue:

During Storage Controller reboot, a decompression issue extends boot time by a few seconds.

Workaround:

None.

Issue:

In the **DMU** (Disk Storage Management Utility v3/GUI), when using Google Chrome, the Turn off LED button in the System topic does not work.

Workaround:

Use a different browser.

Issue:

The CLI reset disk-group-statistics command does not reset the counters for the disk groups.

Workaround:

The counter for disk group statistics can be reset using the reset all-statistics command.

Issue:

The CLI **show snapshots** command may display an incorrect unique-data value. After taking a snapshot of an existing snapshot, the unique-data value will be reset to 0. Despite this, the **show snapshot-space** command will show the allocated space correctly.

Workaround:

None.

Issue:

In the **DMU** (Disk Storage Management Utility v3/GUI), after changing the host port mode from iSCSI to FC, the "replicate now" and "replication set status" fields show incorrect status during replication.

Workaround:

Stop replication before changing the host port mode.

Description**Issue:**

Cannot delete more than 50 volume groups at a time in the **DMU** (Disk Storage Management Utility v3/GUI) or the CLI.

Workaround:

Delete up to 50 volume groups, then delete the remainder.

Issue:

In the **DMU** (Disk Storage Management Utility v3/GUI), when performing a complete tear-down of volumes where you're deleting a volume group and select a check box to delete all the volumes in the group, the volumes are not deleted and the operation fails with error:

Unable to delete the volume as the volume is part of volume group

However, the volume group is not identified.

Workaround:

Perform a rescan to repopulate the internal system with fresh data, and then attempt to delete all the volumes in the group again.

Issue:

In the **DMU** (Disk Storage Management Utility v3/GUI), when I/O is in progress, a failed or removed controller causes the aggregate IOPS and MB/s fields in the Home topic to display 0 values. When the controller is reinstalled, proper values appear.

Workaround:

Both controllers must be present and active for the values to appear again.

Issue:

A replication set was not created correctly but the secondary volume was created on the peer system. Replications cannot run if there is no replication set between the primary and secondary volumes.

Workaround:

None. To remove the secondary volume, contact support.

Description**Issue:**

In the **DMU** (Disk Storage Management Utility v3/GUI), the Temporary License tab mentions Dot Hill in the End User License Agreement.

Workaround:

None

Issue:

The CLI **show power-supplies** command does not display the firmware revision.

Workaround:

Use the **show configuration** command instead.

Issue:

In the **Manage Schedules** panel of the **DMU** (Disk Storage Management Utility v3/GUI), the last scheduled task that was deleted is not grayed out.

Workaround:

Refresh the browser.

Issue:

In the **User Management** panel of the **DMU** (Disk Storage Management Utility v3/GUI), if you try to create a user with a name that already exists, an error message will appear. If you then toggle the SNMPv3 setting the error message will be cleared.

Workaround:

When the **Apply** button is grayed out, check whether the name is already in use.

Issue:

The "Manage" and "Monitor" labels for roles are not consistently translated in the **DMU** (Disk Storage Management Utility v3/GUI).

Workaround:

None.

Issue:

When using Microsoft Edge, **DMU** (Disk Storage Management Utility v3/GUI) help text does not appear.

Workaround:

This is due to a Microsoft Edge issue. Use a different browser.

Description

Issue:

In the **DMU** (Disk Storage Management Utility v3/GUI), an error message does not immediately appear when a user enters invalid characters in the domain name of the e-mail parameters fields of the **Setup Notifications** dialog.

Workaround:

None; the user is informed of the invalid characters when the user clicks **Apply**.

Issue:

In the **DMU** (Disk Storage Management Utility v3/GUI), SNMPv3 user authentication passwords do not accept special characters.

Workaround:

Use the CLI **set user** command instead.

Issue:

After restarting one of the controllers, the **show disks perf** command shows incorrect data for disks that are not part of a disk group and have no jobs running on them.

Workaround:

Ensure both controllers are operating and re-run the **show disks perf** command.

Issue:

In the **Manage Schedules** panel of the **DMU** (Disk Storage Management Utility v3/GUI), when deleting a schedule, a success message will appear but the schedule may remain listed. Trying to delete the schedule again will result in the following error:

```
Unable to delete schedule. The schedule was not found.
```

Workaround:

Refresh the browser.

Issue:

If the user attempts to create a disk group with a name that contains the lowercase string "dga", and if the index ID matches, the command will fail.

Workaround:

Do not create a disk group with a name that contains the lowercase string "dga".

Description**Issue:**

While creating a peer connection, the new peer connection is not shown in the secondary system's **DMU** (Disk Storage Management Utility v3/GUI) until it is refreshed.

Workaround:

Refresh the browser on the secondary system or use the CLI.

Issue:

A single controller crashed during scheduled replication.

Workaround:

None needed. The controller automatically recovers immediately after the crash without user intervention.

Issue:

Replication set status differs in the primary and secondary systems during scheduled replication. The replication set status shows "Ready" for the secondary system and "Running" for the primary system, which is incorrect.

Workaround:

If this condition persists, manually abort the replication and restart it. Any queued replication will start after the abort.

Issue:

Sometimes the Management Controller is not ready for an extended period of time after controller is rebooted.

Workaround:

Wait until the Management Controller restarts and becomes available.

Issue:

An Expander Controller code-load failure signal can occur and the controller reports that it has an invalid bundle.

Workaround:

Reinstall the firmware bundle.

Issue:

In the Chrome browser, the **DMU** (Disk Storage Management Utility v3/GUI) fails to load. It will show as "Loading" until it times out.

Workaround:

A 60-second timeout has been implemented. Refresh the browser after the timeout.

Description**Issue:**

Verify operation returns false success message and success event for Adapt disk group even though the "Verify" operation is not supported on Adapt disk group.

Workaround:

Ignore the false return of success event for an ADAPT disk group.

Issue:

In a dual-controller system in which SFTP protocol is enabled, after replacing a controller module with one having down-level firmware, SFTP may be disabled on the new controller.

Workaround:

Enable the SFTP protocol on the new controller.

Issue:

A scrub error is detected on unallocated pages.

Workaround:

Perform the scrub procedure manually to fix the parity errors.

Issue:

In the **DMU** (Disk Storage Management Utility v3/GUI), the rear-view image of an expansion module shows the center port having a blue label instead of a yellow label.

Workaround:

None.

Issue:

When the **PFU** (Partner Firmware Update) is in progress, a system may temporarily report disk channel events on the controller being updated. These are not an indication of disk or I/O failures, and the events will subside after code load has completed.

Workaround:

None. Wait for code load to complete.

Description**Issue:**

Linear RAID-5 expand disk-group operation will begin but will fail with the following error:

```
Failed to expand disk group. Invalid disk(s) were specified.
```

Workaround:

Ensure that all disks used in the expand operation have the same protection-information (PI) format as the disk-group members. A disk's PI format can be found in the XML output of the CLI **show disks** command.

Issue:

The **DMU** (Disk Storage Management Utility v3/GUI) has an option to copy a volume from linear to virtual, which is not supported.

Workaround:

None. Ignore the option.

Issue:

When connected to an x3650 host with a Lenovo N2225 HBA and SLES 12 SP1, after manually stopping I/O one 5U84 JBOD is no longer displayed. (This is only a reporting issue; I/O is not interrupted.)

Workaround:

Power cycle the chassis.

Issue:

The **GEM driveinv** command does not list the temperature of all the disks in a 5U84 JBOD. (This is only a reporting issue; an overtemperature condition will still be handled by the firmware and fault status will be reported in SES, but will not be device specific.)

Workaround:

None.

Issue:

Cannot see all disks after a 5U84 JBOD is power cycled when connected to a 930-8e HBA (MegaRAID).

Workaround:

Follow the documented power sequence of powering up the expansion chassis first, then the server with the 930-8e HBA.

Update QXS Firmware

Firmware Update Notes

This section lists the known issues that could potentially affect your system.

 **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

- 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 update is performed while host I/Os are being processed, I/O load can impact the update process.
- Select a period of low I/O activity to ensure the update 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 chassis 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 the Update

Monitor the system display during firmware installation to:

- Determine update status
- See when the update is complete

After the Update is Complete

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

- Verify system status in the **Disk Storage Management Utility (DMU)/User Interface (UI)**, 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 .

Additional Firmware Notes

- Windows Server 2012 management integration:
 - Quantum recommends that you update the Windows cache by using a cmdlet command manually, after making any storage provision operations or changes in the QXS **DMU** or the **CLI**). Enter:

```
Update-StorageProviderCache -DiscoveryLevel Full -Name <storageProviderName>
```

- Quantum recommends that you use the QXS **DMU** to modify volume mappings, delete volumes, or modify volume names. Manually update Windows cache by using a cmdlet command . Enter:

```
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\CurrentVersion\Storage Management\EnableIndications value from 1 to 0. Then, restart the Windows Standards Based Storage Management Service.
 - If you want to enable the Indication subscription in 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, there are additional precautions to consider. See the section “Real World: Dynamic versus Basic Disks” 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.

Update the QXS 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 update.

IMPORTANT: In dual-module chassis, 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.

Update Firmware Using the DMU v3 UI

This section describes how to install the firmware using the **DMU**.


Follow these steps to install the firmware package using the **DMU**:

1. Obtain the firmware package from https://www.quantum.com/en/service-support/downloads-and-firmware/qxs-3-4_g2_12g/




Note: A valid QXS serial number for your system is required.

2. Save the downloaded file to a temporary directory on your system. The firmware filename is in the following format: GxxxRyyy-zz.bin.
3. In single-domain environments, stop all I/O to vdisks in the chassis before you initiate the firmware update.
4. Log in to **DMU**.
5. Go to the **Update Firmware** option. There are two ways to do this:
 - a. On the top of the home screen (Figure 1), click the first arrow on the upper right where the System name and Version are shown. Click **Update Firmware**:



- b. Click the **System** tab on the left navigation panel (Figure 2), click the **Action** button at the top of the system panel, and click **Update Firmware**:



i 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 firmware.
8. Wait for the installation to complete. During installation, each updated controller module will automatically restart.
9. In the **DMU**, verify that the expected firmware version is installed on each controller.

Update Firmware Via an FTP Connection

To install the firmware package using FTP:

1. Obtain the firmware package from https://www.quantum.com/en/service-support/downloads-and-firmware/qxs-3-4_g2_12g/

i Note: A valid QXS serial number for your system is required.

2. Save the downloaded file to a temporary directory on your system. The firmware filename is in the following format: GxxxRyyy-zz.bin.
3. In single-domain environments, stop all I/O to vdisks in the chassis before you initiate the firmware update.
4. Log in to the **DMU**, and gather the information you will need to for the FTP update:
 - 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.
5. In single-domain environments, stop I/O to vdisks in the chassis before starting the firmware update.
 - a. Open a command prompt (Windows) or a terminal window (UNIX/Linux/macOS), and navigate to the directory containing the firmware file to load.
 - a. Enter a command with the following syntax:

```
ftp <controller-network-address>
```

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

Example:

```
put GxxxRyyy-zz.bin flash
```

6. Wait for the update to complete. During update, each updated module automatically restarts.
7. If needed, repeat these steps to load the firmware on additional modules.
8. Exit the FTP utility.
9. Verify that the expected firmware version is installed on each module.
 - In the **DMU**, 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. In the **System Overview** panel of the **DMU**, look for the System Version information. If significantly more than an hour has elapsed and the components do not show that they were updated to the new firmware version, refresh your 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, log in to the **CLI**, and 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, log in to the **CLI**, and enter the **show disks** command. 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 from the **CLI**. (Do not restart the controllers; cycle power on the controllers.)
- If the **show disks** command is run from the **CLI** on all controllers, and is not successful, perform the **Firmware Update** again.

Calculating the 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.

The table below provides the calculated values based on the number of hosts being explicitly mapped to a LUN on the array:

of Hosts = Number of hosts explicitly mapped to a LUN

Max IQN = Maximum IQN character length

# of Hosts	Max IQN	# of Hosts	Max IQN	# of Hosts	Max IQN
1-17	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

# of Hosts	Max IQN	# of Hosts	Max IQN	# of Hosts	Max IQN
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

For further assistance, 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:

<https://www.quantum.com/en/company/contact-us/>
