



Exos[®] CORVAULT[™]

S100R013-01 Release Notes

Part Number 204630700-00, B • October 2022

Description

This package delivers firmware for the Exos CORVAULT storage enclosure.

Model	Firmware version
6575	S100R013-01

Update recommendation

This is a recommended firmware update.

Operating systems

Supported operating systems include:

- Microsoft Windows Server 2019
Microsoft Windows Server 2016 Hyper V
Microsoft Windows Server 2016
- Red Hat Enterprise Linux 8.0
Red Hat Enterprise Linux 7.6
- SuSE Linux Enterprise Server 15
SuSE Linux Enterprise Server 12 SP2
- VMware ESXi Server 6.7 U1
VMware ESXi Server 6.5 U2

Installation instructions

To install this firmware by using either the Storage Management Console or the SFTP/FTP interface, see the *Seagate Exos CORVAULT Storage Management Guide*.

New features or enhancements in S100R013-01

- Added support to increase the number of drive failures which are eligible for Autonomous Drive Regeneration (ADR) in CORVAULT storage systems. The new drive firmware must be installed before the improved ADR support is enabled. For

Exos X18 SAS SED 512E 18T HDDs, the required drive firmware is ETA4. For Exos X20 SAS SED 512E 20T HDDs, the required drive firmware is ETA3..

Issues fixed in S100R013-01

- Fixed an issue where the system did not report a direct-attach disk is single ported (that is, not accessible to both controller modules) and did not adjust the system health accordingly.
- Fixed an issue where if the controllers restarted during daylight savings time, then the controllers needed a short duration to sync to system time or NTP (network time protocol).
- Fixed an issue when both the controllers were replaced simultaneously, some disks went into a leftover state.
- Fixed an issue where during reconstruction of a disk group, the LED Status property shown by the `show disks details` command displayed Online instead of Rebuild.
- Fixed an issue where when one of the controllers was down (logical shutdown), then the enclosure's Fault-Application LED located on the lower left corner of the enclosure front panel illuminated as solid amber to show the logic fault.
- Fixed an issue where aggregate health was Degraded or Bad, without any unhealthy components when the CLI `show system` command was run.
- Fixed an issue where the controller lost communication with the host through the in-band interface during a partner firmware update (PFU).
- Fixed an issue where during a failover or failback event/error recovery, the controller hangs and does not automatically recover.
- Fixed an issue where when a system fan module experiences a failure, its fault LED does not illuminate as solid amber.
- Fixed an issue where when a controller module was removed while its Controller ID LED was blinking, upon insertion the LED stopped blinking.
- Fixed an issue where when the controller modules were being swapped or when power cycling the enclosure, the system fan modules LEDs blinked and there was an inconsistent front panel LED operation.
- Fixed an issue where there was no provision in CLI to set identification LED for fans.
- Fixed an issue where the controller showed incorrect firmware status after a connection failure. There was an Ethernet connectivity issue during the upload which left the firmware updated activity progress in an 'in progress' state and not in a 'timed out' state.
- Fixed an issue where the Management Controller does not reboot after a GEM update.

New features or enhancements in S100R011-01

- None

Issues fixed in S100R011-01

- Fixed an issue where after removing the controller module, the auxiliary bay fault LED illuminated instead of the rear side fault LED.
- Fixed an over-scrubbing issue.
- Fixed an issue where a direct attach, single-ported drive was not marked as single-ported.

New features or enhancements in S100R010-01

- Implemented the ability to detect the flash chip variant and to automatically select the appropriate bus timing parameters for the chip. This ability will enable handling of alternative flash memory parts without slowing down the memory timing for the existing parts.

Issues fixed in S100R010-01

- Increased flash mapping memory to accommodate larger CPLD images.
- Fixed an issue where firmware build failed due to incorrect bit bucket handling while rebooting the partner controller.

Known issues and workarounds

<p>Issue: A few event messages include a drawer ID, which is not applicable.</p> <p>Workaround: Ignore the drawer information reported in the event messages.</p>
<p>Issue: Redundant response messages are displayed for the CLI <code>map volume</code> command.</p> <p>Workaround: Ignore the redundant messages.</p>
<p>Issue: When running the CLI <code>set password</code> command, pressing Tab to see the command options displays "<tab>" in the command line instead of showing the options.</p> <p>Workaround: Press Backspace to remove the "<tab>" instances. To see what the command options are, view the command's help.</p>
<p>Issue: Build date does not include year.</p> <p>Workaround: No workaround required. The system build year is shown as part of the bundle build date on the Maintenance > Firmware > System page. It can be also be seen by running the CLI <code>show versions</code> command.</p>
<p>Issue: In the Storage Management Console, Provisioning > Hosts > Create Host, shows the first Initiator ID Nickname field as empty. The value does not appear until you select that initiator ID.</p> <p>Workaround: The nickname value appears after the initiator ID is selected.</p>
<p>Issue: For volume mappings in the CLI, if the <code>access</code> parameter is set to <code>no-access</code>, the CLI <code>show maps</code> command does not show all the volumes.</p> <p>Workaround: Do not map volumes using the 'no-access' option.</p>
<p>Issue: Slow disk performance resulting in a cache I/O stall with an event similar to: "WARNING Killed partner controller. The reason for this is partner I/O stall in cache."</p> <p>Workaround: Replace the disk that is slow and having timeout problems with a functional/compatible disk.</p>
<p>Issue: The CLI <code>show enclosures</code> command shows bus ID and target ID properties that are not applicable to this product.</p> <p>Workaround: Ignore bus ID and target ID fields.</p>
<p>Issue: Controller identification LEDs are not working properly in the Storage Management Console. When identification LEDs are turned on for both controllers and one controller is removed and reinserted, the Storage Management Console continues to show the LED as turned on even though the physical LED is turned off.</p> <p>Workaround: Turn LED off and back on.</p>
<p>Issue: Controller crashes on code load when partner firmware update (PFU) is disabled.</p> <p>Workaround: Use <code>unfail controller</code> CLI command to unkill failed partner controller.</p>
<p>Issue: For an ADAPT disk group, any volume size less than 9GB reports volume too small.</p> <p>Workaround: Use a larger volume size.</p>
<p>Issue: The CLI <code>show disk-statistics</code> command returns "Error: An invalid device was specified".</p> <p>Workaround: None.</p>

<p>Issue: If controllers A and B are connected to different networks, the fully qualified domain name (FQDN) of the partner controller (controller B) is not displayed after setting the hostname.</p> <p>Workaround: Check the partner controller's FQDN from the partner controller's Storage Management Console or CLI.</p>
<p>Issue: The controller can be reached by all protocols using the previous DNS settings, even after the settings are changed.</p> <p>Workaround: Do not manually set DNS settings while running DHCP. Let DHCP automatically set settings for the user.</p>
<p>Issue: Not able to verify the mini-SAS HD port's fault LED on the controller module, as the status is not propagated to illuminate the Fault-Rear Panel LED located with the enclosure front panel's status LEDs.</p> <p>Workaround: Turn on the identified LED of the faulted system.</p>
<p>Issue: Shift+click does not work for check boxes in the Firefox browser.</p> <p>Workaround: Individually left-click on check boxes to select them.</p>
<p>Issue: Amber fault LED on enclosure top lid, does not glow on the faceplate when the two expanders are removed. No fault is generated in Generic Enclosure Management (GEM).</p> <p>Workaround: None.</p>
<p>Issue: After restart, the platform takes 4 minutes to synchronize which is longer than the expected time of 3,5 minutes.</p> <p>Workaround: Wait for 2 more minutes for the post restart synchronization.</p>
<p>Issue: A controller halts during the restart and does not respond.</p> <p>Workaround: Pull and reseal the controller.</p>
<p>Issue: When you use the Safari browser to log in to the Storage Management Console, you cannot establish a stable connection with the CORVAULT system. A warning message says "Communication with the system has been lost. Connection will be automatically restored once communication with the system has been established."</p> <p>Workaround: Use supported browsers such as Google Chrome, Mozilla Firefox, or Microsoft Internet Explorer. Supported versions are listed in the Storage Management Guide and online help.</p>
<p>Issue: After setting up the system successfully, when the user logs in again, the Storage Management Console shows the onboarding wizard instead of the dashboard.</p> <p>Workaround: Click Continue to Dashboard to proceed to the Storage Management Console.</p>
<p>Issue: Commands to clear FDE keys takes more time to execute than usual, and the Management Controller does not respond for small intervals of time.</p> <p>Workaround: Wait for 5 minutes and run the command again.</p>
<p>Issue: A controller kills its partner controller due to reason: Partner I/O stall in cache.</p> <p>Workaround: Unkill the failed controller by running the <code>restart sc<controller-name></code> CLI command.</p>
<p>Issue: After a reboot, a controller rejects enclosure related commands due to the enclosure not being found. The Storage Controller returns error: "CAPI_ERROR_WWN_NOT_FOUND (81)."</p> <p>Workaround: Wait for 5 minutes and run the command again.</p>
<p>Issue: A controller becomes unresponsive and appears to be down even after being unkilld.</p> <p>Workaround: Pull and reseal the controller module to recover the system.</p>
<p>Issue: The CLI <code>show alerts</code> command returns either of the following alerts: "The temperature is below the critical low temperature value." or "The temperature is below the normal operating range."</p> <p>Workaround: Consider this alert to be erroneous when it occurs shortly (that is, within 5 minutes) after power cycling a system. Acknowledge the alerts to clear them. If the alert reoccurs within next 10 minutes, consider it a real alert and follow the normal recommended action for this alert.</p>
<p>Issue: After a server is rebooted, controller IP addresses are not accessible via in-band protocol.</p> <p>Workaround: Disable DHCPv6 on both controllers.</p>
<p>Issue: After a shut down restart, the CLI/WBI may go into a not communicating state for one of the controllers. An error <code>Unable to obtain FRU ID data for the FRU</code> is reported.</p> <p>Workaround: Power cycle the system to recover from this condition.</p>

Issue: Some systems experience transient CRC errors on their host link. This seems to vary in frequency and severity across systems and across time. In general it has no impact on system I/O or throughput as the lower level communication retries and the retries succeed. This issue is visible in the logs on the host. However on Linux hosts when running a ZFS file system on the storage, ZFS is overly sensitive. If the issue is significant, ZFS will drop the storage. Other file systems ignore this low level issue and do not have problems. Seagate has been working with the open source community to update ZFS to ignore this error similar to other file systems (again since low level errors and retries in communication systems are not uncommon). Seagate can provide a ZFS patch until this update has propagated to the open source repository for ZFS.

Below are some details for detecting this issue and the logs that you will see.

Log details:

The event itself first appears in `dmesg/kernel` message traffic on the host. It can be detected by running the following commands:

```
journalctl -k | grep log_info(0x31120303)
or
dmesg -H | grep log_info(0x31120303)
[+0.000009]mpt3sas_cm2: log_info(0x31120303): originator(PL), code(0x12), sub_code(0x0303)
[+0.000004]mpt3sas_cm2: log_info(0x31120303): originator(PL), code(0x12), sub_code(0x0303)
[+0.000004]mpt3sas_cm2: log_info(0x31120303): originator(PL), code(0x12), sub_code(0x0303)
[+0.000007]mpt3sas_cm2: log_info(0x31120303): originator(PL), code(0x12), sub_code(0x0303)
[+0.000007]mpt3sas_cm2: log_info(0x31120303): originator(PL), code(0x12), sub_code(0x0303)
[+0.000005]mpt3sas_cm2: log_info(0x31120303): originator(PL), code(0x12), sub_code(0x0303)
[+0.000004]mpt3sas_cm2: log_info(0x31120303): originator(PL), code(0x12), sub_code(0x0303)
```

NOTE No output from either command implies that the event is not occurring. Output as reflected above means that the event is occurring.

If running ZFS, the `zpool status` command can be used to further confirm the potential presence of the issue. Note the `mpt3sas log_info` events are interpreted by the OpenZFS ZED daemon as write errors, even though they are eventually successful.

```
zpool status
pool: cvt_klara_3
state: ONLINE
status: One or more devices has experienced an unrecoverable error. An attempt was made to correct the error. Applications are unaffected.
action: Determine if the device needs to be replaced, and clear the errors using zpool clear or replace the device with zpool replace.
see: https://openzfs.github.io/openzfs-docs/msg/ZFS-8000-9P
```

Message ID: ZFS-8000-9P — OpenZFS documentation - GitHub Pages

```
# zpool status -x pool: test state: DEGRADED status: One or more devices is currently being replaced. The pool may not be providing the necessary level of replication.
action: Wait for the resilvering operation to complete scrub: resilver in progress, 0.14% done, 0h0m to go
config: NAME STATE READ WRITE CKSUM
        ONLINE 0 0 3 ...
```

`openzfs.github.io`
`config:`

NAME	STATE	READ	WRITE	CKSUM
cvt_klara_3	ONLINE	0	0	0
raidz2-0	ONLINE	0	0	0
wwn-0x600c0ff000512d8225f232630100000	ONLINE	0	0	0
wwn-0x600c0ff000512e8013f232630100000	ONLINE	0	0	0
wwn-0x600c0ff000512e86300033630100000	ONLINE	0	538	0
wwn-0x600c0ff00052075b1e0033630100000	ONLINE	0	0	0
wwn-0x600c0ff000f48b5a10f232630100000	ONLINE	0	0	0
wwn-0x600c0ff000f48b63fff132630100000	ONLINE	0	0	0

<p>Workaround: Apply ZFS patch or upgrade to the newest ZFS open source version. The patch is available in OpenZFS 2.2 and above.</p>
<p>Issue: When a controller recovers from being killed or shut down it may take up to five minutes to become available.</p> <p>Workaround: Wait for five minutes for the MCs to self-recover and synchronize.</p>
<p>Issue: Intermittently, the system fan module LED does not illuminate when identified using the Fan (0-3) Information > Turn on Locator LED action in the WBI (Maintenance > Hardware > Rear View > Fan (0-3) Module) or <code>set led fan-module <0-3></code> CLI command.</p> <p>Workaround: See the <i>Exos® CORVAULT™ Hardware Installation and Maintenance Guide</i> or <i>Storage Management Guide</i> (or help: Maintenance > Hardware) for the location of system fan modules in the enclosure rear panel.</p>
<p>Issue: WBI will sometimes go into not ready state.</p> <p>Workaround: Change the Timeout setting from 720 minutes to default 30 minutes on the Settings > Users > Local Users page.</p>
<p>Issue: During on boarding on the Notification: Syslog page, you cannot proceed further by clicking the Continue button.</p> <p>Workaround: After you click the Send Test Syslog Notification checkbox, the Apply and Continue button gets enabled. You can continue to navigate after clicking the Apply and Continue button.</p>
<p>Issue: When the power cable is disconnected between the PDU and PSU, the Fault-Rear panel LED on the enclosure front panel illuminates as solid amber. If the PSU FRU is removed, the LED illuminates as blinking amber. This actual LED behavior is reversed relative to the functional specification.</p> <p>Workaround: None.</p>
<p>Issue: In the unlikely event that two heads fail on a single drive, the system chooses not to perform ADR on that drive.</p> <p>Workaround: Replace the drive.</p>
<p>Issue: Enclosure and slot numbers are not correct for the events generated for a drive in the system.</p> <p>Workaround: To determine the correct drive number the formula is: (enclosure number * 24) - 1 + slot number = drive number. For example, for enclosure 4 and slot 6 the drive number would be: (4 * 24) - 1 + 6 = drive 101. Refer <i>Exos CORVAULT Hardware Installation and Maintenance Guide</i> for more information on the disk locations.</p>
<p>Issue: After replacing a controller module in a system, the health status is shown as "in progress" when you run the <code>show firmware-update-status</code> CLI command.</p> <p>Workaround: Reload the code to clear the in progress status.</p>
<p>Issue: When you upgrade CORVAULT code to S100R013, one or two expanders may not update to the latest version.</p> <p>Workaround: Revert to S100R011 or an earlier version of code and reload S100R013 to update all the expanders.</p>
<p>Issue: You may see an event that states a drive is in reman but when you run the commands <code>show disks</code> and <code>show disks details</code> the reman status is delayed.</p> <p>Workaround: Check event status for reman state.</p>
<p>Issue: Autonomous Drive Regeneration (ADR) is not supported on drives secured with Full Disk Encryption (FDE).</p> <p>Workaround: If you want to use the ADR feature, do not use drives secured with FDE.</p>
<p>Issue: PFU Codeload and CPLD update codeload fails.</p> <p>Workaround: Retry the codeload.</p>

© 2022 Seagate Technology LLC or its affiliates. All rights reserved. Seagate, Seagate Technology, and the Spiral logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. CORVAULT and Exos are either trademarks or registered trademarks of Seagate Technology LLC or one of its affiliated companies in the United States and/or other countries. All other trademarks or registered trademarks are the property of their respective owners. When referring to disk capacity, one gigabyte (GB) equals one billion bytes, one terabyte (TB) equals one trillion bytes, and one petabyte (PB) equals one thousand terabytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Actual data rates may vary depending on operating environment and other factors, such as chosen interface and disk capacity. The export or re-export of Seagate hardware or software is regulated by the U.S. Department of Commerce, Bureau of Industry and Security (for more information, visit www.bis.doc.gov), and may be controlled for export, import and use in other countries. All coded instruction and program statements contained herein remain copyrighted works and confidential proprietary and trade secret information of Seagate Technology LLC or its affiliates. Any use, derivation, disassembly, reverse engineering, dissemination, reproduction, or any attempt to modify, prepare derivative works, reproduce, distribute, disclose copyrighted material of Seagate Technology LLC, for any reason, in any manner, medium, or form, in whole or in part, if not expressly authorized, is strictly prohibited. Seagate reserves the right to change, without notice, product offerings or specifications.

Regulatory and compliance information

For the latest regulatory and compliance information see www.seagate.com/support. Scroll down to the Compliance, Safety and Disposal Guide link.

Open source third-party licenses and code

Seagate storage products use open source software components. To view information about open source software licenses and open source code used in Seagate storage products, see www.seagate.com/support.