

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.....	23
3.1. Updating Controllers to Latest Firmware.....	23
4. Revision History.....	24
Microchip Information.....	25
The Microchip Website.....	25
Product Change Notification Service.....	25
Customer Support.....	25
Microchip Devices Code Protection Feature.....	25
Legal Notice.....	25
Trademarks.....	26
Quality Management System.....	27
Worldwide Sales and Service.....	28

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

Solutions release	3.3.0
Package release date	October 31, 2023
Firmware version	3.01.23.68
UEFI/Legacy BIOS	2.8.3/2.8.2
Driver versions	<p>Windows Drivers:</p> <ul style="list-style-type: none"> Windows 2022, 2019, 2016, Windows 11, 10: 1010.74.0.1020 <p>Linux SmartPQI:</p> <ul style="list-style-type: none"> RHEL 7/8/9: 2.1.24-046 SLES 12/15: 2.1.24-046 Ubuntu 18/20/22: 2.1.24-046 Oracle Linux 7/8/9: 2.1.24-046 Citrix Xenserver 8: 2.1.24-046 Debian 10/11: 2.1.24-046 <p>VMware:</p> <ul style="list-style-type: none"> VMware ESX 7.0/8.0: 4530.0.104 <p>FreeBSD:</p> <ul style="list-style-type: none"> FreeBSD 12/13: 4410.0.1005
ARCCONF/maxView	4.14.0.26068
PLDM	6.25.9.0

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, 2016, Windows 11, 10	x64

Table 1-5. Linux Drivers

OS	Version
RHEL 9.2 ¹ , 9.1, 9.0 ² , 8.8 ¹ , 8.7, 8.6, 8.5, 7.9	x64
SLES 12 SP5, SP4	x64
SLES 15 SP5 ¹ , SP4, SP3, SP2	x64
Ubuntu 20.04.5, 20.04.4, 20.04, 18.04.5, 18.04.4	x64
Ubuntu 22.04.2, 22.04.1, 22.04	x64
Oracle Linux 7.9 UEK6U3	x64
Oracle Linux 9.2 ¹ , 9.1, 9.0, 8.8 ¹ , 8.7, 8.6, UEK7	x64
Debian 11.6, 10.13	x64
Fedora 38 (inbox)	x64
Citrix XenServer 8.2.1	x64
Rocky Linux 9.1	x64

Notes:

1. New OS support—minimally tested drivers in this release. Fully supported drivers are expected in the next release.
2. Support based off August 2022 RHEL 9.0 ISO refresh.

Table 1-6. FreeBSD and VMware Drivers

OS	Version
ESX 8.0 U1, 7.0 U3/U2	x64
FreeBSD 13.2, 12.4	x64

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 arconf_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 03.01.23.72

This release includes the following fixes and enhancements.

- Added support for Remote Key Management of Managed SED.
- Added support for 256 bytes Key Management Service (KMS) key identifier.
- Added support to improve flash interoperability with UBM backplanes.
- Enabled power sensor monitoring
- Added support for full 15-byte controller serial number through PBSI
- Added support for cascaded expander to uniquely identify the attached enclosure and physical devices.
- Moved VDM message processing to a lower priority thread to allow for quicker responses to PCIe configuration cycles from the Host.
- Added support in PBSI to show maximum and negotiated link rate for physical drives.
- Fixed an issue for Managed SED in Local Key Management (LKM) mode where firmware allows to import the foreign SED while the adapter password is not received yet.
 - Root Cause: Firmware does not check for the received adapter password while processing the request to import a foreign SED. Firmware should fail the request if the adapter password is not provided as the master key is not available until the adapter password is provided. Without the master key, importing a foreign SED cannot be performed.
 - Fix: Firmware will check if the adapter password is not received then fail the request.
 - Risk: Low
- Fixed an issue where failed managed SED-secured fault-tolerant logical drives moved to an OK state on importing all foreign SEDs.
 - Root Cause: All physical SEDs will be foreign SEDs when they are added as replacements in a failed managed SED secured fault-tolerant logical drive. While importing these replaced foreign SEDs, the firmware will try to update the logical drive state every time an SED is imported. Firmware will mark all the imported replacement SEDs to an OK state on importing the last foreign SED but failed to keep the logical drive state to bring in media exchange mode. This resulted in the logical drive state getting updated to the OK state instead of FAILED (Mex) state.
 - Fix: Do not clear the media exchange mode status of the logical drive until the completion of the logical drive state is updated.
 - Risk: Low
- Fixed an issue where the slot number is shown as unknown for a failed physical drive present in SES supported enclosure.
 - Root Cause: The SCSI Enclosure Services (SES) supported enclosure will provide multiple additional status pages(0Ah), which consist of each physical drive's information such as the device type, WWN or SAS address, slot number and so on. This additional status page data

- will be compared against the controller-detected enclosure-specific data. On a successful comparison, the firmware will assign the slot number for the physical drive. For a failed physical drive, the WWN or SAS address comparison failed and resulted in firmware skipping the slot number assignment.
- Fix: If the physical drive is detected by the controller and WWN or SAS address comparison failed, then compare the device slot number. If it matches, assign the slot number to the physical drive.
 - Risk: Low
- Fixed an issue where the controller is reset to factory defaults and RAID is lost after an abrupt reboot.
 - Root Cause: An abrupt reboot during the local controller settings update caused its corruption and the discovery protocol of all the connectors got reset to default (that is, SGPIO). When the UBMx4 backplane is connected to the controller using a bifurcated connector, each bay will have two PHYs. As controller settings were reset, the firmware is not aware that each bay supports multiple PHYs and added both handles from the same PHY as different entries. Now firmware detects identical RAID metadata in two physical drives, fails to identify the right data, and ignores the RAID metadata.
 - Fix: Added logic to check and discard duplicate handles from the same physical drives.
 - Risk: Low
 - Fixed an issue where the degraded managed SED logical drive is moved to the failed state after inserting a foreign physical drive as a replacement when the transformation is in progress.
 - Root Cause: When there is an active spare assigned to a failed SED in managed SED logical drive and the failed SED is replaced with a foreign SED, the firmware will issue transformation writes to both the spare and replaced foreign SED. Since the replaced SED is foreign and is yet to be imported, firmware writes to the foreign SED will fail and result in transformation and logical drive failure.
 - Fix: Firmware writes to the foreign SED will be skipped. Once foreign SED is imported, the data rebuild will occur on the imported SED.
 - Risk: Low
 - Fixed an issue where the spare drive rebuild is not started on the queued managed SED logical drives after inserting a foreign replacement SED.
 - Root Cause: The firmware will start a spare rebuild for a logical drive when the data drive is in the FAILED state. After replacing a failed data SED with a foreign SED in managed SED logical array, the SED state is updated from FAILED to WRONG_REPLACED, and due to this firmware failed to start the spare rebuild on the left over queued managed SED logical drives.
 - Fix: Firmware will start the spare rebuild for the logical drives if the data drive is in FAILED state or WRONG_REPLACED state.
 - Risk: Low
 - Fixed an issue where managed SED mirrored logical drives queued for rebuilding moved to a DEGRADED state after inserting a foreign replacement SED.
 - Root Cause: When a foreign SED is added as a replacement drive, the firmware will mark it as WRONG_REPLACED till it gets imported and will update the logical drive's state. For mirrored logical drives, the firmware has not included the WRONG_REPLACED drive count while updating the logical drive state and wrongly updated it to the DEGRADED state instead of leaving them in the NEEDS_REBUILD state. Due to this, rebuild is not getting started for the mirrored logical drives even though there is a dedicated spare available.
 - Fix: Included the WRONG_REPLACED drive count while updating the logical drive state for mirrored logical drives.

- Risk: Low
- Fixed an issue where transformation does not resume after reboot on a spare assigned managed SED logical drive if a data drive failure is observed before providing the adapter password.
 - Root Cause: After reboot, the managed SED logical drive will be in the SED_DATA_LOCKED state and waits for the adapter password before resuming the transformation. At this stage, firmware is starting spare rebuild on detection of a data drive failure, even though the logical drive is in SED_DATA_LOCKED state. With this, transformation and rebuilding are set together on a logical drive, and none of them are able to proceed after providing the adapter password.
 - Fix: If the logical drive is in the SED_DATA_LOCKED or SED_LOCKED state, then the spare rebuild is blocked.
 - Risk: Low
- Fixed an issue where the foreign unconfigured SED is not exposed to the OS after import.
 - Root Cause: If any foreign unconfigured SED is connected to the controller, the controller firmware will fail to access the RAID metadata region of the SED, as it is locked. The controller firmware will not expose this foreign SED to OS until it gets unlocked, to avoid any operations on the SED. While importing the foreign SED, firmware does not try to access the RAID metadata region and is not exposing the SED to the host.
 - Fix: Firmware will access the RAID metadata region of the SED while importing. If the read is successful, the firmware will expose the SED to the host.
 - Risk: Low
- Fixed an issue where the spare rebuild does not start on the SSD array.
 - Root Cause: During boot, if controller detects any data drive is missing from the logical drive due to loose cable or any other reason, controller will move to an abnormal volume state and will rely on firmware to determine the actual volume state based on available drives. At this stage, all the logical drives in the array are moved to DEGRADED state. On spare availability, firmware moved all the logical drives to NEEDS_REBUILD state and started REBUILDING on one logical drive. For the SSDs, before starting the first rebuild, the firmware will issue a TRIM operation to erase it. Once the operation completes, SSD will be marked as a bad drive to avoid any other I/O operations on it. At this stage, firmware received a bypass volume state request, and all logical drives states got updated again. As the spare SSD is still going through TRIM operation and marked as bad, the firmware updated the logical drives states from NEEDS_REBUILD to DEGRADED. Due to this state change, firmware is not able to start the rebuild even after the TRIM operation is completed.
 - Fix: Added additional checks to treat the drive going through TRIM operation as a replacement drive.
 - Risk: Low
- Fixed an issue where adding a foreign/otherwise-owned SED as a replacement into a degraded non-SED logical drive changed the logical drive state to SED_LOCKED after reboot.
 - Root Cause: During boot, the logical drives which contain LOCKED SED drives will get marked as SED_LOCKED by firmware irrespective of managed SED status of the controller.
 - Fix: Firmware will mark the drive as failed if it fails to pass the SED qualification process.
 - Risk: Low
- Fixed an issue where the consistency check progress status is displayed as “In progress” even after disabling the consistency check on the controller.
 - Root Cause: Firmware is populating consistency check progress status irrespective of consistency check enable/disable status on the controller.

- Fix: Firmware will not populate the consistency check status when it is disabled on the controller.
 - Risk: Low
- Fixed an issue where the managed SED logical drive is going to a Degraded state if a foreign SED is added as a replacement during the heal-array operation.
 - Root Cause: When the host initiated a heal-array operation on a DEGRADED managed SED logical drive and the failed data drive is replaced with foreign SED during this operation, the firmware failed to ignore the foreign drive in old RAID metadata. Due to this, the firmware does not update the managed SED logical drive status to OK after completion of the heal-array operation and left it in the Degraded state.
 - Fix: Firmware will ignore the old RAID metadata configuration and update the managed SED logical drive status according to new RAID metadata on successful completion of the heal-array operation.
 - Risk: Low
- Fixed an issue where all the managed SED logical drives are moved to the REBUILDING state.
 - Root Cause: On any drive failure, if a spare is available, the firmware will assign a spare to the failed drive and starts the data rebuild for all the logical drives one by one. If a foreign SED is added as a replacement and one more data failure happens, then at the end of the current rebuild cycle the firmware updates all the managed SED logical drives status to REBUILDING due to an incorrect check.
 - Fix: Fixed the incorrect check. The firmware will update the managed SED logical drive state as per the configured SED drive's status.
 - Risk: Low
- Fixed an issue where a controller lockup may occur after an interrupted clear configuration operation with Managed SED logical drives.
 - Root Cause: When a clear configuration operation is interrupted due to a panic shutdown, the next boot up results in the controller reading the datastore on the SED that may indicate a RAID metadata range is enabled. The firmware then sets a flag indicating the RAID metadata range already exists. When a new Managed SED logical drive is created, the controller skips creating the RAID metadata range and the next system boot sequence the firmware has a lockup trying to save the RAID metadata because the firmware should not have skipped creating the RAID metadata range.
 - Fix: Firmware saves a flag in NVRAM to indicate a clear configuration process is occurring, and if interrupted, on the next boot up reverts the SED to OFS. Subsequent logical drive creations will ensure the RAID metadata range is created.
 - Risk: Low
- Fixed an issue where the firmware does not block the revert with PSID on a configured foreign SED.
 - Root Cause: Firmware should block the revert with PSID for a configured foreign SED.
 - Fix: Firmware blocks the revert with PSID on a configured foreign SED.
 - Risk: Low
- Fixed an issue to improve the backup power status reporting.
 - Root Cause: Backup power status was not entirely accurate. It could be reported as not present instead of failed.
 - Fix: Added extra error handling to more accurately depict the backup power status.
 - Risk: Low
- Fixed migrating a maxCache when destination controller has existing logical drive.

- Root Cause: maxCache on new controller was getting re-initialized because it detected new logical drives.
 - Fix: Fixed logic to not re-initialize maxCache when existing logical drive on new controller was detected.
 - Risk: Low
- Fixed an issue where the spare rebuild started on foreign SED spare drive.
 - Root Cause: Firmware is not checking the spare drive's OK status before assigning the spare to the degraded logical drive and starts the spare rebuild.
 - Fix: Added a check in the firmware to not activate the spare if it is not in an OK state.
 - Risk: Low
- Fixed an issue where maxCache status reports as destroyed after importing the managed SED logical drive along with maxCache.
 - Root Cause: When managed SED logical drive along with maxCache is moved to another controller with LKM (Local Key Management)/RKM (Remote Key Management) configured, firmware does not load the maxCache configuration after importing the logical drive and maxCache.
 - Fix: Load maxCache configuration when the foreign logical drives are imported along with maxCache.
 - Risk: Low
- Fixed an issue where the spare rebuild is not started for managed SED logical drive when failed SED is replaced with a foreign SED.
 - Root Cause: Firmware will assign the spare only for the failed drives. When the failed drive was replaced with foreign SED, firmware failed to assign the spare to it, and the rebuild does not start.
 - Fix: Added additional checks in firmware to assign the spare to foreign SED drives.
 - Risk: Low
- Fixed an issue where transformation does not resume after moving managed SED logical drive to another controller.
 - Root Cause: When managed SED logical drives are moved to another controller with LKM (Local Key Management)/RKM (Remote Key Management) configured, firmware does not resume the transformation after importing the logical drives.
 - Fix: Added code to resume the transformation after importing the managed SED logical drives.
 - Risk: Low
- Fixed an issue where Predictive Spare Rebuild (PSR) started on the SED_DATA_LOCKED managed SED logical drive.
 - Root Cause: When managed SED logical drive is in SED_DATA_LOCKED state and one of the SEDs is marked as a predictive failure, firmware checks the spare SED status and if it's not foreign, then it assigns the spare and starts a predictive spare rebuild.
 - Fix: Modified the firmware check to consider the SED status of both data and spare drive to start the PSR.
 - Risk: Low
- Fixed an issue where transforming managed SED logical drive is going to the SED_DATA_LOCKED state when one of the SEDs is replaced with a foreign SED.
 - Root Cause: When a single foreign SED is added as a replacement drive to a transforming logical drive, the firmware will mark the foreign SED as a bad drive in the RAID metadata. But

- when updating the logical drive state, the firmware checks the presence of a foreign SED flag irrespective of the bad drive status of the SED and updates the state to SED_DATA_LOCKED.
- Fix: Firmware will ignore the foreign SED flag status if the drive is marked as bad while updating the logical drive state.
 - Risk: Low
- Fixed an issue where an Uncorrectable error PSOD observed after power on Nutanix VM.
 - Root Cause: The firmware logs to the UART every time there is a BME bit change. Logging to the UART adds additional latency. In some corner cases, using the UART logging results in the firmware taking longer than 10 ms to respond to the PCIe Configuration Write to change the BME bit, resulting in an Unsupported Request response back to the originator of the PCIe Configuration Write TLP.
 - Fix: Moved the logging of the BME bit change to the internal logging system rather than to the UART.
 - Risk: Low
 - Fixed an issue where Non-fast path commands stuck in SAT firmware pending queue.
 - Root Cause: In some special cases, non-fast path commands such as INQUIRY can remain stuck in the firmware's SAT command pending queue, if the command cannot be sent out during NCQ traffic.
 - Fix: On completion of NCQ commands, add firmware to service the pending queue.
 - Risk: Medium
 - Fixed an issue where Error on Test Unit Ready (TUR) command for NVMe SEDs.
 - Root Cause: The SCSI-to-NVMe translation for the TUR issues a Read command which the drive rejects if the drive is still locked.
 - Fix: Added a unique flag to mark I/Os as TURs. The NVMe completion path detects the completion for the TUR with the check condition for SCSI sense key MEDIUM ERROR (03) and ASC/ASCQ ACCESS DENIED - INVALID LU IDENTIFIER (20/09) and if the drive is NVMe SED, the firmware changes the check condition status to a good status.
 - Risk: Low
 - Fixed an issue where controller not assigned with EID.
 - Root Cause: In this issue, BME is disabled when the controller attempts to send a Discovery Notify to the MCTP bus owner as per the specification. The controller will set a retry timer to attempt another Discovery Notify after 5 seconds. There is a corner case where BME is enabled, the Set EID comes from the host, and the retry timer expires and the controller sends the Discovery Notify. This sequence violates the specification because the Discovery Notify follows the Set EID. For this reason, the discovery fails and the controller is marked as Unknown by the MCTP Bus owner.
 - Fix: If the Discovery Notify return timer is enabled when a Set EID message is received, disable the Discovery retry mechanism.
 - Risk: Medium
 - Fixed an issue where the Unique ID of different SATA disks is the same in Windows®.
 - Root Cause: The controller is reporting the same unique ID for a SATA drive in the same slot, even if the SATA drives are different family models because the lower level firmware did not register its API to the upper level callback firmware function pointer used to enable retrieving the unique ID from the SATA drive.
 - Fix: Register the lower level firmware API to the upper level firmware callback function pointer to enable retrieving the unique ID from the SATA drive.
 - Risk: Low

- Fixed an issue where controller was incorrectly returning mode page information regarding the Write Cache attribute of NVMe drives.
 - Root Cause: The controller is incorrectly marking the write cache setting as non-changeable.
 - Fix: Add translation to the SCSI-to-NVME translation firmware to ensure that if the NVMe drive supports volatile write cache, then to report the write cache is changeable.
 - Risk: Low
- Fixed an issue where OS fails to see controller due to long boot time due to a locked SED timing out command.
 - Root Cause: When an I/O times out, it takes a long time to recover that I/O. SED drive is timing out lots of I/Os so it takes too long to discover this drive. This failure does go away when the drive is unlocked.
 - Fix: Set a flag when a locked SED fails a command for I/O timeout and stop post spinup operations. When drive is unlocked, check this flag and then do post spinup operations.
 - Risk: Medium
- Fixed an issue where a maxCache logical drive is migrated from one controller to another and maxCache logical drive failed error message could be seen at power up.
 - Root Cause: The RAID metadata on the maxCache logical drive was getting invalidated.
 - Fix: Corrected logic to mark RAID metadata as valid on the maxCache logical drive.
 - Risk: Low
- Fixed an issue when the local mode has managed SED encryption enabled and tries to change the master key identifier without changing the master key does not successfully update the new master key identifier.
 - Root Cause: Logic was not saving the new master key/master key identifier values in the NVRAM.
 - Fix: Updated logic to make sure to check if master key/master key identifier has valid data, so it gets updated in NVRAM.
 - Risk: Low
- Fixed an issue that controller firmware flashed event has random characters at the end of the event message.
 - Root Cause: When logging the event, a local variable that saves the active ROM image is used without being initialized. The variable is a two-byte array. The first byte is used to save "A" or "B", ROM image version. The second byte is expected to be 0, and it is used at the end of the event message. Since the array is not initialized, "A" or "B" is not null terminated, so random characters could appear at the end of the event message.
 - Fix: Initialize the local variable before putting it into use.
 - Risk: Low
- Fixed an issue where taking ownership of enterprise drive was failing on boot after panic shutdown.
 - Root Cause: When changing a master key occurs, several SED authorities are changed to a new key. This SED flow (open the session, perform an SED task, and end session) gets interrupted due to panic shutdown, but drives are not power-cycled (hence, not reset) since the drives are attached to an enclosure which has its own power source. The drive is left in some state and expecting the next SED operation. Instead, now due to reboot, firmware restarts and attempts to open a new session to validate the datastore on the drive and a start session failure occurs.
 - Fix: When start session failure occurs, depending on the failure, error recovery is implemented and then retries a start session.

- Risk: Low
- Fixed an issue where firmware lockup is observed after hot removing a SES device while a LUN reset to the device is in progress.
 - Root Cause: While processing a host-issued LUN reset to the device and if the device is hot removed, LUN reset task management is completed and cleared from the management list for the device. Lockup is observed when firmware attempts to clear LUN reset task management again for the removed device.
 - Fix: Before issuing reset to the device, if the device does not exist and LUN reset task management is not present in the list, then the reset request is already cleared, so firmware should not attempt to clear it again.
 - Risk: Low
- Fixed an issue where the I/O latency value is not as expected for NCQ priority SMR drives.
 - Root Cause: The RAID path did not have support for NCQ priority commands.
 - Fix: Added support for priority bits in messages derived from the driver and propagate to lower layer firmware interface.
 - Risk: Low
- Fixed an issue where firmware fails to capture vendor-specific expander log.
 - Root Cause: This is caused by a code change that firmware relies on the driver to provide data transaction direction. The Linux SCSI layer is providing direction as data-out instead of data-in. This failed in CentOS 7.9 but passed in RHEL 9.1. This kind of incompatibility happened among various flavors of Linux if we depended on the driver for data transaction direction.
 - Fix: For T10 supported commands like inquiry, firmware does not depend on the driver for data transfer direction. It sets it according to the T10 specification. The fix is that firmware sets the data transfer direction as data-in for the read expander log only based on the WDC OEM specification, and does not rely on the driver's input.
 - Risk: Low
- Fixed an issue where the controller locked up on inserting a foreign SED replacement into the degraded volume during disk-based transformation.
 - Root Cause: When the disk-based transformation is progressing on a degraded logical drive, a foreign SED drive is added to the volume. At this stage, the newly added foreign SED drive will be in the WRONG_REPLACED state and the firmware metadata range is in the locked state. As part of the disk-based transformation, transformation progress data will be saved to firmware metadata on all the drives of the logical drive at regular intervals. In this path, we are sending metadata writes to all the drives except failed drives. As the newly inserted foreign drive is still in a locked state and the controller has not taken ownership of it, we are seeing commands failing with aborted command status and also receiving duplicate completions for these commands and causing the controller lockup.
 - Fix: Skip metadata_write to SED drives if the firmware metadata locking range is not enabled yet.
 - Risk: Low
- Fixed an issue where maxView reports for an NVMe SSD Predictive fail but other tools do not report any failure.
 - Root Cause: The NVMe Translation layer is returning Predictive Failure if any of the "critical" bits are set in a drive's SMART/Health Information log page. Specifically in this issue, the bits indicating under or over temperature are set to "1" which will cause a Predictive Failure response.

- Fix: The bit for over or under temperature should be excluded from this check as this condition is not representative of a Predictive Failure.
- 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 2.8.3/Legacy BIOS Build 2.8.2

This release includes the following UEFI fixes and enhancements:

- Added Remote Key Management support for controller-managed SED encryption. The remote key management server is utilized for encryption key generation and storage.
- Added controller password support for the Remote mode controller-managed SED encryption.
- Added multi actuator devices support for EFI SCSI pass through protocol. The EFI SCSI pass through protocol supports device enumeration and pass thru commands to multi actuator devices.
- Fixed an issue where PCIe slot information is not provided in the configuration tools.
 - Root Cause: UEFI driver does not get the PCIe slot information from EFI SMBIOS protocol.
 - Fix: Find PCIe slot number from the connected host root bridge configuration space if the slot information is not found in EFI SMBIOS protocol method.
 - Risk: Low

2.1.3 Driver Fixes

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

2.1.3.1 Windows Driver Fixes

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

2.1.3.1.1 Fixes and Enhancements for Windows Driver Build 1010.74.0.1020

- Added registry value "LunResetBehavior" feature. Setting this registry value changes the SRB_FUNCTION_RESET_LOGICAL_UNIT behavior. The new LUN reset behavior is to return the SRB status after the internal TMF LUN reset command completes. If the TMF does not complete, the driver will let it hang until timeout. The new behavior for the TMF LUN reset timeout is set to what the SRB timeout passes into the miniport. HW_RESET_BUS hardware callback routine will pause the controller I/O for up to 25 seconds while checking to see if controller completes all I/O within 18 seconds. If I/O is still not completed then the callback hardware bus reset will be failed. If the driver does not detect any outstanding I/O after 18 seconds, then the hardware bus reset callback will be marked as successful.

Note: The new reset LUN behavior will only occur if the registry value "LunResetBehavior" is present and set to 1.
- Fixed an issue where the random drives in the system were going offline after a hot plug and reboot.
 - Root Cause: Incorrect logic in traversing the report_physical_lun response while hot adding drives to the system. In the drive hotplug handling path, the driver was using an incorrect size while traversing the list of physical devices without checking the firmware feature support.
 - Fix: Added logic to check the firmware feature set to determine the size of the RPL entry while traversing the RPL response.
 - Risk: Low
- Fixed an issue where an incorrect tag table is assigned for PQI queue groups.
 - Root Cause: The incorrect tag table assignment for the PQI queue groups when there are more than eight NUMA nodes present in the system. The driver was skipping the creation

of IOBypass queues associated with certain queue groups because of the invalid tag table assignment.

- Fix: Resolved issues with the invalid tag table assignment when there are more than eight NUMA nodes present within the system.
- Risk: Medium

2.1.3.2 Linux Driver Fixes

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

2.1.3.2.1 Fixes and Enhancements for Linux Driver Build 2.1.24-046

This release includes the following fixes and enhancements.

- Added support for ABORT handler in the driver in order to avoid I/O stalls across all devices attached to a controller when I/O requests time out.
- Added sysfs entry for NUMA node in /sys/block/sdX/device. NUMA node detail is added for each exposed device similar to NVMe devices.

2.1.3.3 VMware Driver Fixes

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

2.1.3.3.1 Fixes and Enhancements for VMware Driver Build 4530.0.104

This release includes the following fixes and enhancements:

- Fixed an issue when PSOD occurs while attempting to access memory which had already been released.
 - Root Cause: PSOD happened when one CPU released a device and freed memory. Simultaneously, another CPU was attempting to free the same memory triggered by a hot-plug timeout.
 - Fix: Modifications made to avoid the double-freeing of the device memory.
 - Risk: Medium
- Fixed an issue where PSOD is observed during array creation and deletion.
 - Root Cause: The PSOD error occurred as a result of a NULL de-reference within a function. This issue arises when a report logical LUN fails and an external RAID device is connected, leading to a NULL de-reference.
 - Fix: Added changes to avoid the NULL de-reference.
 - Risk: Low
- Fixed an issue where datastore creation using logical drive from VMware client GUI is too slow.
 - Root Cause: In the current SmartPQI driver, when a RAID volume is created, drives associated with RAID volumes are added to the remove device list and the upper layer will be notified after 20 minutes regarding the removal of the drive. These drives will then display a timeout error in the VMware client GUI.
 - Fix: When a RAID volume is created, the drives associated with RAID volumes are removed at once, avoiding the 20 minute wait time.
 - Risk: Low
- Fixed an issue where the Hotswapped HBA drives are detected after 20 minutes or when a manual rescan is done.
 - Root Cause: When a new device is hotswapped with an old device on the same slot, both the new and the old device will have the same `scsi3addr`. Due to this, the new device will be assigned the marked for removal flag status, resulting in not being added to the new device list during device discovery.
 - Fix: The device marked for removal flag status will only be set if `scsi3addr`, model number, and serial number of both devices are equal. If `scsi3addr` is the same but serial or model

number are different, the drive will be detected as new and will be added to the new device list.

- Risk: Low

2.1.3.4 FreeBSD Driver Fixes

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

2.1.3.4.1 Fixes and Enhancements for FreeBSD Driver Build 4410.0.1005

There are no known fixes for this release.

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 26068

Microchip strongly recommends that maxView users update to the latest version of the tools to avoid a security vulnerability that has since been resolved.

2.1.4.1.2 Fixes and Enhancements for maxView Storage Manager/ARCCONF Build 26064

This release includes the following fixes and enhancements for Arccconf/maxView:

- Added remote Key management service (KMS) support for the managed SED.
- Added support to display the CPLD revision and Platform image revision in Arccconf and maxView.
- Added support to display the supercap temperature and voltage information in Arccconf and maxView.
- Blocked conversion from non-secure to secure volume for managed SED.
- Added UBM controller firmware upgrade support in Arccconf and maxView.
- Added SPDM Certificate Storage and Management support.
- Fixed an issue where phantom enclosures are displayed under every connector when there was a VPP backplane in the configuration.
 - Root Cause: maxView/Arccconf was discovering invalid enclosure object per connector when there is a VPP backplane in the configuration.
 - Fix: Implemented changes to skip adding the invalid enclosure objects without a SEP device to the configuration.
 - Risk: Low
- Fixed an issue where invalid enclosure slot count was displayed in maxView.
 - Root Cause: maxView was displaying invalid connector IDs for an enclosure where enclosure has multiple expanders in it, resulting in wrong slot count.
 - Fix: Implemented changes to add the proper connector ID for the enclosure with multiple expanders.
 - Risk: Low
- Fixed an issue where GETSMARTSTATS command is failing in Arccconf.
 - Root Cause: The Arccconf command resolver could not find the associated GETSMARTSTATS command resulting in a failure to execute the command.
 - Fix: Implemented changes to load the GETSMARTSTATS command in Arccconf.
 - Risk: Low
- Fixed an issue where auto discovery function in maxView is not working in a specific configuration.

- Root Cause: The firewall setting was blocking SSDP packets which were used for auto discovery functionality. This resulted in maxView not discovering the specific windows machines during auto discovery.
- Fix: Added firewall inbound rule for the maxView Redfish server port. Also, a discover button in auto discovery dialog to refresh the auto discovered servers in maxView.
- 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 6.25.9.0

This release includes the following fixes and enhancements:

- Added support for self-contained activation of storage enclosure firmware flashed using Type 5 downstream device firmware update.
- Added RDE READ support for the following property annotations to the VolumeCapabilities resource:
 - CapacityBytes@Redfish.AllowableNumbers
 - MediaSpanCount@Redfish.AllowableNumbers
 - StripSizeBytes@Redfish.AllowableNumbers
- Added a new descriptor of Type 0x010A (IEEE EUI-64 Identifier) for NVMe drives appearing in the response to a QueryDownstreamIdentifiers request.
- Changed the Availability state set of the controller composite state sensor to require a rearm in order to transition from a state of Starting to Enabled.
- Changed the Version state set of the controller composite state sensor to reflect changes in firmware version in downstream devices in addition to the controller.
- Updated the Storage resource to use the v1.14.0 schema and added RDE READ support for the following properties:
 - EncryptionMode
 - LocalEncryptionMode
- All drives connected to the controller which are not configured as a data or spare drive for a RAID Volume resource will now have an associated Volume resource, informally referred to as an HBA Volume or JBOD Volume, with RAIDType of "None" automatically created by the controller.
 - These Volumes will have Redfish URIs and PLDM Type 5 resource IDs listed in the Volume PDR published using a GetPDR request for that PDR handle.
 - Configuration changes such as creation and deletion of RAID Volumes and unconfigured drive removal or insertion will result in `pldmPDRRepositoryChgEvent` events being sent to any active event listeners.
 - RDE READ for an unconfigured drive resource will have a Links.Volumes entry for its associated HBA Volume resource.
 - RDE READ for the StorageController resource will have the value of "None" added to its SupportedRAIDTypes value array.
 - RDE READ for the VolumeCollection resource will have entries for HBA Volumes in its Members property array, and Members@odata.count will add these Volume resources to its count value.
- Fixed an issue where PLDM Type 5 downstream device firmware update fails on Microchip (SXP 24G SAS-4 Expander) SEPs.
 - Symptom: PLDM Type 5 GetFirmwareData fails on SXP 24G SAS-4 Expanders.

- Root Cause: PLDM uses 16K buffer chunks; whereas, SXP 24G SAS-4 Expanders will only accept 4K buffers.
 - Fix: For expander SEPs, break the 16K buffer into 4K chunks for flashing.
 - Risk: Low
- Fixed an issue of inappropriate returning Allow equal to POST on Storage and Drive to advertise the actions.
 - Symptom: Redfish clients observe the POST value being returned in the Allow header for Redfish requests for Drive and Storage resources when only GET and HEAD should be returned.
 - Root Cause: The implementation of RDE ACTION operations for these resources erroneously included a change to set the CREATE bit in the `PermissionFlags` bitfield in RDE command responses.
 - Fix: Reverted the setting of the `PermissionFlags` CREATE bit for these resources when ACTION operation support has been negotiated.
 - Risk: Low.
- Fixed an issue when the energy pack is not required, `StorageController[CacheSummary][Status][Health]` shall be OK.
 - Symptom: Users would receive cache and battery alerts on systems where an energy pack is not applicable. Redfish `StorageController[CacheSummary][Status][Health]` would show statuses other than OK when an energy pack was not applicable and there were no ECC errors.
 - Root Cause: Incorrect assumptions on what hardware setups are available to users.
 - Fix: Added checks for read cache percentage and NBWC to determine if a backup power source is applicable. Cache and battery alerts are filtered if a backup power source is not applicable. Redfish `StorageController[CacheSummary][Status][Health]` will be OK if a backup power source is not applicable and there are no ECC errors.
 - Risk: Medium
- Fixed an issue where `MediaSpanCount` allows a volume creation with -1 negative value.
 - Symptom: A user is able to create a volume despite specifying a negative `MediaSpanCount` value.
 - Root Cause: Internal code would treat a -1 value as a valid unsigned integer.
 - Fix: Added a check for `MediaSpanCount` input by checking if the most significant bit is set. If set, will send a `PropertyValueIncorrect` extended message for invalid value.
 - Risk: Low
- Fixed an issue where `MaxMembers` property should be displayed with Odata Type `CollectionCapabilities.v1_2_0`.
 - Symptom: The `VolumeCollection` resource contains an annotation object `@Redfish.CollectionCapabilities`. This annotation object has a child property called `MaxMembers` which was added to the `v1_2_0` `CollectionCapabilities` schema. However, the `@odata.type` for this annotation object is published as `CollectionCapabilities.v1_1_0.CollectionCapabilities`.
 - Root Cause: The annotation's `@odata.type` was not correctly updated when the `MaxMembers` property was added.
 - Fix: Updated the `@odata.type` string to the value `CollectionCapabilities.v1_2_0.CollectionCapabilities`.
 - Risk: Low
- Fixed an issue for possible memory leak in RDE GET on a Drive resource.

- Symptom: An RDE Get operation will have a memory leak if one of the Binary Encoded JSON (BEJ) encoding calls fails while encoding the Identifiers section of the response.
- Root Cause: The macros used to perform the BEJ encoding perform an early return after logging the error. In the case of the Identifiers section there is a buffer that is allocated before the encoding starts which needs to be freed once the encoding completes. The early return skips the code that performs the free.
- Fix: New macros were created that set a flag rather than return early. The flag is used to skip down to the free call early. After the free, if the flag is set the code performs the return.
- Risk: Low
- Fixed an issue where the incorrect WriteCachePolicy was being reported on a RDE READ on a Volume resource.
 - Symptom: Creating a Volume with WriteCachePolicy = UnprotectedWriteBack, then inducing a temporarily disabled condition in the cache results in the Volume's WriteCachePolicy to be reported as ProtectedWriteBack.
 - Root Cause: When the cache is in a temporarily disabled state, no check of the controller NBWC setting was made when setting the WriteCachePolicy for a volume.
 - Fix: Added a check of the controller's NBWC setting when setting the WriteCachePolicy in the Volume resource READ encoder.
 - Risk: Low
- Fixed an issue where PersistentCacheSize value fails to set '0' under Controller Cache Summary after removing the backup power source from the server.
 - Symptom: After removing the backup power source, the "PersistentCacheSizeMiB" field is non-zero when reading the Storage Controller page.
 - Root Cause: The conditional statement which sets the persistent cache size to zero was expecting the battery status to be "power source not present". However, the battery status was set to "power source not applicable". So the conditional statement was bypassed and the persistent cache size was not set to zero.
 - Fix: The battery status value for "power source not applicable" was removed. The battery status will be initialized to a value of zero. And thereafter, will be set to the value returned from SA_GetControllerBackupPowerSourceStatus. Additional changes made to assist with the fix are:
 - The RDA battery_status field was converted from an enumeration to two bitfields.
 - The first bitfield in the battery status contains the various states for the backup power source.
 - The second bitfield in the battery status contains information about the use for the backup power source status. In the second bitfield, a bit was added to determine when alerts would be sent for a "battery missing" event and all write cache events.
 - Risk: Medium
- Fixed an issue where PLDM Type 2 GetPDRRepositoryInfo returns incorrect RepositorySize when no physical drives are present.
 - Symptom: Mismatch in the PDR Repository size and number of records for a PLDM terminus when a user queries the PLDM Type 2 GetPDRRepositoryInfo command for a configuration with zero drives.
 - Root Cause: A Drive Action PDR is still present in the repository despite there not being any drives present on the device.
 - Fix: Deleting Drive Action PDR when all drives are removed if Action is negotiated. Re-adding Drive Action PDR when the first drive gets added if Action is negotiated.
 - Risk: Low

- Fixed an issue where WriteCachePolicy@Redfish.AllowableValues value is incorrect after removing the backup power source from the server.
 - Symptom: Removing the battery after creating a volume which uses the cache displays the patchable values for the write cache as "Off", "ProtectedWriteBack", and "UnprotectedWriteBack". The only patchable value which should be displayed in this instance is "Off".
 - Root Cause: When determining the write cache patchable values, the state of the write cache was not being taken into consideration.
 - Fix: When the write cache is in the "Degraded" state for any reason, the only patchable value will be "Off".
 - 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)

2.2.2 Firmware Limitations

This section shows the firmware limitations for this release.

2.2.2.1 Limitations for Firmware Release 03.01.23.72

- 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.
- Firmware downgrade from firmware version 3.01.23.72 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.
- Power cycle to the enclosure may be needed if connected server goes through abnormal shutdown under the following condition: SED operation on OPAL drives like taking ownership, reverting the ownership, or changing the master key where firmware internally performs open session, performs SED management, and ends session gets interrupted due to abnormal

shutdown on the server. This condition causes firmware to restart on reboot while drives are left off in the middle of performing SED task so drives need to be power cycled also.

- Workaround: Allow the change master key operation to complete before shutting down the server.
- If SEDs are in an external enclosure, power cycle the external enclosure and SEDs before powering up the server with the controller.
- Under certain high-traffic conditions, if SATA drives are present on a port with Dynamic Channel Multiplexing (DCM) 6G SATA operation enabled and the drives are attached to an expander, controller lockup may occur.
 - Workaround: Firmware detects this configuration and will drop the link to 12G from 24G.

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 2.8.3/Legacy BIOS Build 2.8.2

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 Windows Driver Limitations

This section shows the Windows driver limitations for this release.

2.2.4.1.1 Limitations for Windows Driver Build 1010.74.0.1020

There are no known limitations for this release.

2.2.4.2 Linux Driver Limitations

This section shows the Linux driver limitations for this release.

2.2.4.2.1 Limitations for Linux Driver Build 2.1.24-046

This release includes the following limitations:

- This release includes the following limitation when doing a driver injection (DUD) install. On some distributions (RHEL7.9, RHEL8.2, RHEL8.3, SLES15SP2, and SLES15SP3), the 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 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 a 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/.
 - 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
- When multiple controllers are in a system, `udev(systemd)` can timeout during `kdump/kexec` resulting in an incomplete `kdump` operation. The indication of the timeout is the following console log entry: "`scsi_hostX: error handler thread failed to spawn, error = -4`".
 - Workaround: Extend the `udev(systemd)` timeout during a `kdump` operation. Perform the following steps to increase the timeout for `udev(systemd)`:

```
vi /etc/sysconfig/kdump
add udev.event-timeout=300 to KDUMP_COMMANDLINE_APPEND
systemctl restart kdump
systemctl status kdump
```

2.2.4.3 VMware Driver Limitations

This section shows VMware driver limitations for this release.

2.2.4.3.1 Limitations for VMware Driver Build 4530.0.104

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. Lockup has only been seen within a Linux Guest VM. No known workaround at the present time.

2.2.4.4 FreeBSD Driver Limitations

This section shows FreeBSD driver limitations for this release.

2.2.4.4.1 Limitations for FreeBSD Driver Build 4410.0.1005

There are no known limitations for this release.

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 26064

This release includes the following limitations:

- Import foreign drive/logical device operation will fail to import the foreign drive/logical device when the remote master key is in ASCII format and the length is less than 32 characters.

- *Workaround:* To import the foreign drive/logical device with an ASCII format master key which has less than 32 characters length, convert the master key from ASCII format to HEX format and input the HEX value.

2.2.5.2 PLDM Limitations

This section shows the PLDM limitations for this release.

2.2.5.2.1 Limitations for PLDM Release 6.25.9.0

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.

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.

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
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/. 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 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

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.

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.

© 2023, Microchip Technology Incorporated and its subsidiaries. All Rights Reserved.

ISBN:

Quality Management System

For information regarding Microchip's Quality Management Systems, please visit www.microchip.com/quality.

Worldwide Sales and Service

AMERICAS	ASIA/PACIFIC	ASIA/PACIFIC	EUROPE
<p>Corporate Office 2355 West Chandler Blvd. Chandler, AZ 85224-6199 Tel: 480-792-7200 Fax: 480-792-7277 Technical Support: www.microchip.com/support Web Address: www.microchip.com</p> <p>Atlanta Duluth, GA Tel: 678-957-9614 Fax: 678-957-1455</p> <p>Austin, TX Tel: 512-257-3370</p> <p>Boston Westborough, MA Tel: 774-760-0087 Fax: 774-760-0088</p> <p>Chicago Itasca, IL Tel: 630-285-0071 Fax: 630-285-0075</p> <p>Dallas Addison, TX Tel: 972-818-7423 Fax: 972-818-2924</p> <p>Detroit Novi, MI Tel: 248-848-4000</p> <p>Houston, TX Tel: 281-894-5983</p> <p>Indianapolis Noblesville, IN Tel: 317-773-8323 Fax: 317-773-5453 Tel: 317-536-2380</p> <p>Los Angeles Mission Viejo, CA Tel: 949-462-9523 Fax: 949-462-9608 Tel: 951-273-7800</p> <p>Raleigh, NC Tel: 919-844-7510</p> <p>New York, NY Tel: 631-435-6000</p> <p>San Jose, CA Tel: 408-735-9110 Tel: 408-436-4270</p> <p>Canada - Toronto Tel: 905-695-1980 Fax: 905-695-2078</p>	<p>Australia - Sydney Tel: 61-2-9868-6733</p> <p>China - Beijing Tel: 86-10-8569-7000</p> <p>China - Chengdu Tel: 86-28-8665-5511</p> <p>China - Chongqing Tel: 86-23-8980-9588</p> <p>China - Dongguan Tel: 86-769-8702-9880</p> <p>China - Guangzhou Tel: 86-20-8755-8029</p> <p>China - Hangzhou Tel: 86-571-8792-8115</p> <p>China - Hong Kong SAR Tel: 852-2943-5100</p> <p>China - Nanjing Tel: 86-25-8473-2460</p> <p>China - Qingdao Tel: 86-532-8502-7355</p> <p>China - Shanghai Tel: 86-21-3326-8000</p> <p>China - Shenyang Tel: 86-24-2334-2829</p> <p>China - Shenzhen Tel: 86-755-8864-2200</p> <p>China - Suzhou Tel: 86-186-6233-1526</p> <p>China - Wuhan Tel: 86-27-5980-5300</p> <p>China - Xian Tel: 86-29-8833-7252</p> <p>China - Xiamen Tel: 86-592-2388138</p> <p>China - Zhuhai Tel: 86-756-3210040</p>	<p>India - Bangalore Tel: 91-80-3090-4444</p> <p>India - New Delhi Tel: 91-11-4160-8631</p> <p>India - Pune Tel: 91-20-4121-0141</p> <p>Japan - Osaka Tel: 81-6-6152-7160</p> <p>Japan - Tokyo Tel: 81-3-6880-3770</p> <p>Korea - Daegu Tel: 82-53-744-4301</p> <p>Korea - Seoul Tel: 82-2-554-7200</p> <p>Malaysia - Kuala Lumpur Tel: 60-3-7651-7906</p> <p>Malaysia - Penang Tel: 60-4-227-8870</p> <p>Philippines - Manila Tel: 63-2-634-9065</p> <p>Singapore Tel: 65-6334-8870</p> <p>Taiwan - Hsin Chu Tel: 886-3-577-8366</p> <p>Taiwan - Kaohsiung Tel: 886-7-213-7830</p> <p>Taiwan - Taipei Tel: 886-2-2508-8600</p> <p>Thailand - Bangkok Tel: 66-2-694-1351</p> <p>Vietnam - Ho Chi Minh Tel: 84-28-5448-2100</p>	<p>Austria - Wels Tel: 43-7242-2244-39 Fax: 43-7242-2244-393</p> <p>Denmark - Copenhagen Tel: 45-4485-5910 Fax: 45-4485-2829</p> <p>Finland - Espoo Tel: 358-9-4520-820</p> <p>France - Paris Tel: 33-1-69-53-63-20 Fax: 33-1-69-30-90-79</p> <p>Germany - Garching Tel: 49-8931-9700</p> <p>Germany - Haan Tel: 49-2129-3766400</p> <p>Germany - Heilbronn Tel: 49-7131-72400</p> <p>Germany - Karlsruhe Tel: 49-721-625370</p> <p>Germany - Munich Tel: 49-89-627-144-0 Fax: 49-89-627-144-44</p> <p>Germany - Rosenheim Tel: 49-8031-354-560</p> <p>Israel - Ra'anana Tel: 972-9-744-7705</p> <p>Italy - Milan Tel: 39-0331-742611 Fax: 39-0331-466781</p> <p>Italy - Padova Tel: 39-049-7625286</p> <p>Netherlands - Drunen Tel: 31-416-690399 Fax: 31-416-690340</p> <p>Norway - Trondheim Tel: 47-72884388</p> <p>Poland - Warsaw Tel: 48-22-3325737</p> <p>Romania - Bucharest Tel: 40-21-407-87-50</p> <p>Spain - Madrid Tel: 34-91-708-08-90 Fax: 34-91-708-08-91</p> <p>Sweden - Gothenberg Tel: 46-31-704-60-40</p> <p>Sweden - Stockholm Tel: 46-8-5090-4654</p> <p>UK - Wokingham Tel: 44-118-921-5800 Fax: 44-118-921-5820</p>