

# SmartRAID 3200 and SmartHBA 2200 Software/Firmware Release Notes



# Table of Contents

- 1. About This Release..... 3
  - 1.1. Release Identification..... 3
  - 1.2. Files Included in this Release..... 3
- 2. What's New?..... 5
  - 2.1. Fixes and Enhancements..... 5
  - 2.2. Limitations..... 19
- 3. Updating the Controller Firmware..... 24
  - 3.1. Updating Controllers to Latest Firmware..... 24
- 4. Revision History..... 25
- Microchip Information..... 26
  - The Microchip Website..... 26
  - Product Change Notification Service..... 26
  - Customer Support..... 26
  - Microchip Devices Code Protection Feature..... 26
  - Legal Notice..... 26
  - Trademarks..... 27
  - Quality Management System..... 28
  - Worldwide Sales and Service..... 29

# 1. About This Release

The release described in this document includes firmware, OS drivers, tools, and host management software for the SmartRAID 3200 and SmartHBA 2200 solutions from Microchip.

## 1.1 Release Identification

The firmware, software, and driver versions for this release are shown in the following table.

**Table 1-1.** Release Summary

<b>Solutions release</b>	3.4.0
<b>Package release date</b>	July 17, 2024
<b>Firmware version</b>	03.01.30.106
<b>UEFI/Legacy BIOS</b>	2.14.4/2.14.2
<b>Driver versions</b>	<p><b>Windows Drivers:</b></p> <ul style="list-style-type: none"> <li>Windows 2022, 2019, Windows 11, 10: 1010.108.0.1015</li> </ul> <p><b>Linux SmartPQI:</b></p> <ul style="list-style-type: none"> <li>Rocky Linux 9: 2.1.30-031</li> <li>RHEL 7/8/9: 2.1.30-031</li> <li>SLES 12/15: 2.1.30-031</li> <li>Ubuntu 20/22: 2.1.30-031</li> <li>Oracle Linux 7/8/9: 2.1.30-031</li> <li>Citrix Xenserver 8: 2.1.30-031</li> <li>Debian 10/11/12: 2.1.30-031</li> </ul> <p><b>VMware:</b></p> <ul style="list-style-type: none"> <li>VMware ESX 7.0/8.0: 4672.0.104</li> </ul> <p><b>FreeBSD:</b></p> <ul style="list-style-type: none"> <li>FreeBSD 14/13: 4540.0.1005</li> </ul>
<b>ARCCONF/maxView</b>	4.18.00.26842
<b>PLDM</b>	6.40.6.0

## 1.2 Files Included in this Release

This section details the files included in this release.

**Table 1-2.** Firmware Files

Component	Description	Pre-Assembly Use	Post-Assembly Use
SmartFWx200.bin	Production-signed programmable NOR Flash File. Use to program NOR Flash for boards that are already running firmware.		X

**Table 1-3.** Firmware Programming Tools

Tool	Description	Executable
ARCCONF	ARCCONF CLI Utility	ARCCONF BXXXXX.zip
maxView	maxView Utility	MAXVIEW XXX BXXXXX.zip

### Driver Files

**Table 1-4.** Windows Drivers

OS	Version
Server 2022, 2019, Windows 11, 10	x64

**Table 1-5.** Linux Drivers

OS	Version
RHEL 9.4, 9.3, 8.10, 8.9, 8.8, 7.9	x64
SLES 12 SP5	x64
SLES 15 SP6 <sup>1</sup> , SP5, SP4	x64
Ubuntu 20.04.6, 20.04.5, 20.04	x64
Ubuntu 24.04, 22.04.4, 22.04.3, 22.04	x64
Oracle Linux 7.9 UEK6U3	x64
Oracle Linux 9.4, 9.3, 8.9, 8.8, UEK7U2	x64
Debian 12.5, 11.9	x64
Fedora 39 (inbox)	x64
Citrix XenServer 8.2.1	x64
Rocky Linux 9.4, 9.3	x64
SLE-Micro 6.0, 6.5 (inbox only)	x64

**Note:**

1. New OS is minimally tested with inbox driver. Full support is expected in the next release.

**Table 1-6.** FreeBSD and VMware Drivers

OS	Version
VMware 8.0 U3 (inbox only) <sup>1</sup>	x64
ESX 8.0 U2/U1, 7.0 U3/U2	x64
FreeBSD 14.0, 13.3	x64

**Note:**

1. New OS is minimally tested with inbox driver. Full support is expected in the next release.

**Host Management Software****Table 1-7.** maxView™ and ARCCONF Utilities

Description	OS	Executable
ARCCONF Command Line Utility	Windows x64 Linux x64 VMware 7.0 and above XenServer UEFI support	See the arccconf_B#####.zip for the installation executables for the relevant OS.
maxView™ Storage Manager	Windows x64 Linux x64 VMware 7.0 and above XenServer	See the maxview_linux_B#####.zip, maxview_win_B#####.zip, and the maxview_vmware_B#####.zip for the installation executables.
maxView™ vSphere Plugin	VMware 7.0 and above	See the maxview_vmware_B#####.zip for the installation executables.
Boot USB (offline or pre-boot) for ARCCONF and maxView Storage Manager	Linux x64	See the maxview_offline_bootusb_B#####.zip for the .iso file.

## 2. What's New?

This section shows what's new in this release.

### 2.1 Fixes and Enhancements

This section shows the fixes and enhancements for this release.

#### 2.1.1 Firmware Fixes

This section shows the firmware fixes and enhancements for this release.

##### 2.1.1.1 Fixes and Enhancements for Firmware Release 03.01.30.106

This release includes the following fixes and enhancements:

- Added RDE Drive Form Factor support for SAS/SATA.
- Added support for saving Managed SED master key and reset key in cipher text in controller NVRAM instead of clear text.
- Changed temperature polling period from 1 minute to 5 minutes for HBA mode drives that are in low power state.
- Added support for clearing UREs in unmapped LBA range of fault-tolerant logical drives.
- Added support for additional IO stats in controller serial logs when a LUN reset occurred.
- Removed SAS PHY downshift from 24G to 12G when expander attached SATA devices are present.
- Fixed a controller lockup from DMA failure after enabling cache with IO ongoing.
  - *Root cause:* IO was put into buffer waiting to be coalesced. When cache is then enabled, this IO tried to do another DMA causing lockup.
  - *Fix:* Send commands that have already been parsed before enabling cache.
  - *Risk:* Low
- Fixed an issue where after a VM performs a soft reset, backup power source status shows "Recharging".
  - *Root cause:* Soft reset cleared variable values.
  - *Fix:* Store backup power source status when issuing soft reset.
  - *Risk:* Low
- Improved performance on read workloads for HDD-based cache-enabled logical drives.
  - *Root cause:* Found that sequential read workloads with multiple simultaneous reads over same LBA range was slow.
  - *Fix:* Optimized code for single stream read performance workloads.
  - *Risk:* Low
- Fixed an issue where foreign SED used as a replacement drive causes Windows Blue Screen Of Death (BSOD).
  - *Root cause:* When a foreign SED replaces a drive in a RAID volume, it can get some read requests that fail because the foreign SED is locked. Firmware's error handling cannot fail the locked SED, which results in LUN Resets and eventually a BSOD.
  - *Fix:* When a foreign SED encounters Aborted Command error, firmware will fail the drive to avoid the LUN Resets and BSOD. The foreign SED can be moved to a free slot to import it later after it is failed.
  - *Risk:* Low
- Fixed an issue where host management software is showing the drive status as optimal when an Otherwise Owned SED drive is inserted as a replacement drive in a degraded logical drive.

- *Root cause:* When the firmware tries to fail the Otherwise Owned SED drive, it fails with the reason code of SED Qualification failed. The firmware prevented the failure of the SED drive for this reason code (SED Qualification fail).
  - *Fix:* Allow the drive to be failed for SED Qualification failed reason code.
  - *Risk:* Low
- Fixed an issue where the import operation was successful for a failed foreign encrypted logical drive.
  - *Root cause:* There was no condition to verify if the logical drive was in a failed state or not.
  - *Fix:* Added a condition to check the state of the logical drive before proceeding to import the logical drive with the master key.
  - *Risk:* Low
- Fixed an issue where foreign drives do not get imported when Managed SED change key is processed during system shutdown.
  - *Root cause:*
    - i. Admin rights are changed on the drive but datastore of the drive is not updated since the drive has failed the commands during system shutdown. During the subsequent boot, the drive is marked as foreign.
    - ii. Some admin rights are changed and some are not since the drives have failed the commands during system shutdown. During the subsequent boot, the drives are marked as foreign.
  - *Fix:*
    - i. During import, when unlock fails with the key provided by the user, use existing master key to unlock the drive.
    - ii. During import, when change pin fails as authentication error with the old key provided by the user, attempt to authenticate with the new key provided by the user.
  - *Risk:* Low
- Fixed an issue where the media exchange status was not updating properly for fault tolerant logical drives with spare drives.
  - *Root cause:* When the spare drives are activated for multiple logical drives across the array, the spare drives can be still activated for the failed physical drives in the failed fault tolerant logical drive. However, the firmware does not consider the spare drives while updating the media exchange status of the respective failed fault tolerant logical drive.
  - *Fix:* Added additional checks to consider the active spare drives in the failed logical drive to update the media exchange status.
  - *Risk:* Low
- Fixed an issue where there is a possible null pointer exception while creating the logical drive.
  - *Root cause:* When a host tries to create a new logical drive, the firmware will try to allocate controller memory to store the logical drive metadata. If the controller's internal memory is completely occupied at this stage, the API will return with a NULL pointer. Firmware fails to safeguard this NULL check, tries accessing the provided address, and ends up with a NULL pointer exception.
  - *Fix:* Added check for memory allocation status to avoid NULL address access.
  - *Risk:* Low
- Fixed an issue where host management software is showing SED encryption status as Foreign SED present when there is no foreign SED on the controller.
  - *Root cause:* When a foreign-managed SED logical drive is detected on the controller, the firmware will update its metadata with the foreign drive information. When this foreign-

- managed SED logical drive is deleted, the firmware fails to clear the metadata of foreign drive information. This results in the wrong reporting of the host management software with foreign SED presence by the firmware.
- *Fix:* Firmware will clear the foreign drive information in its metadata while clearing the foreign-managed SED configuration.
  - *Risk:* Low
- Fixed an issue where rebuild status in host management software is showing as predictive spare rebuild while standard failure rebuild is in progress on Managed SED logical drive.
    - *Root cause:* When a failed SED drive is replaced, the new SED drive state gets updated to WRONG\_REPLACED till the drive is unlocked. During this time, the firmware will continue the rebuild process on the spare drive. The firmware will check if the original drive is a failure (DRIVE\_BAD), then it will update the rebuild status to standard failure rebuild. Otherwise, it will be updated to predictive spare rebuild for all other drive states.
    - *Fix:* Updated the firmware check to include "DRIVE\_BAD" and "WRONG\_REPLACED" for updating the rebuild status to standard failure rebuild.
    - *Risk:* Low
  - Fixed an issue where failed drive was reported as bay 255 in event logs.
    - *Root cause:* Event was being logged before drive bay was discovered correctly.
    - *Fix:* Log event after drive bay is read properly.
    - *Risk:* Low
  - Fixed an issue where the surface scan did not clear fixable Unrecoverable Read Errors (UREs) on a RAID 50 logical drive.
    - *Root cause:* In a RAID 50 logical drive, Unrecoverable Read Errors (UREs) can occur in multiple parity groups on the same blocks. During a surface scan cycle, if the UREs cannot be corrected in the first parity group, the firmware will fail to check for UREs in subsequent parity groups. The firmware is designed to check each stripe individually for UREs and attempts to fix them before moving on to the next parity group's stripe. However, if a READ ERROR occurs on a stripe and cannot be corrected, the firmware did not proceed to check for UREs in the subsequent parity group.
    - *Fix:* During every surface scan cycle, if a URE is encountered that cannot be corrected, the firmware will log an event for the bad block and proceed to check the stripe in the next parity group.
    - *Risk:* Low
  - Fixed a possible operating system hang issue on cache enabled logical drive followed by multiple LUN resets.
    - *Root cause:* An I/O workload on cache enabled logical drive, which can send multiple requests to the same block will be queued to a retry queue if the corresponding block is locked due to an internal cache operation such as cache flush. Once the internal operation is completed, the firmware will process all the pending commands. The number of pending commands is tracked using a BYTE variable in the firmware, which will overflow after queuing 255 entries. Due to this BYTE overflow, firmware will not process the pending commands, resulting in system hang and eventual operating system crash.
    - *Fix:* The pending command counter is moved from BYTE to HWORD.
    - *Risk:* Low
  - Fixed an issue where a BSOD occurs while running I/O on a fault-tolerant managed SED logical drive with a foreign SED present.
    - *Root cause:* When a foreign SED drive is replaced in place of a failed SED in a managed SED logical drive, the I/O targeted to the foreign SED is supposed to be regenerated from

- other SED drives. Due to a missing conditional check, the firmware repeatedly queued the I/O requests to the locked foreign SED drive even after failed completions and never marked the drive as failed, since the drive was in a locked state. This resulted in commands getting timed out and an eventual BSOD error.
- *Fix:* Added a conditional check in the firmware to skip queuing I/O operations to the foreign SED and use a data regenerative method to serve the I/O commands.
  - *Risk:* Low
- Fixed an issue where the controller gets hung when unlocking a Managed SED logical drive.
    - *Root cause:* When unlocking the Managed SED logical drive, the physical drives failed to be unlocked causing the logical drive to be failed. The controller becomes hung while trying to save the RAID metadata to the failed logical drive.
    - *Fix:* Resolved the controller hang by not writing the RAID metadata to the locked SEDs.
    - *Risk:* Low
  - Fixed a possible lockup issue when running I/O on a degraded volume.
    - *Root cause:* When two drives were removed, there was a brief removal and hot add processed, but the firmware event layer did not process the hot removal flag.
    - *Fix:* Firmware will make sure the hot removal flag is cleared.
    - *Risk:* Low
  - Fixed an issue where host management software cannot differentiate between a drive in a rebuilding LUN that has been rebuilt and a drive that was not rebuilt.
    - *Root cause:* Firmware sets the drive rebuild percentage to 100 once a drive has been rebuilt and provides no way to determine if a drive in the LUN has been rebuilt while the rebuild operation is ongoing.
    - *Fix:* Save a list of the drives that have been rebuilt to the RAID metadata.
    - *Risk:* Low
  - Fixed possible lockup issue while hot plugging an enclosure from one connector port to another.
    - *Root cause:* Deadlock scenario occurs where a command is frozen and is being sent to enclosure but another thread is waiting for this command to be sent.
    - *Fix:* Skip sending command to a frozen enclosure (when it is hot removed).
    - *Risk:* Medium
  - Fixed an issue where server is taking more time to boot into OS and unable to perform any controller operation in presence of ATA Locked drive.
    - *Root cause:* Server bootup was taking time when controllers are directly attached to a ATA Locked SATA drive. Controller firmware had an issue in dispatching the next queued command to the drive while processing Error D2H response from the drive and the same sequence is repeated for a long time that caused time out for one of the SCSI ATA passthrough requests. The "*Byte 0, bits 4: 0 of NCQ error log\_info*" issue from the RLE response was not correct (that is, set to 0x80) even for all the NCQ tags 0-31 that received an error D2H response from the drive.
    - *Fix:* Modified firmware to send an errored SCSI response for those SCSI requests which are supposed to be aborted by the drive when it's in a ATA Locked drive. Preserve the failed Receive FPDMA Queued request at "*dev\_ptr->failed\_request*" before sending the RLE request to the drive as part of NCQ error handling. So that the SCSI check condition response will be sent for the failed request after the RLE command was getting processed successfully.
    - *Risk:* Low (Only with SATA Locked Drive)
  - Fixed an issue where controller lockup with code=0x00001=No module.



- *Root cause:* When an Uncorrectable ECC occurred at a particular address with 64-bit address, its firmware handler reads higher address word from wrong location which resulted in fatal assert without proper lockup or recovery path.
  - *Fix:* Updated firmware to read the correct higher address word when processing the Uncorrectable ECC error to execute recovery or lockup handling rather than asserting when reading the wrong address.
  - *Risk:* Low
- Fixed an issue where rediscovery is triggered after BME is enabled even if the channel is already discovered.
  - *Root cause:* When BME is enabled for the first time, firmware requests MCTP rediscovery, even if it had already responded to any initial discovery messages.
  - *Fix:* Add option to ignore rediscovery requests for channels where discovery has already happened.
  - *Risk:* Low
- Fixed an issue where controller lockup with code=0x3D037=MSGU.
  - *Root cause:* The controller internal RAM is only partially initialized for PCIE physical function 0 and virtual function 0. A read access to an uninitialized address may cause an ECC error and lead to a firmware lockup.
  - *Fix:* Initialize memory for the entire scratchpad areas and the ATU configuration table, including the unused physical function1, to initialize ECC for all addresses during MSGU PQI initialization sequence.
  - *Risk:* Low
- Fixed an issue where the device ready timeout is increased from 10s to 45s after reset.
  - *Root cause:* Some large-capacity SATA drives (4 TB–8 TB) take upto 30s to recover after a reset. Large-capacity SATA drives would be marked as failed in some cases as they had not recovered in the 10s allowed by the firmware before drive recovery.
  - *Fix:* Increased link reset time for SATA drives from 10s to 45s. The link reset time for SAS drives remains at 10s.
  - *Risk:* Medium
- Fixed an issue controller lockup with code=0x01ABD.
  - *Root cause:* A SATA loss of signal condition causes a target reset. As a result SCSI-ATA-Translation requires that on ATA reset other than due to a TM function, the translation layer must terminate processing of all commands for each unit affected by the reset and establish a unit attention condition. Previously, if connection with a SATA drive was re-established during LOS recovery processing the firmware would not terminate processing of all commands.
  - *Fix:* Terminate processing of all commands on LOS recovery of SATA device.
  - *Risk:* Medium
- Fixed an issue where EID is not getting set through MCTP over PCIe.
  - *Root cause:* During boot after UEFI driver load, PQI reset is being issued by host for driver loading. During PQI reset, Bus Master Enable (BME) will be disabled for some time and then re-enabled. During this small interval of time if MCTP Bus owner sends any commands to get the information about MCTP Endpoint which is discovered, the response messages are blocked in firmware as BME is disabled.
  - *Fix:* Allow MCTP messages irrespective of Bus Master Enable (BME) bit.
  - *Risk:* Low

## 2.1.2 UEFI/Legacy BIOS Fixes

This section shows the UEFI/Legacy BIOS fixes and enhancements for this release.

### 2.1.2.1 Fixes and Enhancements for UEFI Build 2.14.4/Legacy BIOS Build 2.14.2

This release includes the following fixes and enhancements:

- Added an enhancement to display inoperative MaxCache status message for hardware security update. The HII controller firmware flash menu will show MaxCache inoperative status if applicable when a firmware update is attempted that includes hardware security update.
- Added an enhancement to show drive write cache status as unknown when it is not configurable. The HII disk information will show the drive write cache status as unknown when it detects configuring write cache as not supported.
- Fixed an issue where the Managed SED Unlock Controller password in driver health form is not shown when lockout period is reset.
  - *Root cause:* The unlock option in driver health menu for remote mode Managed SED controller password remains disabled even when the lockout timer is expired or reset.
  - *Fix:* Always show the unlock option if controller is waiting for password input. Error message will be shown after the attempt, if there is a password lockout period.
  - *Risk:* Low
- Fixed an issue where the parity count migration fails for RAID 50.
  - *Root cause:* Migration processing modules only consider RAID level change and does not consider the parity number change.
  - *Fix:* Start the migration process if there is a change in parity number.
  - *Risk:* Low
- Fixed an issue where the driver health form shows the Managed SED controller locked message even after successful unlock.
  - *Root cause:* Driver health form content is not updated if controller is unlocked successfully.
  - *Fix:* Refresh the driver health form with applicable messages after the controller unlock operation.
  - *Risk:* Low

## 2.1.3 Driver Fixes

This section shows the driver fixes and enhancements for this release.

### 2.1.3.1 Linux Driver Fixes

This section shows the Linux driver fixes and enhancements for this release.

#### 2.1.3.1.1 Fixes and Enhancements for Linux Driver Build 2.1.30-031

This release includes the following fixes and enhancements.

- Fixed an issue where during the processing of a TMF on a device, the driver is stuck while waiting for I/O to be drained from the driver's internal request queue.
  - *Root cause:* During heavy I/O load, a data path request on a particular device is added to the driver's request list but encounters a conditional check where it needs more elements than are currently free in the inbound queue on a particular queue group. This request remains in the request list until the request submission function is triggered by the IRQ handler, but unfortunately no completions arrive on that queue group for a significantly long period of time. During the same time, a LUN reset is triggered on another device. While processing the TMF on the other device, attempts are made to fail I/Os queued on the device undergoing reset. However, I/Os queued on other devices are not failed due to the TMF condition check. Because of the stuck command in the request list and the failure to fail I/O on the other devices, the system experiences a LUN reset and system hang.

- *Fix:* For devices which are not undergoing reset, return DID\_QUEUE and complete the requests that are stuck in the driver's internal request queue. This adds the request back to the mid-layer queue, ensuring that it is resubmitted after a short period without decrementing the retry count in the OS SCSI-mid-layer.
- *Risk:* Low

### 2.1.3.2 Windows Driver Fixes

This section shows the Windows driver fixes and enhancements for this release.

#### 2.1.3.2.1 Fixes and Enhancements for Windows Driver Build 1010.108.0.1015

There are no known fixes and enhancements for this release.

### 2.1.3.3 FreeBSD Driver Fixes

This section shows the FreeBSD driver fixes and enhancements for this release.

#### 2.1.3.3.1 Fixes and Enhancements for FreeBSD Driver Build 4540.0.1005

There are no known fixes and enhancements for this release.

### 2.1.3.4 VMware Driver Fixes

This section shows the VMware driver fixes and enhancements for this release.

#### 2.1.3.4.1 Fixes and Enhancements for VMware Driver Build 4672.0.104

This release includes the following fixes and enhancements:

- Fixed a potential risk of data loss for MaxCache Logical Drive when configured in WriteBack mode.
  - *Root cause:* When a MaxCache logical drive is configured in WriteBack mode, a persistent notification "1931-Slot "X" - Data in write cache has been lost." is generated at each system reboot. The root cause of this POST message is attributed to the SmartPQI driver issuing dual flush cache commands, each signaling a shutdown occurrence, which leads to the loss of metadata within the SSD Cache. Consequently, there is a potential risk of data loss for logical drives that are set up with MaxCache.
  - *Fix:* To address this issue, only send a single flush cache command during shutdown.
  - *Risk:* Medium

### 2.1.4 Management Software Fixes

This section shows the management software fixes and enhancements for this release.

#### 2.1.4.1 maxView Storage Manager/ARCCONF Fixes

This section shows the maxView Storage Manager/ARCCONF fixes and enhancements for this release.

##### 2.1.4.1.1 Fixes and Enhancements for maxView Storage Manager/ARCCONF Build 26842

This release includes the following fixes and enhancements for ARCCONF/maxView.

- Added support to display the NAND and NOR Flash type as part of the controller properties in the ARCCONF and the maxView.
- Fixed an issue where the ARCCONF was not displaying the "Task" command help for the HBA controller
  - *Root cause:* Because of the HBA controller check for the ARCCONF help menu in the task command, "Task" was missing from the ARCCONF help menu for the HBA controller.
  - *Fix:* Changes added to display "Task" in the arconf help menu for the HBA controller.
  - *Risk:* Low
- Fixed an issue where the maxView was displaying a cache error alert for a controller which does not support controller cache.

- *Root cause:* maxView didn't check the "FeatCacheSupport" feature bit before enabling show device alert for controllers that don't support controller cache and have logical devices.
  - *Fix:* Added a feature bit check "FeatCacheSupport" before generating the device alert for the controllers that don't support the controller cache and have logical devices.
  - *Risk:* Low
- Fixed an issue where the maxView was not displaying the system name consistently across the reboot.
  - *Root cause:* Usually, the Redfish server starts first, followed by the maxView webserver and subscribes to events with the configured hostname. This hostname will be added to maxView events. If the maxView web server starts before the Redfish server, maxView webserver tries to subscribe to events from the Redfish server using the loopback hostname 'localhost'. Here, the server is configured with the hostname 'localhost.localdomain', which will be used for event subscription under normal scenario when the Redfish server starts first, followed by the maxView webserver. If the maxView webserver starts first, it takes the loopback hostname 'localhost'. Due to this, two different hostnames are present in the events.
  - *Fix:* The code change has been made to always use the configured hostname for event subscription. The configured hostname will be used to subscribe to events if the Redfish server starts first. If the maxView web server starts first, it will attempt to subscribe to events using the configured hostname instead of 'localhost'. If, even after retrying for 20 seconds, it is not able to connect to the Redfish server using the configured hostname, then the maxView web server will attempt to establish a connection with the Redfish server using 'localhost'.
  - *Risk:* Low
- Fixed an issue where the maxView installer was setting the write permissions for the .txt files in the installation folder.
  - *Root cause:* maxView installer was setting the write permissions for the .txt files in the installation folder.
  - *Fix:* Removed the write permissions for the .txt files in the installation folder and provided only the read-only permissions for these .txt files.
  - *Risk:* Low
- Fixed an issue where the maxView was not displaying a warning message when the Redfish server is connected through loopback network address.
  - *Root cause:* maxView was not displaying a warning message when the Redfish server is connected through loopback network address.
  - *Fix:* Added a warning message tooltip on the IP Address attribute when Redfish server is connected through loopback network address
  - *Risk:* Low
- Fixed an issue where the arconf was creating the raidzeroarray with the default stripe size of 128 kB when the raidzeroarray is created with the SAS and SATA drives.
  - *Root cause:* Instead of creating the raidzeroarray with 256 Kb as default stripe size, the ARCCONF was creating the raidzeroarray with the default stripe size of 128 kB when the raidzeroarray is created with the SAS and SATA drives.
  - *Fix:* Changed 256 kB as default stripe size for SAS and SATA drives and 128 kB for NVMe drives when raidzeroarray is used.
  - *Risk:* Low

#### 2.1.4.2 PLDM Fixes

This section shows the PLDM fixes and enhancements for this release.

### 2.1.4.2.1 Fixes and Enhancements for PLDM Release 6.40.6.0

This release includes the following fixes and enhancements:

- Added PLDM Type 6 support for Redfish Chassis for UBM Backplanes. This feature adds representation of UBM backplanes and their contained drives using Redfish resources and PDRs.
- Added PLDM Type 6 Redfish Resource:
  - UBM Chassis resource
    - It represents a single physical UBM backplane.
    - PLDM Type 6 RDE READ and HEAD operations are supported.
  - UBM Attached Drive resource
    - Will represented a single physical Drive attached to a UBM.
    - PLDM Type 6 RDE READ, UPDATE, UPDATE and HEAD operations are supported.
  - DriveCollection Resource
    - Will be the lone subordinate of a UBM Chassis resource.
    - Will publish properties as per the Redfish DriveCollection schema
    - PLDM Type 6 RDE READ and HEAD operations are supported
  - Absent Drive resource
    - Will represent a single drive bay of a UBM.
    - PLDM Type 6 RDE READ and HEAD operations are supported.
- PLDM Type 2 Command Changes:
  - GetPDR will now retrieve additional Drive and DriveCollection Redfish PDRs for each attached UBM.
  - Drive and DriveCollection Redfish PDRs will be static and will not generate PDR Change/Add/Modified alerts.
- Add Descriptor Type 0xFFFF for All Devices in Downstream Device ID Record PLDM Type 5 QueryDownstreamIdentifiers command response for downstream drives will now publish a Vendor Defined Descriptor which holds the ServiceLabel in the following format:  
Slot=<Slot#>:Port=<Port#>:Box=<Box#>:Bay=<Bay#>
- Added support to perform RDE UPDATE on Drive.HotspareReplacementMode. This feature allows Redfish clients to change the spare type of a given drive. This feature does not allow Redfish clients to modify non-Spare Drive resources.
  - Make the following changes to RDE READ on a drive resource:
    - i. Added @Redfish.WriteableProperties property to indicate if "HotspareReplacementMode" can be updated.
    - ii. Added HotspareReplacementMode@Redfish.AllowableValues if "HotspareReplacementMode" can be updated. The values are:
      - If current spare drive type is dedicated and if the mode is allowed to change to either the NonRevertible or Revertible equivalents, it publishes "NonRevertible" and "Revertible" in the AllowableValues.
      - If current spare drive type is auto-replace and if the mode is allowed to change to either the NonRevertible or Revertible equivalents, it publishes "NonRevertible" and "Revertible" in the AllowableValues.
  - RDE UPDATE on Drive resource: If the current spare drive type requested is already set and the update is allowed, then the UPDATE will succeed.

To change spare type to an auto-replace/roaming spare, the RDE operation payload should be as follows:

```
{
  "HotspareReplacementMode": "NonRevertible"
}
```

To change spare type to a shareable spare, the RDE operation payload should be as shown below:

```
{
  "HotspareReplacementMode": "Revertible"
}
```

- Added UPDATE support for Drive.HotspareType.
  - RDE UPDATE on a Drive resource will now support updating the HotspareType property. The only allowable value for UPDATE on Drive.HotspareType is None, that is, RDE clients can only remove a spare from their configuration.
  - RDE READ on a Drive resource will now publish the following property:  
HotspareType@Redfish.AllowableValues
- PLDM Type 6 Add Drive.FirmwareVersion support for RDE READs on Drive resources
  - RDE READ on a Drive resource will now publish FirmwareVersion property. The FirmwareVersion property will have the same value as the Revision property.
- AllowablePattern annotations for Storage SetEncryptionKey request properties. This feature adds AllowablePattern fields to the SetEncryptionKey action on a Redfish Storage resource. The added fields are:
  - CurrentEncryptionKey@Redfish.AllowablePattern
  - EncryptionKey@Redfish.AllowablePattern
  - EncryptionKeyIdentifier@Redfish.AllowablePattern
- PLDM Type 6 READ and UPDATE Support for Storage.HotspareActivationPolicy. Added support to perform RDE UPDATE on Storage.HotspareActivationPolicy. This feature allows Redfish clients to manage the controller's spare activation mode. Below are the changes to RDE READ on a storage resource:
  - a. On RDE READ, the HotspareActivationPolicy is published.
  - b. Added @Redfish.WriteableProperties property to indicate if "HotspareActivationPolicy" can be updated
  - c. Added HotspareActivationPolicy@Redfish.AllowableValues if "HotspareActivationPolicy" can be updated.

The allowable values are:

If current HotspareActivationPolicy is OnDriveFailure and if the mode is allowed to change to either the OnDriveFailure or OnDrivePredictedFailure equivalents, "OnDriveFailure" and "OnDrivePredictedFailure" will be published in the AllowableValues.

If current HotspareActivationPolicy is OnDrivePredictedFailure and if the mode is allowed to change to either the OnDrivePredictedFailure or OnDriveFailure equivalents, "OnDrivePredictedFailure" and "OnDriveFailure" will be published in the AllowableValues.

If the current Spare Activation Policy requested is already set and the update is allowed, then the RDE UPDATE on Drive resource will succeed.

To change Spare Activation Policy to on drive failure, the RDE operation payload should be as follows:

```
{
  "HotspareActivationPolicy": "OnDriveFailure"
}
```

To change spare activation policy to on drive predicted failure, the RDE operation payload should be as follows:

```
{
  "HotspareActivationPolicy": "OnDrivePredictedFailure"
}
```

- Added PLDM Type 6 READ support for StorageController.Links.PCIeFunction. RDE READ on the StorageController resource will now publish the following properties:
  - Links.PCIeFunctions@odata.count
  - Links.PCIeFunctions@odata.id

This property gets published as a deferred binding string of %PF0, where 0 is the PCIe Function of the controller.

- Updated permission flags on resources if operation cannot be performed. The PermissionFlags returned for an RDE operation have been brought into accordance with the PLDM Type 6 specification.
  - RDE READ and HEAD operations will return non-zero PermissionFlags in the response.
  - Operations failing with ERROR\_NOT\_ALLOWED will have a response with non-zero PermissionFlags.
  - All other operations and failures will result in PermissionFlags of 0x00 in the response.
- Added support for Drive.SlotCapableProtocols. This property will contain the drive protocols that are supported in the slot being read. It will now be published as part of the RDE READ response on all Drive resources, including those representing empty UBM bays. The following values are supported:
  - SAS
  - SATA
  - NVMe

The Drive.SlotCapableProtocols property will be published as an array containing some or all of the above values depending on hardware specifications. If no protocols are supported, then an empty array will still be published.

- Added PLDM Type 6 READ Support for Drive.DriveFormFactor. RDE READ on a SAS/SATA Drive resource will now include the property DriveFormFactor in the response. The following values are supported:
  - Drive2\_5
  - Drive3\_5
  - M2

For NVMe drives and all SAS and SATA drives with a form factor not listed above, the Drive.DriveFormFactor property will not be published.

- Added DriveMetrics support for all controllers, which allows reporting of some basic drive performance statistics information.
- Added PLDM Type 6 Redfish DriveMetrics Resource
  - Each DriveMetrics resource will be mapped to a single Drive resource.
  - PLDM Type 6 RDE READ and HEAD operations are supported.
- Updated RDE READ on Drive resources to add the Metrics property which will hold a link to the Drive's related DriveMetrics resource.
- PLDM Type 2 Command Changes

- GetPDR will now retrieve additional DriveMetrics Redfish Resource PDRs for each Drive Resource.
  - DriveMetrics Redfish Resource PDRs will generate PDRRepoChange events and the PDRs will be added and deleted as drives are hot-plugged and hot-removed.
  - The GetSchemaURI and GetSchemaDictionary commands can fetch information related to the new DriveMetrics resource schema dictionary.
- Added support for StorageController.AssetTag to allow Redfish clients to assign an asset tag string to a StorageController resource which will persist across boot cycles.
- Added the RDE READ operation support for Drive resources which represent a Revertible (that is, shareable) spare drive.
- Modified the ordering of Volume links in various collections.
  - Modified the ordering of volume links in the VolumeCollection resource's Members array to group Volumes by the array of Drives on which they are housed, then sorting each array's Volumes in order of increasing starting LBA offset.
  - Modified the ordering of volume links in the Drive resource's Links.Volumes array to group the Volumes by the array of Drives on which they are housed in the case of a shareable standby spare drive, then sorting each array's Volumes in order of increasing starting LBA offset.
- Fixed an issue where battery count is reported as one in storagecontroller for the system that supports only cache but does not have a battery.
  - *Root cause:* The specific check for a battery was not being performed. Instead a general check for the existence of a cache was performed assuming that a backup power source existed.
  - *Fix:* In locations where the backup power source information is being reported, perform a specific check to determine if a backup power source exists. While modifying the checks, it was noticed that the `WriteCachePolicy@Redfish.AllowableValues` indicated `ProtectedWriteBack`. Therefore modifications were made to disallow `ProtectedWriteBack` when no backup power source exists for the cache.
  - *Risk:* Low
- Fixed an issue where incorrect IOPerfModeEnabled/ReadCachePolicy/WriteCachePolicy after volume creation.
  - *Root cause:* The read cache percentage wasn't defaulting to the correct value.
  - *Fix:* The default read cache percentage is being set correctly in all cases now.
  - *Risk:* Low
- Fixed an issue where `NegotiateTransferParameters` command allows requester part size that is not a power of two.
  - *Root cause:* While issuing `NegotiateTransferParameters` command with requester part size that is not a power of two, the command returns success instead of `ERROR_INVALID_DATA`. A code mistake caused only the first byte of the requester part size to be checked for power of two requirement.
  - *Fix:* Modified the code so that requester part size will be checked properly.
  - *Risk:* Low
- Fixed an issue where there are missing fields in `FileDescriptorPDR` for crash dump log.
  - *Root cause:* The `FileDescriptorPDR` is missing the `SuperiorDirectoryFileIdentifier` and `FileMaximumFileDescriptorCount` fields.



- *Fix:* Added the `SuperiorDirectoryFileIdentifier` and `FileMaximumFileDescriptorCount` fields to the crash dump log's `FileDescriptorPDR` to be in compliance with the PLDM Type 2 specification.
- *Risk:* Low
- Fixed an issue where the cache status is degraded due to a missing or a failed battery, the `WriteCacheDegraded` alert will have a severity of "Critical" instead of "Warning".
  - *Root cause:* The `WriteCacheDegraded` alert was hard-coded to only be sent with a severity of "Critical" when the battery was missing or failed.
  - *Fix:* Modified the `WriteCacheDegraded` alert logic to send a severity of "Warning" for certain controllers when the battery is missing or failed.
  - *Risk:* Low
- Fixed an issue where the `CacheSummary.Status.State` property in the `StorageController` resource is published as "Disabled" when the cache is not supported on the controller.
  - *Root cause:* The case of the controller cache not being supported was not being directly handled when encoding the `StorageController` resource.
  - *Fix:* The `CacheSummary` health and state will be set to OK and absent, respectively, when the controller cache is not supported.
  - *Risk:* Low
- Fixed an issue where `OperationExecutionFlags` is incorrect in the `RDEOperationStatus` response for a completed operation.
  - *Root cause:* When an operation status is completed, the `OperationExecutionFlags` should not have the `HaveResultPayload` or the `HaveCustomResponseParameters` bits set. A previous change that updated the response handler for the `RDEOperationStatus` command to properly set the `OperationExecutionFlags` bits based on the actual current state of the operation exposed that `OperationExecutionFlags` bits were not properly being reported.
  - *Fix:* Unset the `HaveResultPayload` and `HaveCustomResponseParameters` bits in `OperationExecutionFlags` for completed RDE operations.
  - *Risk:* Low
- Fixed an issue where the controller serial log contains repeated prints related to an error with handling a `GetSchemaDictionary` request.
  - *Root cause:* The request was being sent to fetch the registry schema class dictionary for various Redfish resources, but this dictionary is not supported. The response completion code is erroneously being set to `ERROR_INVALID_DATA` instead of `ERROR_UNSUPPORTED` as directed in the PLDM Type 6 specification.
  - *Fix:* Modified the `GetSchemaDictionary` and `GetSchemaURI` command handler functions to return `ERROR_UNSUPPORTED` if the registry schema class is passed in the request.
  - *Risk:* Low
- Fixed an issue when a `StorageController` resource has a `CacheSummary.Status.Conditions` entry with `MessageId` of `StorageDevice.1.1.WriteCacheDataLoss`, the associated severity is warning instead of critical. Additionally, some controllers expect a severity of Ok for the `StorageDevice.1.1.WriteCacheTemporarilyDegraded` condition but are seeing warning instead.
  - *Root cause:* The RDE READ encoder for the `StorageController` resource had an incorrect severity hard-coded for `WriteCacheDataLoss` and was missing the logic to determine the correct severity for `WriteCacheTemporarilyDegraded`.

- *Fix:* Updated the RDE READ encoder to publish the expected severity values for the `WriteCacheDataLoss` and `WriteCacheTemporarilyDegraded` `CacheSummary.Status.Conditions` entries.
  - *Risk:* Low
- Fixed an issue where `DurableNameFormat` for RAID volumes are reported as UUID in RDE READ operations, but should be NAA.
  - *Root cause:* The original implementation set the `DurableNameFormat` to be UUID. Also changes were made to add dashes to make the `DisplayName` conform to the UUID format.
  - *Fix:* Change `DurableNameFormat` to be NAA and revert the previous hyphenation changes since UUID is no longer the format.
  - *Risk:* Low
- Fixed an issue where hot removal of a drive that is part of a RAID array resulted in the `Status.Conditions` in a RDE Drive READ operation to have `DriveFailure` and `DriveRemoved` entries. They should be `DriveRemoved` and `DriveMissing`.
  - *Root cause:* The code was incorrectly treating the event as a failure followed by a removal.
  - *Fix:* The code now treats the event correctly as a removal resulting in a missing drive.
  - *Risk:* Low
- Fixed an issue where a `GetSensorReading` request with the `rearmEventState` option set for the version state sensor did not move the version state sensor `presentState` or `previousState` to normal.
  - *Root cause:* The `rearmEventState` option was not being accounted for when a controller flash was detected.
  - *Fix:* If a controller flash is detected, the `rearmEventState` option will be considered when processing a `GetSensorReading` request.
  - *Risk:* Low
- Fixed an issue where on certain controllers, `GetSensorReading` on a `NumericSensor` of `EntityType` `DEVICE_FILE` returns an error completion code of `INVALID_SENSOR_ID`.
  - *Root cause:* The logic that validates the requested `SensorId` from the `GetSensorReading` request was incorrectly returning that the `SensorId` is not valid on controllers that do not support individual drive temperature numeric sensors.
  - *Fix:* Corrected the logic to validate the requested `SensorId` regardless of whether or not the controllers support individual drive numeric sensors.
  - *Risk:* Low
- Fixed an issue when RDE UPDATE on a SED HBA volume resource to take ownership of the SED or to revert a SED fails.
  - *Root cause:* Controller firmware requires more than 6 seconds to take ownership or revert a SED. The time needed by controller firmware is more than PLDM's allowed response time and hence the RDE UPDATE operation fails.
  - *Fix:* The following changes have been made:
    - RDE UPDATE on a SED HBA volume resource will now be performed by a long running task.
    - If the MC has not negotiated TASK support by using `RDENegotiateRedfishParameters` command, RDE UPDATE on a SED HBA volume will return a completion code of `ERROR_NOT_ALLOWED`.
  - *Risk:* Medium

- Fixed an issue where the StorageController.Status.State property would not be 'Updating' when controller firmware had been flashed and was pending activation.
  - *Root cause:* When encoding the StorageController.Status.State property, no check was made for the presence of controller firmware pending activation.
  - *Fix:* The presence of controller firmware pending activation is now considered when encoding StorageController.Status.State. The StorageController.Status.State property will be 'Updating' if pending controller firmware is present.
  - *Risk:* Low
- Fixed an issue when inducing a failure condition in a Volume, for example, removing two drives from a RAID 5 volume's array, the volume's Status.Conditions array contains entries for both VolumeFailure and VolumeOffline conditions when only VolumeFailure is expected.
  - *Root cause:* The logic governing whether or not to publish the VolumeOffline condition did not contain a check of whether or not a VolumeFailure condition had been previously established.
  - *Fix:* Modified the logic for publishing the VolumeOffline condition to only publish if the Volume is not experiencing a VolumeFailure condition in addition to the existing checks.
  - *Risk:* Low
- Fixed an issue when a drive hosting a RAID 0 Volume is removed and reinserted, a VolumeOK event is not generated.
  - *Root cause:* VolumeOfflineCleared condition doesn't check the volume's previous health state. This results VolumeOfflineCleared event is generated instead of VolumeOK when the logical drive is re-enabled.
  - *Fix:* Modified VolumeOfflineCleared condition to check the volume's previous health state. If the previous health state is a failure, do not generate VolumeOfflineCleared event.
  - *Risk:* Low
- Fixed an issue where an RDE UPDATE operation on a volume resource to modify DedicatedSpareDrives is removing existing spares even if the operation fails to add new spare drives.
  - *Root cause:* The DedicatedSpareDrives logic is set up to attempt to remove all existing spares before adding new ones. In the existing logic, if there is a failure at any point in the removal/addition of spares, the code would stop the RDE UPDATE without attempting to revert to the existing spare configuration. Instead of reverting, the code would leave the spare configuration in whatever state it was in during the failure.
  - *Fix:* Modified the logic to revert the spare setup to the existing configuration if there is a failure during the PATCH operation.
  - *Risk:* Low
- Fixed an issue when a drive resource hosting a RAID volume with a FailurePredicted value of true is hot-removed, then RDE READ on the drive resource will show FailurePredicted as false.
  - *Root cause:* When a drive has a Status.State of Absent, no check of the drive's predictive failure state is made, and its default value of false is published.
  - *Fix:* Added a check of a drive's predictive failure status in all circumstances regardless of its Status.State value.
  - *Risk:* Low

## 2.2 Limitations

This section shows the limitations for this release.

### 2.2.1 General Limitations

This release includes the following general limitations.

- The following are the limitations of Multi-Actuator:
  - Supports only:
    - HBA drive
    - Windows/Linux/VMware
    - Intel/AMD
    - UEFI mode (for multi-LUN display)
- The NCQ Priority feature is currently not supported in this release.

## 2.2.2 Firmware Limitations

This section shows the firmware limitations for this release.

### 2.2.2.1 Limitations for Firmware Release 03.01.30.106

This release includes the following limitations:

- If a boot volume is secured by Managed SED Remote Key Management (RKM) or Managed SED Adapter Password enabled Local Key Management (LKM), it will fail to write Windows memory dump file during Windows OS crash dump.
  - *Workaround:* Do not use secured volumes as described above with an OS boot logical drive.
- Persistent Event Logs (PEL) will be cleared under the following conditions:
  - Upgrading from firmware releases prior to 03.01.17.56 to 03.01.17.56 or later firmware releases.
  - Downgrading from firmware releases 03.01.17.56 or later to firmware releases prior to 03.01.17.56.
- Firmware downgrade is blocked if disk-based transformation is in-progress.
  - *Workaround:* Wait for the transformation to complete and retry the firmware downgrade.
- Transformation is blocked if a reboot is done after the firmware update is pending, and the flashed new firmware version is older than 03.01.17.56.
  - *Workaround:* Reboot the system.
- Logical drive is not detected when disk-based transformation is in-progress during logical drive movement to a different controller and the different controller has a firmware version older than 03.01.17.56, or, the firmware downgrade occurred while internal-cache based transformation was in progress, but the Backup Power Source failed before firmware activation.
  - *Workaround:* Move the logical drive to a controller with firmware version 03.01.17.56 or later.
- Firmware downgrade from firmware version 3.01.30.106 to any older firmware version is blocked if Managed SED is enabled.
  - *Workaround:* Disable Managed SED and try firmware downgrade.
- Managed SED cannot be enabled on the controller when reboot is pending after firmware downgrade from firmware version 3.01.23.72 to any older firmware version.
  - *Workaround:* Reboot the controller and enable the Managed SED.
- Flashing from 3.01.30.106 back to 3.01.28.82 or 3.01.26.36 may result in the spin down spare policy being changed to the default setting specified in the board configuration file.
  - *Workaround:* If the default board configuration file specified setting is not the required setting, then re-apply the spin down spare policy using host management software.

## 2.2.3 UEFI/Legacy BIOS Limitations

This section shows the UEFI/Legacy BIOS limitations for this release.

### 2.2.3.1 Limitations for UEFI Build 2.14.4/Legacy BIOS Build 2.14.2

There are no known limitations for this release.

## 2.2.4 Driver Limitations

This section shows the driver limitations for this release.

### 2.2.4.1 Linux Driver Limitations

This section shows the Linux driver limitations for this release.

#### 2.2.4.1.1 Limitations for Linux Driver Build 2.1.30-031

This release includes the following limitations:

- SL-Micro 6.0 fails to boot after installation on 4Kn drives.
  - *Workaround:* This is a SUSE issue and only workaround is to use non-4Kn drives.
- On some distributions (RHEL7.9, RHEL8.2, RHEL8.3, SLES15SP2, SLES15SP3, OpenEuler 20.03LTS, and 22.03LTS including SP releases), the driver injection (DUD) install will hang if an attached drive (either HBA mode or Logical Volume) has Write Cache enabled.
  - *Workaround:* There are two workarounds for this issue:
    - Ensure that the Write Cache is disabled for any attached drive.
    - For RHEL7.9/8.2/8.3 and OpenEuler 20.03LTS, 22.03LTS, add `rd.driver.blacklist=smartpqi` to the grub entry along with `inst.dd`.
- RHEL driver injection (DUD) install where OS ISO is mounted as virtual media on BMC based servers (non-ILO). Installer will hang after driver injection. It is reported on RHEL 8.5, 8.6, 9.0, and 9.1.
  - *Workaround:*
    - Load the OS from USB device instead of virtual media.
    - Load the OS from virtual media but initiate ISO verification (media test) during the installation followed by ESC to cancel the media test.
    - Edit grub to include the boot argument "nompath". Replace "inst.dd" with "nompath inst.dd" for DUD install.
- Oracle 9 UEK 7 kernel causes SmartPQI rpm dependency failures. This is an issue with how the kernel package was created by Oracle. Correct UEK7 kernel for Oracle 9, which is expected in the mid-October UEK7 release, version number is still pending.
 

**Note:** This does not affect Oracle 8 UEK 7.

  - *Workaround:* Install the rpm using "--nodeps" when dependency failures occur.
    - Update:
      - For SmartPQI driver versions > 2.1.20-020 and UEK7 kernels >= 5.15.0-3.60.2.el9uek.x86\_64, the SmartPQI rpm will install normally.
      - For UEK7 kernels < 5.15.0-3.60.2.el9uek.x86\_64, the SmartPQI rpm needs to be installed using the "--nodeps".
- On AMD systems, the system might crash or hang due to a bug in the IOMMU module. For details, see [lore.kernel.org/linux-iommu/20191018093830.GA26328@suse.de/t/](https://lore.kernel.org/linux-iommu/20191018093830.GA26328@suse.de/t/).
  - *Workaround:* Disable the IOMMU setting option in BIOS.
- Depending on hardware configurations, the SmartPQI `expose_ld_first` parameter may not always work consistently.
  - *Workaround:* None
- On some distributions (including RHEL 9.0/Oracle Linux 9.0), you are unable to inject the OOB driver (DUD) during install when a multi-actuator drive is attached.
  - *Workaround:* Install using the inbox driver, complete OS installation, then install the OOB driver.

### 2.2.4.2 Windows Driver Limitations

This section shows the Windows driver limitations for this release.

### 2.2.4.2.1 Limitations for Windows Driver Build 1010.106.0.1028

This release includes the following limitations:

- The Windows driver issues an internal flush cache command for flushing the controller cache to the drives before changing the power state of the system (during shutdown/reboot/hibernate ). Due to many factors, example of speed of drives, size of cache, type of data in cache, and so on, the time taken by the controller to flush the cached data can exceed the operating system specified timeout values. A system crash can be expected in those scenarios. Controller cache flushing will continue and complete while the system is in the BSOD state. In general, it is advised not to do heavy write operations on logical drives composed of slow drives while initiating a system shutdown in Windows 10 environments.
- In certain circumstances, the installation may fail on Windows Server 2016 and Windows 2012 R2 after selecting drives.
  - *Workaround:* Follow these steps to ensure drives are clean and all partitions are removed before beginning a new installation:
    - a. Hit Shift + F10 to open the command prompt
    - b. Type `Diskpart`
    - c. Type `List Disk`
    - d. Select the disk you want to clean. For example, to select Disk 0 type `select disk 0.`
    - e. Type `Clean`
- A system crash may occur when hibernating a system installed on a Dual Actuator drive.
  - *Workaround:*
    - Avoid hibernating the system while running heavy I/Os to multiple Dual Actuator drives.
    - Stop running the I/Os to the drives and then hibernate the system.
    - Reboot the server to recover the system.

### 2.2.4.3 FreeBSD Driver Limitations

This section shows FreeBSD driver limitations for this release.

#### 2.2.4.3.1 Limitations for FreeBSD Driver Build 4540.0.1005

This release contains the following limitations:

- FreeBSD 13.2 and later OS installations will fail with the out of box driver.
  - *Workaround:* Install with inbox driver then update to latest.

### 2.2.4.4 VMware Driver Limitations

This section shows VMware driver limitations for this release.

#### 2.2.4.4.1 Limitations for VMware Driver Build 4672.0.104

This release includes the following limitation:

- If the controller SED Encryption feature is “On” and locked, Datastores created from secured logical drives on the controller are not automatically mounted even after unlocking the controller, they are not visible through the ESXi hypervisor client.
  - *Workaround:* Use the command `vmkfstool -V` or ESXCLI storage filesystem rescan. Alternatively, use the Rescan option from the Devices tab in the Hypervisor’s Storage section. Any of these options solve the issue by forcing a rescan, causing the datastore to mount.
- A controller lockup may occur when using VMDirectPath on a single-processor AMD system. These lockups have been seen with VMs running Linux and Windows. No known workaround at the present time. If a lockup of a passed-through controller occurs, a reboot of the ESXi server may be required to clear the lockup condition and restore the virtual machine to working condition.
- Customers may encounter failures when attempting to add new Logical Drives (LD), particularly in cases involving a dead path.

- *Workaround:* To facilitate recovery of new LD, customers are required to clear the dead path initially. Following the clearance of the dead path, if the newly created LD is still not exposed, then it is required to initiate a driver level rescan using the appropriate management tool. While clearing the dead path fails, a host reboot is required.

## 2.2.5 Management Software Limitations

This section shows management software limitations for this release.

### 2.2.5.1 maxView Storage Manager/ARCCONF Limitations

This section shows the maxView Storage Manager/ARCCONF limitations for this release.

#### 2.2.5.1.1 Limitations for maxView Storage Manager/ARCCONF Build 26842

There are no known limitations for this release.

### 2.2.5.2 PLDM Limitations

This section shows the PLDM limitations for this release.


#### 2.2.5.2.1 Limitations for PLDM Release 6.40.6.0

There are no known limitations for this release.

### 3. Updating the Controller Firmware

This section describes how to update the controller firmware to the latest release.

---

 **Important:** When downgrading firmware, there may be cases when newer hardware is not supported by an older version of firmware. In these cases, attempting to downgrade firmware will be prevented (fail). It is recommended to regularly qualify newer firmware versions, to ensure that newer hardware is supported in your system(s)

---

#### 3.1 Updating Controllers to Latest Firmware

If running firmware is 3.01.00.006 or lower, please contact Adaptec Apps team at [ask.adaptec.com](https://ask.adaptec.com).

##### 3.1.1 Upgrading to 3.0X.XX.XXX Firmware

1. For controllers running 3.01.02.042 or higher firmware, flash with 3.0X.XX.XXX version of firmware "SmartFWx200.bin" provided in this package using maxview or ARCCONF utility.
2. Power cycle the server.



## 4. Revision History

**Table 4-1.** Revision History

Revision	Date	Description
Q	07/2024	Updated for SR 3.4.0 release.
P	02/2024	Updated for SR 3.3.4 release.
N	11/2023	Updated for SR 3.3.2 release.
M	10/2023	SR 3.3.0 patch release with maxView™ version B26068.
L	10/2023	SR 3.2.0 patch release with maxView™ version B25339.
K	08/2023	Updated for SR 3.3.0 release.
J	03/2023	Updated for SR 3.2.4 release.
H	11/2022	Updated for SR 3.2.2 release.
G	07/2022	Updated for SR 3.2.0 release.
F	02/2022	VMware driver version changed from 4250.0.120 to 4252.0.103.
E	02/2022	Updated for SR 3.1.8 release.
D	12/2021	Updated for SR 3.1.6.1 release. Updated Fixes and Enhancements for maxView Storage Manager/ARCCONF section for log4j vulnerabilities.
C	11/2021	Updated for SR 3.1.6 release.
B	08/2021	Updated for SR 3.1.4 release.
A	06/2021	Document created.

## Microchip Information

### The Microchip Website

Microchip provides online support via our website at [www.microchip.com/](http://www.microchip.com/). This website is used to make files and information easily available to customers. Some of the content available includes:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user’s guides and hardware support documents, latest software releases and archived software
- **General Technical Support** – Frequently Asked Questions (FAQs), technical support requests, online discussion groups, Microchip design partner program member listing
- **Business of Microchip** – Product selector and ordering guides, latest Microchip press releases, listing of seminars and events, listings of Microchip sales offices, distributors and factory representatives

### Product Change Notification Service

Microchip’s product change notification service helps keep customers current on Microchip products. Subscribers will receive email notification whenever there are changes, updates, revisions or errata related to a specified product family or development tool of interest.

To register, go to [www.microchip.com/pcn](http://www.microchip.com/pcn) and follow the registration instructions.

### Customer Support

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Embedded Solutions Engineer (ESE)
- Technical Support

Customers should contact their distributor, representative or ESE for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in this document.

Technical support is available through the website at: [www.microchip.com/support](http://www.microchip.com/support)

### Microchip Devices Code Protection Feature

Note the following details of the code protection feature on Microchip products:

- Microchip products meet the specifications contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is secure when used in the intended manner, within operating specifications, and under normal conditions.
- Microchip values and aggressively protects its intellectual property rights. Attempts to breach the code protection features of Microchip product is strictly prohibited and may violate the Digital Millennium Copyright Act.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of its code. Code protection does not mean that we are guaranteeing the product is “unbreakable”. Code protection is constantly evolving. Microchip is committed to continuously improving the code protection features of our products.

### Legal Notice

This publication and the information herein may be used only with Microchip products, including to design, test, and integrate Microchip products with your application. Use of this information in any other manner violates these terms. Information regarding device applications is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure

that your application meets with your specifications. Contact your local Microchip sales office for additional support or, obtain additional support at [www.microchip.com/en-us/support/design-help/client-support-services](http://www.microchip.com/en-us/support/design-help/client-support-services).

THIS INFORMATION IS PROVIDED BY MICROCHIP "AS IS". MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE, OR WARRANTIES RELATED TO ITS CONDITION, QUALITY, OR PERFORMANCE.

IN NO EVENT WILL MICROCHIP BE LIABLE FOR ANY INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, OR CONSEQUENTIAL LOSS, DAMAGE, COST, OR EXPENSE OF ANY KIND WHATSOEVER RELATED TO THE INFORMATION OR ITS USE, HOWEVER CAUSED, EVEN IF MICROCHIP HAS BEEN ADVISED OF THE POSSIBILITY OR THE DAMAGES ARE FORESEEABLE. TO THE FULLEST EXTENT ALLOWED BY LAW, MICROCHIP'S TOTAL LIABILITY ON ALL CLAIMS IN ANY WAY RELATED TO THE INFORMATION OR ITS USE WILL NOT EXCEED THE AMOUNT OF FEES, IF ANY, THAT YOU HAVE PAID DIRECTLY TO MICROCHIP FOR THE INFORMATION.

Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

## Trademarks

The Microchip name and logo, the Microchip logo, Adaptec, AVR, AVR logo, AVR Freaks, BesTime, BitCloud, CryptoMemory, CryptoRF, dsPIC, flexPWR, HELDO, IGLOO, JukeBlox, KeeLoq, Kleer, LANCheck, LinkMD, maXStylus, maXTouch, MediaLB, megaAVR, Microsemi, Microsemi logo, MOST, MOST logo, MPLAB, OptoLyzer, PIC, picoPower, PICSTART, PIC32 logo, PolarFire, Prochip Designer, QTouch, SAM-BA, SenGenuity, SpyNIC, SST, SST Logo, SuperFlash, Symmetricom, SyncServer, Tachyon, TimeSource, tinyAVR, UNI/O, Vectron, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

AgileSwitch, APT, ClockWorks, The Embedded Control Solutions Company, EtherSynch, Flashtec, Hyper Speed Control, HyperLight Load, Libero, motorBench, mTouch, Powermite 3, Precision Edge, ProASIC, ProASIC Plus, ProASIC Plus logo, Quiet- Wire, SmartFusion, SyncWorld, Temux, TimeCesium, TimeHub, TimePictra, TimeProvider, TrueTime, and ZL are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, Augmented Switching, BlueSky, BodyCom, Clockstudio, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, Espresso T1S, EtherGREEN, GridTime, IdealBridge, In-Circuit Serial Programming, ICSP, INICnet, Intelligent Paralleling, IntelliMOS, Inter-Chip Connectivity, JitterBlocker, Knob-on-Display, KoD, maxCrypto, maxView, memBrain, Mindi, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICkit, PICtail, PowerSmart, PureSilicon, QMatrix, REAL ICE, Ripple Blocker, RTAX, RTG4, SAM-ICE, Serial Quad I/O, simpleMAP, SimpliPHY, SmartBuffer, SmartHLS, SMART-I.S., storClad, SQL, SuperSwitcher, SuperSwitcher II, Switchtec, SynchroPHY, Total Endurance, Trusted Time, TSHARC, USBCheck, VariSense, VectorBlox, VeriPHY, ViewSpan, WiperLock, XpressConnect, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

The Adaptec logo, Frequency on Demand, Silicon Storage Technology, and Symmcom are registered trademarks of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.

© 2024, Microchip Technology Incorporated and its subsidiaries. All Rights Reserved.

ISBN: 978-1-6683-4924-3

## **Quality Management System**

For information regarding Microchip's Quality Management Systems, please visit [www.microchip.com/quality](http://www.microchip.com/quality).

# Worldwide Sales and Service

AMERICAS	ASIA/PACIFIC	ASIA/PACIFIC	EUROPE
<b>Corporate Office</b> 2355 West Chandler Blvd. Chandler, AZ 85224-6199 Tel: 480-792-7200 Fax: 480-792-7277 Technical Support: <a href="http://www.microchip.com/support">www.microchip.com/support</a> Web Address: <a href="http://www.microchip.com">www.microchip.com</a>	<b>Australia - Sydney</b> Tel: 61-2-9868-6733 <b>China - Beijing</b> Tel: 86-10-8569-7000 <b>China - Chengdu</b> Tel: 86-28-8665-5511 <b>China - Chongqing</b> Tel: 86-23-8980-9588 <b>China - Dongguan</b> Tel: 86-769-8702-9880 <b>China - Guangzhou</b> Tel: 86-20-8755-8029 <b>China - Hangzhou</b> Tel: 86-571-8792-8115 <b>China - Hong Kong SAR</b> Tel: 852-2943-5100 <b>China - Nanjing</b> Tel: 86-25-8473-2460 <b>China - Qingdao</b> Tel: 86-532-8502-7355 <b>China - Shanghai</b> Tel: 86-21-3326-8000 <b>China - Shenyang</b> Tel: 86-24-2334-2829 <b>China - Shenzhen</b> Tel: 86-755-8864-2200 <b>China - Suzhou</b> Tel: 86-186-6233-1526 <b>China - Wuhan</b> Tel: 86-27-5980-5300 <b>China - Xian</b> Tel: 86-29-8833-7252 <b>China - Xiamen</b> Tel: 86-592-2388138 <b>China - Zhuhai</b> Tel: 86-756-3210040	<b>India - Bangalore</b> Tel: 91-80-3090-4444 <b>India - New Delhi</b> Tel: 91-11-4160-8631 <b>India - Pune</b> Tel: 91-20-4121-0141 <b>Japan - Osaka</b> Tel: 81-6-6152-7160 <b>Japan - Tokyo</b> Tel: 81-3-6880-3770 <b>Korea - Daegu</b> Tel: 82-53-744-4301 <b>Korea - Seoul</b> Tel: 82-2-554-7200 <b>Malaysia - Kuala Lumpur</b> Tel: 60-3-7651-7906 <b>Malaysia - Penang</b> Tel: 60-4-227-8870 <b>Philippines - Manila</b> Tel: 63-2-634-9065 <b>Singapore</b> Tel: 65-6334-8870 <b>Taiwan - Hsin Chu</b> Tel: 886-3-577-8366 <b>Taiwan - Kaohsiung</b> Tel: 886-7-213-7830 <b>Taiwan - Taipei</b> Tel: 886-2-2508-8600 <b>Thailand - Bangkok</b> Tel: 66-2-694-1351 <b>Vietnam - Ho Chi Minh</b> Tel: 84-28-5448-2100	<b>Austria - Wels</b> Tel: 43-7242-2244-39 Fax: 43-7242-2244-393 <b>Denmark - Copenhagen</b> Tel: 45-4485-5910 Fax: 45-4485-2829 <b>Finland - Espoo</b> Tel: 358-9-4520-820 <b>France - Paris</b> Tel: 33-1-69-53-63-20 Fax: 33-1-69-30-90-79 <b>Germany - Garching</b> Tel: 49-8931-9700 <b>Germany - Haan</b> Tel: 49-2129-3766400 <b>Germany - Heilbronn</b> Tel: 49-7131-72400 <b>Germany - Karlsruhe</b> Tel: 49-721-625370 <b>Germany - Munich</b> Tel: 49-89-627-144-0 Fax: 49-89-627-144-44 <b>Germany - Rosenheim</b> Tel: 49-8031-354-560 <b>Israel - Ra'anana</b> Tel: 972-9-744-7705 <b>Italy - Milan</b> Tel: 39-0331-742611 Fax: 39-0331-466781 <b>Italy - Padova</b> Tel: 39-049-7625286 <b>Netherlands - Drunen</b> Tel: 31-416-690399 Fax: 31-416-690340 <b>Norway - Trondheim</b> Tel: 47-72884388 <b>Poland - Warsaw</b> Tel: 48-22-3325737 <b>Romania - Bucharest</b> Tel: 40-21-407-87-50 <b>Spain - Madrid</b> Tel: 34-91-708-08-90 Fax: 34-91-708-08-91 <b>Sweden - Gothenberg</b> Tel: 46-31-704-60-40 <b>Sweden - Stockholm</b> Tel: 46-8-5090-4654 <b>UK - Wokingham</b> Tel: 44-118-921-5800 Fax: 44-118-921-5820