

QXS 12G (Gen1) G2xxxxxx

Release Notes

Product/Software Release Date	December, 2019
-------------------------------	----------------

Contents

What's New	2
Update QXS Firmware	15
Contacting Quantum Support	21
Submit a Ticket (Service Request)	22
Use MyQuantum Service Delivery Platform	22
Use Cloud Based Analytics (CBA)	23
Escalate a Case	23
Contact Quantum Sales	23

© 2024 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. Myriad, 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.

What's New

This package delivers firmware for 5005/4005/3005 Series controller enclosures and expansion enclosures. Expansion-enclosure firmware is embedded in GN280R014-01/GT280R014-01 firmware bundles and is updated when attached to one of the following controller enclosures.

Model	Firmware version
5005	GN280R014-01
4005	GT280R014-01
3005	GT280R014-01

Update recommendation

This is a recommended firmware update for 5005/4005/3005 Series products.

Operating Systems

Supported operating systems include the following.

Controller/expansion enclosures

- Microsoft Windows Server 2022, 2019
- Red Hat Enterprise Linux 8.5, 7.9, 7.8
- SuSE Linux Enterprise Server 15.3
- Ubuntu 22.04, 20.04

JBOD enclosures

- Microsoft Windows Server 2022, 2019
- Red Hat Enterprise Linux 8.5, 7.9, 7.8
- SuSE Linux Enterprise Server 15.3
- Ubuntu 22.04, 20.04
- VMware ESXi Server 7.0 U3

Features or Enhancements Introduced in GN280R014-01/GT280R014-01

- Security enhancements implemented.
- In the CLI, the `clear degraded-disk` command can now be used to set a degraded disk's health to OK and clear the disk's error count. After this command is run, error counts used to measure disk medium errors are restarted. The minimum role to use the command is Standard. The syntax is `clear degraded-disk <disks>`. For information about specifying disks, see the command syntax topic in the CLI Reference Guide.
- Added firmware support for 764W AC PCMs and updated firmware for 764W DC PCMs.

Issues fixed in GN280R014-01/GT280R014-01

- Fixed an issue where an EMP reported an alert condition on a voltage sensor.
- Fixed an issue where a pool was offline due to inconsistent metadata.
- Fixed an issue where, after a brief power loss, disks in an expansion enclosure became inaccessible though disks in the controller enclosure remained accessible.
- Fixed an issue with ADAPT spare capacity after disk replacement.
- Fixed an issue with replication in a single-controller configuration.
- Fixed an issue where the customer could not repurpose a foreign FDE-capable disk in a secured system.
- Fixed an issue where SED disks went into an Unusable/Leftover state after the system was power-cycled.
- Fixed an issue where FDE disks were not detected when a system was power-cycled.
- Fixed an issue described as "A virtual pool was offline because a metadata volume is unreadable or missing."
- Fixed an issue related to upgrading the GEM version in 5U84 enclosures having Lattice sideplanes.
- Fixed an issue that caused the system to display an FDE protocol error.
- Fixed an issue where missing VPD values caused an improper health warning to appear for 764W DC PCMs.
- Fixed an issue where storage performance was degraded and the controller restarted automatically.
- Fixed an issue where a dual-controller crash occurred while performing a rebalance operation.
- Fixed an issue where physical PHY mapping was not enabled for Lattice sideplanes.
- Fixed an issue where firmware update failed for a system with a mixed Lattice and Intel IOM

configuration.

- Fixed an issue where a misleading health message led to the wrong part being replaced. For health reason "A midplane SGPIO bus failure was detected" the health recommendation has changed to "Contact technical support."
- Fixed an issue where the disk count shown for a disk group was wrong after replacement of a failed disk.
- Fixed a VPD issue that prevented zoning for 2U enclosures using certain midpanes.
- Fixed an issue where both virtual pools were offline with health reason "The virtual pool is too large."
- Fixed an issue where, after disk replacement and disk-group rebuild and rebalance, system performance was degraded.
- Fixed an issue where, after changing the password of a default user, the old password still worked.
- Fixed an issue where attempting to update 5U84 JBOD IOM firmware failed due to a firmware mismatch.
- Fixed an issue where updating to UUT 6.23 on a 5U84 system changed PHY mappings.
- Fixed an issue where a disk group remained in a degraded state after replacement of a faulty disk.
- Fixed an issue where a system using mixed-size disks in an ADAPT disk group had a failure that left multiple disks in a leftover state.
- Fixed an issue where, after disk firmware updates, an ADAPT disk group became quarantined.
- Fixed an issue recorded as "Unexpected HW power state or mismatch between EBOD and Midplane CPLD power states for drive" in controller logs that caused a disk group to be quarantined.

Features or Enhancements Introduced in GN280R011-01/GT280R011-01

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.

- The new Lattice CPLD is not supported by firmware version GN280R010-01 or earlier. Attempting to use that firmware with the new hardware will result in unexpected behavior. The system must be updated to firmware version GN280R011-01.
- If the firmware update fails with the following code load error: `SP: Error: Fatal`, contact

your Seagate account manager for a process to resolve it.

- . 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 error `Failed to push bundle to partner 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

Issues fixed in GN280R011-01/GT280R011-01

- . Fixed issue where SNMP service stopped responding.
- . Fixed issue where SMI-S API broke after firmware upgrade to GT280R010-01.
- . Added Lattice CPLD support.

Features or Enhancements Introduced in GN280R010-01/GT280R010-01

- . Added support to the new Ecodesign 2.2kW Titanium PSUs for use in 5U84 enclosures.

Issues fixed in GN280R010-01/GT280R010-01

- . Fixed `Overlapped command Check condition`.
- . Fixed an issue of `Historical Performance Statistics` in Storage Management Console or Command Line Interface.
- . Fixed a controller crash issue when global spare was assigned.
- . 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 partner controller due to heartbeat loss.
- Fixed an issue where a data-unavailability situation occurred due to rebooting of a controller.
- Fixed incorrect status reported in SupportAssist logs and upgrade notifications.
- In the Storage Management Console, fixed inconsistencies between English and non-English languages in the confirmation dialog box when modifying a volume.
- Fixed an issue where a root expander experienced a communication fault.
- Fixed an issue where a PCIe crash exceeded an I/O synchronization timeout.
- Fixed an issue where a controller crashed in the NOCP area.
- Fixed an issue where a pool went offline after an ADAPT disk group became quarantined.
- Fixed an issue where sensitive information was shown in two system logs.
- Fixed an issue where a controller crashed due to Double Deletion of Host IO in ATS abort path.
- Fixed an issue where an incorrect string was shown for a temperature alert event on an enclosure.
- Fixed an issue where customer failed to configure an email account using smtp.office365.com.
- Fixed an issue where a user was unable to login to Storage Management Console if the username or password contains special characters like £.
- Fixed an issue where firmware upgrade failed on storage systems until the MC restarted.
- Fixed an issue where the CLI commands `show license` and `query peer-connection` showed inconsistent license information.
- Fixed an issue where the Management Controller failed to be ready.
- Fixed an issue where UNMAP feature for HDD drive incorrectly enabled during enclosure renumbering.
- Fixed an issue where Engineering lab storage systems failed to work using DHCP.
- Fixed an issue where a controller crashed while multiple copyback operation was running on the same disk group.
- Fixed an issue where disks were becoming degraded during firmware update.
- Fixed an issue where incompatible IOM firmware was detected on IOM A after a firmware upgrade.

Features or Enhancements Introduced in GN280R009-02/GT280R009-02

- . Enclosure Management port 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 Seagate Lyve Pilot agent (Unified Data Service).
- . Added Unified Data Service commands `set protocol` and `show protocol`.
- . 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.

Issues fixed in GN280R009-02/GT280R009-02

- . 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` on the system.
- . Fixed an issue where clear expander-status fails to clear expander-status for 5U84 enclosure.
- . 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 failed during a bad-block recovery.
- . Fixed an issue of offline disk groups remaining quarantined even after the controller was live.
- . Implemented **Salvage** mode for field recovery red-alerts.
- . Fixed issue of reading incorrect enclosure configuration page and swapping of incorrect enclosure configurations, 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 GN280R002-03/ GT280R002-03.
- . 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 enclosure 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.
- Seagate 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 page fault crash in Inband Management Read/Write path.
- Fixed an issue where after upgrading to GN280R002-03/GT280R002-03, system is asking for 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 GEM CLI `ddump_pwrmgr` to show correct power cooling module /power supply 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.
- Security enhancements implemented.

Issues fixed in GN280R008-04/GT280R008-04

- Enabled retrieval of drive logs from Seagate drives.
- Fixed invalid string references on the Install License screen.
- Fixed the 24-hour start time format conversion.
- Enabled retrieval of drive logs from Seagate drives in Seagate storage systems.
- Security enhancements implemented.
- Increased default media scrub interval from 24 hours to 360 hours.
- Fixed failure in ADAPT disk movement in the chassis.
- Updated warning and critical temp thresholds.
- Ambient temperature sensor readings corrected on 2U12/2U24/4U24.
- Updated midplane CPLD low density to 15V.
- Enabled meter scrub duration.
- Supported single step upgrade from older release to latest G280 release.
- Upgraded OpenSSH to openssh8.2p1 version.
- Updated the USB core driver.
- Removed Dequarantine_Vdisk and Trust_Vdisk from Jenkins Silver list.
- Fixed read cache from running out of CEs for internal reads.
- Initiated invocation of nvdevice UID check for RSCU case.
- Added clearing of corresponding NV UID quarantine bit mask when container number is freed.
- Fixed incorrect reporting of `All global spare deleted` event when global spare still available.
- Changed severity of event 608 (`backend miscabled`) from `Informational` to `Error`.
- Updated logic to retrieve volume groups by name.
- Modified `Validate_Sysinfo` to skip `lastANotified` and `lastBNotified`.

- Corrected handling of disk group when associated disk is pulled.
- Corrected `Mappings Verification` text box.
- Fixed `Create Initiator` issue on SAS controllers.
- Removed unnecessary error prints from lone zone container.
- Added a confirmation message for the shutdown command.
- Improved `reset all-statistics` for the CLI handler.
- Fixed SEEPROM read from FRU event generation.
- Prevented unwanted removal of in-use disk group.
- Fixed page fault during `restore defaults factory` testing with SAS host interface.
- Improved the DMA-to-cache transfer function.
- Fixed shutdown stall when flushing unwritten cache data to a quarantined disk group.
- Fixed non-PI errors and prevented data corruptions.
- Fixed sideplane firmware downgrade issue for 5U84 enclosures.
- Set default storage type as paged for supported virtual disk groups.
- Fixed timing issue in SAFTE initialization.
- Implemented proper identification strings in disk reports.
- Fixed improper RAID-10 and RAID-50 preemptive reconstruct.
- Fixed controller hangs during ADAPT preemptive reconstruct.
- Added check for cache and page storage initialization before starting ADAPT.
- Modified code to avoid continual rebuilding.
- Prevented unnecessary update in missing drive mask function when drives are missing.
- Fixed shutdown stall when flushing unwritten data when the controllers are shutdown and then rebooted using MC CLI or MUI.
- Fixed hangs in BBR recovery for copyback and fenced data detection.
- Fixed a crash for write I/O operations when the pool becomes full.
- Fixed a crash due to non-availability of memory resources.
- Improved handling of out of NV_CPU_BLPTS function.
- Fixed `Cancel Replication` message crash.
- Changed replication set queue policy from discard to queue-latest if either side's pool is over the high threshold.
- Changed a failed-over system to finish in-progress snapshot operations before attempting a failback.
- Improved the debug logs for better engineering analysis.
- Added support for new CAPI error codes.
- Fixed PSC removal during initialization.
- Resolved compilation error and removed `userArray` variable usage.
- Change made to mark a dead drive with an error without reading it to avoid accessing the

physical failed drive.

- Modified controller shutdown to occur only once by sensor monitoring after detecting the error threshold.
- Fixed in-band management issue.
- Security enhancements implemented.

Issues fixed in GT280R007-03

- Added patch for CVE_2015_5364.
- Added patch for CVE-2015-1465.
- Prevented removal of a Paged Storage Component that is being used as the source of a metadata copy.
- Fixed page fault during restore defaults factory testing with SAS host interface.
- Fixed ATS state machine to return DMA to cache by checking for no Outstanding Worker IO's.
- Added patch for CVE-2016-7117.
- Added patch for CVE 2016-10229.
- Added patch for CVE 2014-2523.
- Fixed an issue of shutdown stall due to flushing of dirty data to a quarantined disk group.
- Handle non-PI WIO errors by retrying the WIO and Prevent Data Corruptions by propagating WIO error status into any associated coalesced WIOs.
- Sideplane firmware downgrade issue.
- For configurations where only paged disk-groups are supported, set the defaultarray type as paged.
- Fixed a timing issue in SAFTE init.
- Fixed the drives report `Degraded health (UNUSABLE)` after changing FDE state from `Unsecured` to `Secure`.
- Fixed RAID-10 and RAID-50 preemptive reconstruct if multiple drives are degraded.
- Fixed controller hangs during ADAPT preemptive reconstruct.
- Added a check for `Cache and Page Storage` initialized before starting ADAPT.
- Changed the code to not rebuild continually.
- When SC receives a manual dequarantine command, it needs to check the `missingDriveMask` of the container. If missing mask is non zero, then do not set a partition as dead and fail the manual dequarantine command. This will avoid an unnecessary update in the dead map when drives are missing.
- Initialize the drive down list earlier, so that down drives are correctly considered during the RAID initialization.
- When the issue of shutdown stall due to flushing of dirty data to a quarantined disk group is fixed (TT-51665/FMW-33107), avoid loss of dirty data when the controllers are shutdown and then rebooted using MC CLI or MUI.

- Fix hangs in BBR recovery running copyback and detecting fenced data.
- G280 CI: build#466 & build#499 Test - Schedules_Remain_After_Kill_Unkill.
- Fixed assert in the PS I/O error path when the pool is full.
- Improve handling of out of NV_CPU_BLPTS in Read Tree Actor and in Read BLPT actor.
- PS isBlpteLocked() - move check for TLPTE lock to outside of non-null blpte clause.
- Fixed a crash which was caused when a `Cancel Replication` message is sent to a secondary system that doesn't have the Replication Set.
- Queue policy changed for pool high threshold reached.
- Finish active snapshots before failback.
- Improved NOPM tracing for debug.
- Fixed a NMI kill due to shutdown hang during SDR with replications.
- Added support for new CAPI error codes.
- Added Check pool limits before expanding a PSC.
- Fix for removing a PSC during initialization.
- Resolved compilation error. remove use of userArray variable.
- Fixed a crash due to excessive media errors.

Issues fixed in GT280R007-02

- Set the `default password changed` flag correctly for new user.
- Dual controller crash observed.
- [Codeload] Controller is down after codeload. Fix to only shutdown once when we have a temperature failure.
- Controller A crashed while running I/O on system with snapshot schedules.
- [AR] Controller A crash.
- Fix possible split write I/O volop deadlock.

Issues fixed in GT280R007-01

- Fix issue related to in-band management.
- Chelsio IOLibs changes - Fix a TX queue resource leak on TMF responses.
- Fix memory leak in PHClient.
- Rearranged the order in which commands to enable or disable both `CloudIQ` and `SupportAssist` are executed.
- Added missing UTC time zones in the list of valid parameters of `set support-assist-info` command.
- Increased width of container widget and removed some margins so that contained text on

Advanced Settings screen is visible properly.

- Create Snapshot Scheduler does not convert time specified in PM to 24 hour format correctly.

Known Issues and Workarounds

Issue: Default values in Help section differ from actual output after performing `restore defaults factory`.

Workaround: None.

Issue: While adding a new disk group, user cannot change the disk group's type in the WBI.

Workaround: Use the CLI parameter `type linear|virtual|read-cache` to change a disk group's type while running the `add disk-group` command.

Issue: Firmware downgrade from G280 to G265 fails with the following error:

Firmware update blocked because the bundle contains firmware that does not support the system's storage pools. Attempt to down-grade firmware on controller that doesn't support the system's storage pools.

Workaround: Use the WBI or the CLI to delete the pools created with G280 firmware on the

controller and retry the firmware downgrade. If you are unsure which disk groups were created using G280 firmware, please contact support for assistance.

Issue: The WBI does not support the `Copy Volume` action for linear volumes.

Workaround: You can use the CLI `copy volume` command to copy a linear volume if the destination volume is a virtual volume and is owned by the same controller as the linear volume.

Issue: Controller A crashed during a controller pull/push operation.

Workaround: None. This was a rare case.

Issue: GN280R007-02/GT280R007-02 has no midplane VPD/CPLD and PSU files included in the build.

Workaround: Use the FTP process to upload the midplane VPD/CPLD and PSU files.

Issue: Disk groups went offline as linear disk groups while upgrading from GN280R007-02/GT280R007-02 to GN280R007-03/GT280R007-03.

Workaround: Use the CLI `Trust` command.

Issue: The WBI 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: No workaround necessary.

Issue: MC (partner) is not ready while downgrading to GN275R003-01/GT275R003-01 with PFU disabled.

Workaround: None. Firmware must be upgraded/downgraded per

the code load sheet. **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: While upgrading from GN260R008-03 to GN280R008-09, firmware upgrade fails and the controllers starting crashing with page faults.

Workaround: Replace faulty disk before upgrading from G260 to G280.

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.

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

Workaround: None.

Issue: When using the Storage Management Console 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 the `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 (').

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.

i Note: The disk management utility (**GUI**) and **CLI** allow you to enable or disable **Partner Firmware Update** for the partner controller. To enable or disable the setting via the **GUI**, see the **Configuring Partner Firmware Update** topic within the [QXS 12G Disk Management Utility User Guide](#). To enable or disable the setting using the **CLI**, use the set **advanced-settings** command, and set the **partner-firmware-upgrade** parameter (see the [QXS 12G CLI Reference Guide](#) for more information about the command parameter syntax).

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 [QXS 12G Events Description Guide](#).

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\Current Version\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**:



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/

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.

Contacting Quantum Support


Below is information related to contacting Quantum Support as well as steps to improve your Quantum customer journey.

- [Submit a Ticket \(Service Request\) on the next page](#)
- [Use MyQuantum Service Delivery Platform on the next page](#)
- [Use Cloud Based Analytics \(CBA\) on page 23](#)


- [Escalate a Case on the next page](#)
- [Contact Quantum Sales on the next page](#)

Submit a Ticket (Service Request)

If you need to submit a ticket or speak to Quantum technical support, go to the Support page at <https://www.quantum.com/en/service-support/>



To start the process with Quantum Technical Support, click **Submit a Ticket**. From here, sign in to the MyQuantum Service Delivery Platform or create an account. For more information, refer to the [Use MyQuantum Service Delivery Platform below](#) section below.




Use MyQuantum Service Delivery Platform

MyQuantum is a single portal for everything Quantum. You can view assets, open support cases, receive real-time updates, and search the Knowledge Base and documentation, all through a secure, online

portal.

1. Create an account and log in to the [MyQuantum Service Delivery Platform](#).
2. Register a product on [MyQuantum](#).



3. Request site access to the Cloud-Based Analytics (CBA) monitoring portal and follow the instructions to set up product(s) to connect to CBA. You can use CBA to monitor Quantum products remotely, from a single dashboard, and Quantum Support can use it to help troubleshoot products more efficiently.

Refer to product documentation for product-specific information related to CBA.

Use Cloud Based Analytics (CBA)

Quantum products are equipped with a Cloud Based Analytics (CBA) agent that can provide log files and snapshots to Quantum CBA servers that are running in the cloud.

CBA enables Quantum systems to collect data regarding system and environment performance. The collected data is bundled and uploaded to the remote CBA server for analysis. You can access Quantum system performance and health results on the CBA dashboard (at <https://insight.quantum.com>) or through the MyQuantum Service Delivery Platform.

The CBA dashboard displays the analytic results of the uploaded CBA data using flexible charting tools, along with an overall health score of each Quantum system configured for the CBA account.

Escalate a Case

To escalate a case, follow the process documented here: <https://www.quantum.com/en/service-support/resources/escalation/>

Contact Quantum Sales

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

