



Seagate 5005/4005/3005 Series Release Notes for Firmware Version GN265R009-03/GT265R009-03

Part Number: 83-00007498-10-01

Version: A

Published: February 2019

© 2019 Seagate Technology LLC or its affiliates.

The information contained herein is subject to change without notice. The only warranties for Seagate products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Seagate shall not be liable for technical or editorial errors or omissions contained herein.

Acknowledgments

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

Description

This package delivers firmware for Seagate 5005/4005/3005 Series controller enclosures and expansion enclosures. Note that expansion-enclosure firmware is embedded in GN265R009-03/GT265R009-03 firmware bundles and is updated when attached to one of the following controller enclosures.

Model	Firmware version
5005	GN265R009-03
4005	GN265R009-03
3005	GN265R009-03

Update recommendation

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

Operating systems

Supported operating systems include:

- Microsoft Windows Server 2016
Microsoft Windows Server 2012 R2
Microsoft Windows Server 2019
- Red Hat Enterprise Linux 6 Update 9
Red Hat Enterprise Linux 7 Update 4
- SuSE Linux Enterprise Server 12 SP2
- VMware ESXi 6.0 U3
VMware ESXi 6.5

Installation instructions

To install this firmware by using either the Storage Management Console or the FTP/SFTP interface, see the Seagate 5005/4005/3005 Series Storage Management Guide.

Issues fixed in GN265009-03/GT265R009-03

- Occasionally, log messages on an expander will show repeated Nexus timeout messages. This was remedied by correcting an internal storage-system communication issue where faults can be reported against the array device, expander, and temperature sensor elements associated with the affected sideplanes. The data path itself remains unaffected.
- Fixed an issue in 2U enclosures where certain internal operations caused the IOM(s) to unexpectedly restart, leading to data unavailability.
- Fixed an issue in 2U enclosures where in some cases, the ambient temperature sensor inaccurately reported a high temperature causing the Ops panel to indicate a fault. If your storage system has experienced this issue, use this procedure to perform a firmware update:
 - Using the CLI:
 - Run: **check firmware-upgrade-health** If the only failure condition relates to temperature sensors then proceed to the next step. Otherwise, resolve any other failure conditions and then repeat this step.
 - Run: **set advanced-settings partner-firmware-update disable**
 - Using SFTP or FTP:
 - On controller A, run: **put filename.bin flash:force**
 - Verify that controller A firmware update completes successfully.
 - On controller B, run: **put filename.bin flash:force**
 - Verify that controller B firmware update completes successfully.
 - Verify that the two controllers are communicating.
 - Using the CLI:
 - Run: **set advanced-settings partner-firmware-update enable**
- Corrected an issue where a controller module's expansion port status LED remained lit (Green) even when SAS cable was not inserted.
- Fixed an issue where in a dual-controller system, after replacing a controller module, PFU failed repeatedly with the message "PFU was unable to build a package, UNKNOWN".
- [JBOD] Fixed an issue related to queue-depth commands where an overlapped command Check Condition was seen with no overlapped command.

Issues fixed in GN265008-04/GT265R008-04

- Fixed a problem with the slot-affinity feature for ADAPT disk groups.
- Fixed an issue where disks and enclosures were not reported in numerical order.
- Removed the unsupported "Verify" operation for ADAPT disk groups.
- Improved performance when using an ADAPT disk group with semi-sequential workloads and other ADAPT improvements.
- Fixed a data draining bug.
- Fixed a condition in an unique use case where WBI access was getting lost.
- Improved messaging when adding a 5U84 expansion enclosure.
- Improved the message and recommended action for event 565. Because this relates to a transient hardware problem, the recommended action is to restart the controller module that logged the event, and only if the problem persists to then replace that controller module.
- Corrected a page fault that could occur on a system running heavy I/O.

Issues fixed in GN265R007/GT265R007

- Fixed an issue where invalid HTTP headers caused the Management Controller Web server to crash and occasionally hang.
- Fixed an issue where the same host mapping was shown twice in volume-mapping output.
- Corrected a code load failure caused by duplicate enclosure entries. The code fix helps to eliminate duplicate enclosure entries and eliminate the need for manual rescan to get correct ordering of enclosures.
- Corrected a controller crash caused from host I/Os greater than 1 MB.
- Corrected a code load issue occurred when more than 9 enclosures were connected together.

- Enhanced EasyStart by changing the default SSD RAID level to RAID 1 or RAID 10 depending on number of SSDs.
- Corrected SNMP default to "disabled" after a **restore defaults** CLI command is issued.

Issue fixed in GN265R006/GT265R006

- Fixed an issue where explicit mapping information was not displayed in the Storage Management Console while upgrading firmware, which led the user to believe maps are no longer valid. If the user remapped the volumes and did so incorrectly, it could cause data unavailability.

Issue fixed in GN265R004/GT265R004

- Corrected the blank 7-segment display on the 2U12 controller enclosure.

Features or enhancements introduced in GN265R003/GT265R003

- Added the ADAPT data protection level for linear and virtual disk groups. This feature maximizes flexibility, provides built-in spare capacity, and allows for very fast rebuilds, large storage pools, and simplified expansion.
- Added support for 5U84 controller and expansion enclosures (a 5U enclosure containing up to 84 disks). A storage system can have up to three 5U84 expansion enclosures.
- Added a JSON API output option to the CLI.
- Increased the maximum virtual pool size per 5005 or 4005 controller module to 1 PiB (pebibyte), with the large-pools feature enabled.
- In the Storage Management Console, introduced a guided-setup interface to speed initial configuration and provisioning of the storage system.
- Enabled DHCP as the default IP addressing mode for controller module network ports.
- Disabled FTP and enabled Secure FTP (SFTP) by default.
- Re-enabled support for SMI-S for storage systems with 2U enclosures only.
- Re-enabled support Microsoft VSS 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.
- Increased performance to 600K IOPS for 5005 systems and to 300K IOPS for 4005 systems.
- For virtual storage, increased performance to 600K IOPS for 5005 systems and to 300K IOPS for 4005 systems.
- Added support for OpenStack with Active/Active high-availability.
- Added an option to enable/disable CHAP in the Storage Management Console Configure CHAP panel.
- For use with X-Clarity, top-level-assembly serial number and part number properties are accessible via SLP, the CLI, and the Storage Management Console.
- Added new virtual-pool statistics to CLI API output. These statistics include page allocations/deallocations, page maps/unmaps, hot/cold page moves, and blocked SSD promotions.

Issues fixed in GN265R003/GT265R003

- Fixed strings that were not translated in the Storage Management Console, CLI, and event messages.
- Fixed an issue where, in the Storage Management Console 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 Storage Management Console.
- 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 script failed with error `XML Parse Failed`.
- Fixed an issue where the Storage Management Console 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 Storage Management Console's Configuration Wizard 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 in the Storage Management Console's Home topic.
- 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 Storage Management Console, the old version of code was still shown.
- Changed the firmware so that if an unsupported (OneStor) expansion enclosure is attached to the system, the enclosure 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 Storage Management Console. 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, in the Storage Management Console's System topic, 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, in the Storage Management Console's Replications topic, hovering on a replication set name did not display replication progress.
- Fixed an issue where both controllers crashed with the following error while running I/O with a power cycle of one controller:

```
unilog set CMOS crashed bit in CriticalErr Controller Crashed - Entered CriticalErr().*****Double entry to CriticalErr()*****
```
- Fixed an issue where, after loading the same firmware bundle onto a system with PFU enabled, the Storage Management Console 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 Storage Management Console, 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, on a 3005 system, the Management Controller crashed with "MC not ready" and "internal timeout error " while aborting a replication set running staggered schedule.
- 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, in the Storage Management Console's Map panel, using filters caused the Next and Previous arrow buttons to become obscured.
- Fixed an issue where, while pulling logs via FTP or the Storage Management Console, an error message stated that logs can't be pulled because a code load is in progress.
- Fixed an issue where, on a 3005 system, replication could not be aborted.
- Fixed an issue where, on a 3005 system, after creating a replication set for a volume group in the Storage Management Console, a prompt to begin replication could not be dismissed.
- 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'. It remains necessary to wait until the code load is complete before logging in, but the message has been corrected to: Login unavailable. Please try again later.`
- Fixed an issue where in management interfaces, FRU data was not available for a power supply.
- Removed the incorrect information shown by event 362 when the task completed successfully.
- Changed events 39, 40, 307, 469, 476, 477, and 524 to specify that the operating temperature range for a controller enclosure is 5°C–35°C (41°F–95°F) and for an expansion enclosure is 5°C–40°C (41°F–104°F).
- Fixed an issue that prevented using the Storage Management Console to switch from static to DHCP addressing for the management port, in single controller mode.
- The CLI **set host-port-mode** command's **FC-and-SCSI** option is not applicable on 3005 systems and is no longer accessible.
- For IOMs in expansion enclosures, the CLI **show configuration** command no longer shows part number and serial number values, which were duplicated.
- Fixed an issue in which unwritten cache data occurred after code load.
- Fixed an issue in the Storage Management Console that prevented the "Enable managed logs" option from remaining selected.
- Fixed an issue where 9 of 12 disks in an expansion enclosure showed as failed with reason: `Illegal Request, logical block address out of range.`

Known issues and workarounds

Issue: The Storage Management Console panel for installing a license, incorrectly shows a Temporary License tab. The capability to create a temporary license is not supported.

Workaround: None.

Issue: With I/O running when the first ADAPT component is added to a non-ADAPT pool, or when the last ADAPT component is removed from a pool; leaving only non-ADAPT components, a controller may crash.

Workaround:

1. Mitigate this issue by ensuring that the system has no dirty cache before adding the first ADAPT component to the pool as follows.
 - i. Halt/quiesce host I/O.
 - ii. Wait for dirty cache to flush to disk (nominal time is 1 minute; while recommended is 5 minutes).
 - iii. Add ADAPT component to pool.
 2. When initially setting up the pool, always include an ADAPT component.
-

3. Do not remove the last ADAPT component from a pool.

To recover if this crash occurs:

Reseat the controller or restart the controller from the partner controller. The system will have completed adding or removing the pool component

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 partner controller.

Issue: A controller's management port stops working.

Workaround: Restart the affected controller's Management Controller from the partner controller.

Issue: Some strings are not translated in the Storage Management Console 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 Storage Management Console, 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.

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 Storage Management Console, 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 Storage Management Console, 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.

Issue: Cannot delete more than 50 volume groups at a time in the Storage Management Console or the CLI.

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

Issue: In the Storage Management Console, 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" but the volume group is not identified.

Workaround: Perform a rescan to repopulate the internal system with fresh data, and then rerun the command.

Issue: In the Storage Management Console, 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.

Issue: In the Storage Management Console, 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.

Issue: In the Storage Management Console's Manage Schedules panel, the last scheduled task that was deleted is not grayed out.

Workaround: Refresh the browser.

Issue: In the Storage Management Console's User Management panel, 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 Storage Management Console.

Workaround: None.

Issue: When using Microsoft Edge, Storage Management Console help text doesn't appear.

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

Issue: In the Storage Management Console, 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 Storage Management Console, 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 command.

Issue: In the Storage Management Console, when deleting a schedule, a success message will appear but the schedule may remain listed in the Manage Schedules panel. Trying to delete the schedule again will result in 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".

Issue: While creating a peer connection, the new peer connection is not shown in the secondary system's Storage Management Console 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: Customer sees an Expander Controller code-load failure signal and the controller reports that it has an invalid bundle.

Workaround: Reinstall the firmware bundle.

Issue: In a Chrome browser the Storage Management Console 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.

Issue: Verify operation returns false success message and success event for ADAPT disk group even though 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 downlevel 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 Storage Management Console, 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 PFU 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.

Issue: Linear RAID-5 expand disk-group operation will begin but will fail with 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 Storage Management Console 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 enclosure.

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 over-temperature 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 enclosure first, then the server with the 930-8e HBA.

Effective date

February 2019