

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.....	18
3. Updating the Controller Firmware.....	22
3.1. Updating Controllers to Latest Firmware.....	22
4. Revision History.....	23
Microchip Information.....	24
The Microchip Website.....	24
Product Change Notification Service.....	24
Customer Support.....	24
Microchip Devices Code Protection Feature.....	24
Legal Notice.....	24
Trademarks.....	25
Quality Management System.....	26
Worldwide Sales and Service.....	27

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	3.2.0
Package release date	October 31, 2023
Firmware version	3.01.14.62
UEFI/Legacy BIOS	2.2.4/2.2.2
Driver versions	<p>Windows Drivers:</p> <ul style="list-style-type: none"> Windows 2022, 2019, 2016, Windows 11, 10: 1010.42.0.1020 <p>Linux SmartPQI:</p> <ul style="list-style-type: none"> RHEL 7/8/9: 2.1.18-045 SLES 12/15: 2.1.18-045 Ubuntu 18/20/21/22: 2.1.18-045 Oracle Linux 7/8: 2.1.18-045 Citrix XenServer 8: 2.1.18-045 Debian 10/11: 2.1.18-045 CentOS 7/8: 2.1.18-045 <p>VMware:</p> <ul style="list-style-type: none"> VMware ESX 7: 4330.0.116 <p>FreeBSD/Solaris:</p> <ul style="list-style-type: none"> FreeBSD 12/13: 4280.0.1007 Solaris: 11: 11.4120.0.1005
ARCCONF/maxView	B25339
PLDM	6.10.14.0

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 9.0 ¹ , 8.6 ¹ , 8.5, 8.4, 7.9, 7.8	x64
CentOS 8.4, 8.3, 8.0, 7.9, 7.8	x64
SLES 12 SP5, SP4	x64
SLES 15 SP4 ¹ , SP3, SP2	x64
Ubuntu 20.04.4, 20.04.3, 20.04, 18.04.5, 18.04.4	x64
Ubuntu 22.04, 21.04	x64
Oracle Linux 7.9 UEK6U3	x64
Oracle Linux 8.5, 8.4 UEK6U3	x64
Debian 11.1, 10.10	x64
Fedora 35 (inbox)	x64
XenServer 8.2	x64

Note:

1. New OS support—minimally tested drivers in this release. Fully supported drivers are expected in the next release.

Table 1-6. FreeBSD, Solaris, and VMware Drivers

OS	Version
ESX7.0 U3/U2	x64
FreeBSD 13, 12.3	x64
Solaris 11.4	x64

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

Description	OS	Executable
ARCCONF Command Line Utility	Windows x64 Linux x64 VMware 6.5 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 6.5 and above XenServer UEFI support	See the maxview_linux_B#####.zip, maxview_win_B#####.zip, and the maxview_vmware_B#####.zip for the installation executables.
maxView vSphere Plugin	VMware 6.5 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 3.01.14.62

This release provides the following fixes and enhancements.

- Added Passive SED support.
- Added support for logical drive transformation without supercap.
- Added ability to coalesce multiple streams of sequential IO to a single RAID logical drive.
- Added support for multi-actuator drives.
- Added support to send SCSI timestamp to drives to sync host system time (RTC).
- Updated maximum power required for controllers so that the device can operate in PCIe 25W slots.
- Added SPDM Protocol for Attestation and Authentication Credentials.
- Added display ASIC minor version.
- Fixed a performance issue with a RAID 5 logical drive writing sequential IOs that resulted in IOs taking a long time and LUN resets.
 - *Root cause:* Commands were getting stuck in queue waiting too long.
 - *Fix:* Added time check to avoid commands waiting too long.
 - *Risk:* Medium
- Fixed an issue sending a PLDM event when controller password is entered and no encrypted logical drives are present.
 - *Root cause:* Firmware was assuming an encrypted logical drive was available to send an event when the controller password is entered.
 - *Fix:* Send an event even if no encrypted logical drive is present.
 - *Risk:* Low
- Improved IO performance by speeding up internal manufacturing data command.
 - *Root cause:* A request to get manufacturing data was taking a long time.
 - *Fix:* Cache the data needed for this command so that it can be returned quicker.
 - *Risk:* Low
- Fixed an issue where logical drives may be incorrectly failed after a rebuild.
 - *Root cause:* An array consisting of multiple logical drives and multiple spare drives will undergo a rebuild when one of the physical drives has failed. Each logical drive is sequentially rebuilt using a spare drive. If some of the logical drives have completed the rebuild and a second physical drive fails, then firmware incorrectly reports the previously rebuilt logical drives as failed instead of degraded and the logical drive that is currently being rebuilt is correctly marked as failed. The firmware incorrectly failed the previously rebuilt logical drives because it removed the spare drives and the failed drive from those logical drives.
 - *Fix:* Firmware keeps the spare drives as part of the rebuilt logical drives so they will correctly report a degraded state.

- *Risk:* High
- Fixed an issue with two UBM backplanes connected to one connector and one UBM backplane incorrectly identified itself as not on a bifurcated cable.
 - *Root cause:* The logic for each connector was trying to take ownership of the drives, which resulted in controller locking up since drives should only be owned by one box.
 - *Fix:* Do not discover drives on second connector if drives have already been marked.
 - *Risk:* Low
- Fixed an issue where Real Time Clock (RTC) timestamp was not being sent to SES-based storage enclosure SEPs attached to internal connectors of the enclosure.
 - *Root cause:* Firmware was only sending the RTC timestamp to SEPs that were connected to external connectors of the controller.
 - *Fix:* Removed check for externally connected enclosure so timestamp now gets sent to all SEPs.
 - *Risk:* Low
- Fixed a potential UBM backplane flash issue where image start index is not aligned with the start sector.
 - *Root cause:* Controller was not calculating start index, and assumed the start index aligns with the start sector.
 - *Fix:* Calculate the start index correctly.
 - *Risk:* Low
- Fixed an issue where encrypted data is not accessible for a RAID 50/60 logical drive when it was failed and healed using "Heal Array".
 - *Root cause:* During the "Heal Array" operation on a failed encrypted logical drive, firmware will create a new Data Encryption Key (DEK) for the healed encrypted logical drive. This causes the data to be inaccessible because the new DEK is different from the original DEK used to encrypt the data for those physical drives.
 - *Fix:* During a "Heal Array" operation, firmware will keep the previously used DEK so the data will be accessible.
 - *Risk:* Low
- Fixed an issue where hang detection detects a hang while an IO is in the error task being spun up.
 - *Root cause:* A drive is being spun up in error task. That can take a long time.
 - *Fix:* Hang timer starts when a drive starts to be spun up. This timer was incrementing with no resets while all steps are performed. Fix was to restart this hang timer it moves to next step.
 - *Risk:* Low
- Fixed an issue where the drive slot number shown for a failed drive is wrong.
 - *Root cause:* If the SES enclosure has a valid additional status page, then the bay number for each drive in the enclosure is used by the firmware from the additional status page information. If the drive has an invalid SAS address in the additional status page, the bay number is assigned from the expander drive slot number by the firmware. When the firmware is traversing through physical drives to validate the additional status page, the validation flag is overwritten with the latest drive's additional status page validity status. This overwriting resulted in firmware not updating the bay number if the drive had an invalid SAS address in the additional status page.
 - *Fix:* Stop overwriting the validation flag of the drive's additional status page when any drive in the list has an invalid status.
 - *Risk:* Low

- Fixed an issue that when the backup power source is removed, cache status shows "temporary disabled."
 - *Root cause:* After completing the boot time cache flush, cache module cleanup has never been performed as an internal flag is not clear.
 - *Fix:* Cleanup cache module when power-up flush is completed.
 - *Risk:* Low
- Fixed continuous Predictive Failure LED blinking during an ongoing Predictive Spare Rebuild.
 - *Root cause:* When the drive is hot removed the `update_led` function is only called for the drive which was hot removed.
 - *Fix:* Update LED behavior for all drives when the drive is hot removed. This ensures the Predictive Failure LED is turned off during a rebuild operation.
 - *Risk:* Low
- Fixed an issue where an unconfigured drive still reports as StandbyOffline.
 - *Root cause:* After a Clear Controller Configuration was completed the PLDM event callback was not invoked.
 - *Fix:* After Clear Controller Configuration completes, make call to wake up the PLDM task thread.
 - *Risk:* Low
- Fixed an issue where the encrypted single drive RAID 0 logical drive was reported with state `NEED_CRYPTO_KEY` after system reboot.
 - *Root cause:* When the physical drive was replaced and the encrypted single drive RAID 0 logical drive was re-enabled, the key that is used to encrypt the data was not written back to the logical drive media. The key was used from within the controller's internal memory. When the system was rebooted, the key is no longer in the controller memory and the key is not on the logical drive either, so it is not present in the current system.
 - *Fix:* Write the key back to the media of all physical drives contained in the logical drive when the logical drive is re-enabled.
 - *Risk:* Low
- Fixed long `SATA SSD TRIM` causing hang.
 - *Root cause:* While creating a logical drive with SATA SSDs, if an Over-Provisioning Optimization (OPO) operation is performed on SSDs, no host I/Os are allowed to complete. This makes firmware assume there is a hang and firmware may trigger a lockup.
 - *Fix:* Resolved the lockup by considering OPO operations as normal I/O. Also, to speed up the OPO operation for SATA SSDs, increase the number of bytes being trimmed per operation to 512 MiB, which has been observed to reduced the OPO operation time to 22 seconds per terabyte.
 - *Risk:* Medium
- Fixed an IO freeze timeout after drive removal.
 - *Root cause:* IO freeze was performed when drive was removed. The routine that needed to handle drive fail functionality was then suspended because of IO freeze. This same routine should have handled IO unfreeze but did not since it was suspended waiting to fail the drive. This deadlock situation caused an IO freeze timeout.
 - *Fix:* Prevent the routine from saving RAID metadata since that causes it to get suspended.
 - *Risk:* Low
- Fixed an issue where firmware continuously retries a command to the SEP device that was completed with the `TASK SET FULL` status.

- *Root cause:* Firmware retries indefinitely on any command to the SEP device that was completed with TASK SET FULL status and the request's retry count was still greater than 1.
- *Fix:* Firmware decreases the request's retry count for all the retry commands to the SEP device regardless of the SCSI error status.
- *Risk:* Low
- Fixed an issue where the Predictive Fail (PF) LED did not blink during predictive spare rebuild.
 - *Root cause:* When rebuild started, only a call to blink the target rebuilding drive was called.
 - *Fix:* When rebuild starts, blink all drive LED states instead of just rebuilding the drive. PF LED will only blink if volume is fault tolerant.
 - *Risk:* Low
- Fixed an issue to set controller PHY rate to 12G during hot-plug of an expander with IO running.
 - *Root cause:* While running heavy IO to SATA drives behind an expander, bringing down PHYs and re-enabling resulted in PHYs not being set to 12G. This was a result of IOs timing out, which resulted in the controller phy rate not to be set to 12G.
 - *Fix:* Ensure the controller PHY rate is set to 12G even if the IO times out.
 - *Risk:* Low
- Fixed an issue where the product ID of an enclosure was not showing correctly.
 - *Root cause:* Product IDs that had a space in the middle of the string were being truncated.
 - *Fix:* Corrected code to ensure that any characters after the space are displayed.
 - *Risk:* Low
- Fixed an issue with a UBM backplane where not all drives were discovered.
 - *Root cause:* The backplane supported 10 DFC connectors but the firmware assumed no more than eight.
 - *Fix:* Increased the maximum DFC connectors to 15.
 - *Risk:* Low
- Fixed an issue where random sense data could be returned when the target is a host or logical drive.
 - *Root cause:* The command responsible for filling out the sense data was only filling specific fields and assumed that the rest of the fields were zero.
 - *Fix:* Zero out sense data before filling it.
 - *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 2.2.4/Legacy BIOS Build 2.2.2

This release provides the following fixes and enhancements.

- Added a new HII menu that will attempt to re-enable a previously failed logical drive whose physical drives are back online later.
- Driver health error codes are consolidated from 0x17xx and 0x18xx series to a 0x19xx series.
- Added an Unlock Controller option in the HII menu when controller password is set for controller-based encryption (CBE).
- Fixed an issue where the controller could lock up during transition from preboot to OS.

- *Root cause:* The UEFI driver sends I/O command to the controller without using a timeout value. During the transition from preboot to an OS, if an I/O command is still pending and the UEFI driver triggers a command interface change, the controller encounters a lockup.
- *Fix:* Timeout value added for I/O commands.
- *Risk:* Low
- Fixed an issue where there was an incorrect location representation for non-disk devices such as SES and expander devices.
 - *Root cause:* Location represented as `port:box:bay` for non-disk devices. Bay number is not applicable for non-disk devices.
 - *Fix:* Location for non-disk devices updated to only show `port:box` format.
 - *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 1010.42.0.1020

This release provides the following fixes and enhancements.

- Added support for Multi-Actuator drives.
- Added driver internal ring buffer logging that allows the driver to log important messages to a driver allocated ring buffer memory.
- Fixed an issue where driver accesses Command Descriptor Block's (CDB) NULL pointer and BSOD occurs.
 - *Root cause:* The SmartPQI driver parses the SRB with CDB length = 0 and sets the CDB pointer to NULL. When the driver accesses the NULL CDB pointer an invalid memory access occurs.
 - *Fix:* The build SCSI request in the driver's SRB routine will return invalid command status when the CDB pointer is NULL.
 - *Risk:* Low
- Fixed an issue where driver might fail to load intermittently after a dirty system shutdown.
 - *Root cause:* There are two problems that can occur after a dirty system shutdown.
 - The first is that the driver can get confused about which mode the controller is in and request an unnecessary soft reset.
 - The second problem is that after requesting the unnecessary soft reset, the driver can access controller registers prematurely before the controller has completed the soft reset. This can result in the driver misinterpreting the state of the controller firmware.
 - *Fix:* Added additional checks to the driver initialization logic to make its Controller mode detection more robust and to prevent misinterpreting controller registers when the controller firmware is not fully up and running.
 - *Risk:* Low
- Fixed a BSOD due to driver and controller hardware not supporting greater than 16 byte CDB.
 - *Root cause:* BSOD is caused by driver copying greater than 16 bytes to the IOBypass request, thus overwriting critical error index field. Therefore, the IOBypass error index field used is out of range and causes the driver to access a bad address.
 - *Fix:* Added check in build I/O command path to check for CDBs greater than 16 byte and if found, reject the command at the driver level as an invalid command or a command that is not supported.

- *Risk:* Low

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 2.1.18-045

This release includes the following fixes and enhancements.

- Added support for Multi-Actuator disk drives.
- Added support for displaying controller firmware version in the OS message log. The controller firmware version is printed to OS message log during driver initialization.
- Added support for a controller ready timeout module parameter (`ctrl_ready_timeout`). The valid range is 0 or 30–1800 seconds. The default value is 0, which causes the driver to use a timeout of 180 seconds (3 minutes).
- Added support for deleting a LUN via sysfs using the following syntax:


```
echo 1 > /sys/block/sdX/device/delete
```
- Added module parameter to disable managed interrupts (`disable_managed_interrupts=1`).
- Fixed a race condition where the driver can access the RAID map when using IOBypass during a RAID configuration change.
 - *Root cause:* A race condition in the driver might cause it to access a stale RAID map when a logical drive is reconfigured.
 - *Fix:* Modified the driver logic to
 - Invalidate a RAID map at an early stage when a RAID configuration change is detected
 - Switch to a new RAID map only after the driver detects that the RAID map has changed
 - *Risk:* Low
- Fixed an issue where the `sg_map` tool issues `SCSI READ BLOCK LIMITS (0x5)` command and the firmware never completes it, causing a system call trace and `sg_map` hang.
 - *Root cause:* Driver is sending an incorrect data direction flag for the RAID path request.
 - *Fix:* Corrected the data direction flag for the RAID path request.
 - *Risk:* Low
- Fixed an issue where PQI Reset might fail with an error “– 6” if firmware takes more than 100 ms to complete Reset.
 - *Root cause:* Method used by the driver to detect controller firmware crash during PQI Reset was incorrect in some cases.
 - *Fix:* Changed method used by the driver to detect controller firmware crash during PQI Reset.
 - *Risk:* Low

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 4330.0.116

This release provides the following fixes and enhancements.

- Added support for Multi-Actuator disk drives.
- Added support for multiple tag table to improve performance by optimizing tag and requesting structure allocation.
- Added support for `ScsiAdapterCheckTarget` which performs a device lookup and returns true only when specified adapter/channel/target exists and is exposed as a SCSI device.
- Added a module parameter `CtrlReadyTimeoutSecs` for controller ready timeout. The valid range is 30–1800 seconds. The default value is 120 seconds.

- Fixed an issue where the `SCSI READ BLOCK LIMITS (0x5)` command is never completed by firmware and a `TMF ABORT` is issued in ESXi.
 - *Root cause:* Driver is sending an incorrect data direction flag for the RAID path request.
 - *Fix:* Corrected the data direction flag for the RAID path request.
 - *Risk:* Low
- Fixed an issue where during driver initialization a warning message appears about the DMA alignment setting, instead of an informational message.
 - *Root cause:* Improper flag on the message.
 - *Fix:* Change message level `WARN` to `INFO`.
 - *Risk:* Low
- Fixed an issue where PSOD is observed during driver load.
 - *Root cause:* Driver was creating a greater number of completion worlds than allowed, which resulted in a SCSI layer deadlock.
 - *Fix:* Set the SCSI completion world to the maximum supported value.
 - *Risk:* Medium
- Fixed an issue with cards that are running on older firmware where firmware features previously were potentially disabled due to incorrect placement in the firmware features table.
 - *Root cause:* Firmware feature parsing logic on the driver's side was incorrectly skipping some features that must have been enabled regardless of firmware's maximum known features. Firmware does not have the maximum known feature bit set, and the code stops evaluating the feature list when it finds this issue.
 - *Fix:* Lean on already implemented method of traversing valid feature entries using the `num_elements` field in the PQI firmware feature table, and reject support for any bit positions outside the valid bytes indicated by firmware.
 - *Risk:* Low
- Fixed an issue where the driver reports an error when the unsupported `SCSI Maintenance IN (0xA3)` command with service action "report supported opcode" (`0xC`) is sent to the logical drive.
 - *Root cause:* Firmware does not support command `0xA3` with service action `0xC` and it returns an error to the driver.
 - *Fix:* Suppress error logs reported for `SCSI Maintenance IN (0xA3)` with service action `0xC`.
 - *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 4280.0.1007

This release includes the following fixes and enhancements.

- Fixed an issue where logical drives created on the hard disk drives were not listed with the "geom disk list" command.
 - *Root cause:* The `bus_dmamap_sync` operation which is performed before device access was being handled incorrectly.
 - *Fix:* Handle the `bus_dmamap_sync` operation correctly.
 - *Risk:* Low
- Fixed an issue of the expanded capacity of a logical drive not reflecting in the OS after modifying the size using `Arcconf`.
 - *Root cause:* Current FreeBSD code detects a capacity expansion by checking the logical drive status.

Firmware may cycle through the logical drive states (OK, NEEDS_EXPAND, EXPANDING, OK) faster than the driver can send commands to check for the logical drive state changes. This results in the driver seeing the logical drive as OK and failing to detect the logical drive expansion.

- *Fix:* Driver detects the event type PQI_EVENT_TYPE_LOGICAL_DEVICE (0x5) due to logical drive expansion and sets the rescan flag. Later, it triggers reprobe of all logical drives that includes sending SCSI READ CAPACITY command and updating the size that is visible to the OS.
- *Risk:* Low
- Fixed an issue where OS commands are getting hung after executing on logical drives that are created using SSD.
 - *Root cause:* Driver was not properly handling the data direction flag for IOBypass.
 - *Fix:* Corrected the data direction flag for IOBypass request.
 - *Risk:* Low
- Fixed an issue where the SCSI READ BLOCK LIMITS (0x5) command is never completed by firmware and a TMF ABORT is observed.
 - *Root cause:* Driver is sending an incorrect data direction flag for the RAID path request.
 - *Fix:* Corrected the data direction flag for the RAID path request.
 - *Risk:* Low

2.1.3.4.2 Fixes and Enhancements for Solaris Driver Build 11.4120.0.1005

There are no fixes and enhancements for this version.

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 25339

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 25335

This release provides the following fixes and enhancements:

- Removed the Log4J library usage completely from maxView.
 - Note:** Microchip strongly recommends users of maxView update to the latest version of the tool to avoid the security vulnerabilities with the previous releases.
- Added TASK and GETSTATUS commands support in UEFI ARCCONF.
- Added ASIC minor version display in maxView and ARCCONF.
- Added support in maxView and ARCCONF to sort the controllers based on the bus number, if the slot ID is not valid.
- Fixed an issue where maxView create logical drive selection check box was invisible in older version of the Edge browsers.
 - *Root cause:* maxView create logical drive selection table had alignment issues in older version of the Edge browser when the column width was specified in percentage.
 - *Fix:* The column width of the table is changed to pixel instead of percentage to resolve the alignment issue in the older edge browser.
 - *Risk:* Low

- Fixed an issue where maxView was failing to modify the maxCache logical device's cache policy to write-through in ESXi 7.xx.
 - *Root cause:* ESXi Redfish server had issues in accepting the HTTP PATCH request from maxView.
 - *Fix:* Updated the maxView HTTP PATCH request call to comply with ESXi redfish server.
 - *Risk:* Low
- Fixed an issue where ARCCONF was not displaying the controller PHY error log information.
 - *Root cause:* PHY error log was disabled in ARCCONF.
 - *Fix:* Added changes to enable the PHY error log from ARCCONF.
 - *Risk:* Low
- Fixed an issue where ARCCONF results in segmentation Fault when collecting the support archive.
 - *Root cause:* ARCCONF resulted in segmentation Fault when retrieving invalid vendor specific diagnostic page from the drive.
 - *Fix:* Added changes in ARCCONF to skip reading the invalid vendor specific diagnostic page.
 - *Risk:* Low
- Fixed an issue where ARCCONF was not displaying the updated expander firmware version after upgrading the expander firmware.
 - *Root cause:* The controller is returning the older expander firmware version instead of recently updated expander firmware version until the next controller power cycle.
 - *Fix:* Added changes in ARCCONF to send the SCSI inquiry command to the expander for retrieving the updated expander firmware version.
 - *Risk:* Low
- Fixed an issue of I²C devices not being detected when "SMBUSCHANNEL" is set to "ENABLE".
 - *Root cause:* ARCCONF disabled the SMBUSCHANNEL when user enabled it.
 - *Fix:* Added changes in ARCCONF to set the correct user input value for SMBUSCHANNEL.
 - *Risk:* Low
- Fixed an issue where ARCCONF fails to set ATAPASSWORD for a SATA device.
 - *Root cause:* ARCCONF ATAPASSWORD command was disabled.
 - *Fix:* Added changes in ARCCONF to enable the ATAPASSWORD command.
 - *Risk:* Low
- Fixed an issue where ARCCONF did not allow changing the write cache setting to write-through for RAID 5 maxCache logical device.
 - *Root cause:* Write-through setting for maxCache write cache setting was aborted due to invalid condition check.
 - *Fix:* Added changes to remove the invalid condition check and allow write-through setting for maxCache logical device.
 - *Risk:* Low
- Fixed an issue where in UEFI ARCCONF where remove hot spare command was failing when array ID was specified.
 - *Root cause:* Parsing of array ID in the UEFI ARCCONF remove hot spare command was not proper.
 - *Fix:* Added changes to fix the array ID parsing issue.
 - *Risk:* Low

- Fixed an issue in ARCCONF where an invalid SMART attribute was displayed in the drive SMART STATS.
 - *Root cause:* ARCCONF was retrieving invalid attributes beyond the SMART attributes range.
 - *Fix:* Added changes in ARCCONF to block the retrieval of invalid attributes beyond the SMART attributes range.
 - *Risk:* Low
- Fixed an issue where ARCCONF was not displaying the connector information for the enclosure device.
 - *Root cause:* Connector ID for the enclosure device was missing.
 - *Fix:* Added changes to display the connector ID for the enclosure device in ARCCONF.
 - *Risk:* Low
- Fixed an issue where maxView was displaying invalid supercap alert message for the controller that does not support Supercap.
 - *Root cause:* Supercap alert message was displayed for the controller that does not support Supercap.
 - *Fix:* Added changes to display the Supercap alert message only for controllers that has Supercap support.
 - *Risk:* Low
- Fixed an issue where ARCCONF failed to execute SLOTCONFIG sub-command.
 - *Root cause:* ARCCONF SLOTCONFIG sub-command was disabled.
 - *Fix:* Added changes to enable SLOTCONFIG sub-command from ARCCONF.
 - *Risk:* Low
- Fixed an issue where UEFI ARCCONF was not displaying array information for physical device.
 - *Root cause:* Array property on the physical device was not implemented in the UEFI ARCCONF
 - *Fix:* Added changes to display the array property in GETCONFIG command for the physical device.
 - *Risk:* Low
- Fixed an issue where ARCCONF was not displaying the negative temperatures for certain drive models.
 - *Root cause:* ARCCONF failed to decode the negative temperature for the specific drives.
 - *Fix:* Added changes to display the correct negative temperature for the supported drives.
 - *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 6.10.14.0

This release provides the following fixes and enhancements.

- Added support for the following Redfish ACTION requests:
 - Drive.#SecureErase
 - Drive.#Reset
 - Storage.#ResetToDefaults

Note: ResetToDefaults does not support a ResetType of ResetAll when encrypted volumes exist on the controller. The user must first either delete or decrypt any encrypted volumes prior to issuing such an ACTION request.
- Added support for Redfish PATCH requests for the following properties:

- Volume.DisplayName
- Volume.Links.DedicatedSpareDrives
- Volume.IOPerfModeEnabled
- Volume.ReadCachePolicy
- Volume.WriteCachePolicy
- Drive.LocationIndicatorActive
- Drive.WriteCacheEnabled
- StorageController.ControllerRates.ConsistencyCheckRatePercent
- StorageController.ControllerRates.RebuildRatePercent
- StorageController.ControllerRates.TransformationRatePercent
- Added the following Redfish alerts:
 - DriveOffline
 - DriveMissing
 - DriveMissingCleared
 - DriveOfflineCleared
 - VolumeOffline
 - VolumeOfflineCleared
 - BatteryMissing
 - BatteryFailure
 - BatteryCharging
 - BatteryOK
 - ControllerDegraded
 - ControllerFailure
 - ControllerPreviousFailure
 - ControllerPasswordRequired
 - ControllerPasswordEntered
- Added the following property to the Redfish GET response for VolumeCollection resources:
 - VolumeCollection.@Redfish.CollectionCapabilities.MaxMembers
- Added support for firmware updates for physical drives.
- Added a new OperationName value of 'Reverting' which is used for Redfish GET responses targeting self-encrypting drives undergoing a revert.
- Added a new EncryptionType of 'NativeDriveEncryption' for logical drives which are secured using SED-based encryption.
- When targeting the physical drives temperature numeric sensor, GetSensorReading will now return the temperature of the hottest drive instead of a normalized temperature.
- Certain controller temperature sensor numeric sensors have had their EntityType changed from "I/O Controller" to "Add-in card".
- RDE READ on a Drive resource will now omit the Vendor from Drive.Name property on some controllers.
- The Type 5 commands QueryDownstreamDevices, GetDownstreamFirmwareParameters, and QueryDownstreamIdentifiers will now report information for physical drives.
- The Volume.WriteCachePolicy property in a Redfish GET response for a Volume resource will have a value of ProtectedWriteBack when the controller battery is removed or goes missing.

- Updated all resource schema dictionaries to the latest version available in the 2021.4 schema bundle. The individual dictionary versions are:
 - Annotations: v1.1.1
 - Drive: v1.14.0
 - Event: v1.7.0
 - Port: v1.6.0
 - PortCollection: unversioned
 - RedfishError: v1.0.1
 - Storage: v1.12.0
 - StorageController: v1.5.0
 - StorageControllerCollection: unversioned
 - Volume: v1.6.2
 - VolumeCollection: unversioned
- On controllers that support managed SED encryption:
 - Redfish GET responses for a self-encrypting drive resource will publish the following 'EncryptionStatus' values:
 - Unencrypted
 - Locked
 - Unlocked
 - Foreign
 - 'Drive.Status.State' for a self-encrypting drive (SED) resource will be set to 'StandbyOffline' in the following conditions:
 - SED is Foreign
 - SED is Locked (only for controller owned SEDs)
 - SED is controller owned and controller is waiting on SED adapter password.
 - 'Encrypted' property will be set to True on Redfish GET responses for Volume resources which are secured using SED-based encryption.
- Fixed and issue where events would be sent continuously if the host does not respond to PlatformEventMessage.
 - *Root cause:* In case of an async event receiver, there was no logic to cap the number of retries when an event subscriber never responds to a PlatformEventMessage request.
 - *Fix:* Added logic to cap the number of retries to 3. After the max limit of retries is reached, no further events will be sent until new events are pushed in or event subscriber resets the event queue.
 - *Risk:* Medium
- Fixed and issue where Redfish CREATE requests for logical drives fails when using 4Kn drives as the data drives.
 - *Root cause:* The API which creates a logical drive with 4Kn drives returned an error causing the CREATE operation to fail.
 - *Fix:* Updated the API to create a logical drive on 4Kn drives correctly.
 - *Risk:* Low
- Fixed and issue where Storage.Status.HealthRollup would erroneously report a value of Warning instead of OK when a logical drive has an InitializationMethod of Foreground.

- *Root cause:* The calculation of the controller's health rollup considers a logical drive undergoing foreground initialization to have a health value of 'Warning' as it is unavailable to the host OS during that time.
- *Fix:* Revised the calculation of the controller health rollup to publish a value of 'OK' if a child logical drive is queued for or undergoing foreground initialization and no other factors would cause a health status other than 'OK'.
- *Risk:* Low
- Fixed and issue where the controller firmware update progress is underestimated.
 - *Root cause:* On controllers that require firmware to be flashed twice, the percent complete stalls at 95% after the first flash and doesn't update until the second flash completes.
 - *Fix:* The number of flashes required is now checked when initializing the percent complete increment.
 - *Risk:* Low
- Fixed and issue where StorageController.CacheSummary.Status.State would have a value of Enabled on controllers lacking a cache module.
 - *Root cause:* The mechanism for determining if a cache module was attached was not checking the right fields in the right conditions.
 - *Fix:* Changed the code to correctly determine the presence of a cache module and correctly report it. During the investigation of this issue it was also determined that other fields in the cache summary also needed to be changed.
If no cache module is attached, the following fields should be as shown here:
 - TotalCacheSizeMiB: 0
 - PersistentCacheSizeMiB: 0
 - Status.Health: OK
 - Status.State: Absent
 - *Fix:* Changed the code to correctly determine the presence of a cache module and correctly report it. During the investigation of this issue it was also determined that other fields in the cache summary also needed to be changed.
If the controller has no cache module but supports RAID modes, the following fields should be as shown here:
 - TotalCacheSizeMiB: The total memory size of the controller
 - PersistentCacheSizeMiB: 0
 - Status.Health: Warning
 - Status.State: Disabled
 - *Risk:* Low
- Fixed and issue where a Port's ServiceLabel did not contain the parent StorageController's slot number.
 - *Root cause:* The Port READ response function was duplicating the Port Name property when publishing the ServiceLabel property.
 - *Fix:* Revised the ServiceLabel value string to contain both the controller slot number and port name in the format "Slot=x:Port=y".
 - *Risk:* Low
- Fixed and issue where an incorrect completion code is sent for Redfish requests which encounter an error.
 - *Root cause:* RDEOperationInit requests which encounters an error returns ERROR_OPERATION_FAILED. This does not conform to the DMTF PLDM Type 6 spec DSP0218.
 - *Fix:* Modified the processing of RDEOperationInit requests to conform to the DMTF PLDM Type 6 spec DSP0218, where any Type 6 operation which encounters an error

and responds with extended error info should return ERROR_UNSUPPORTED instead of ERROR_OPERATION_FAILED.

- *Risk:* Low
- Fixed an issue where an incorrect strip size is applied when creating logical drives resources on 4Kn drives.
 - *Root cause:* The calculation of the strip size in blocks used when creating the logical drive assumed a physical drive block size of 512 bytes.
 - *Fix:* Modified the logical drive creation encoder logic to use the actual physical drive block size for the array instead of a hard coded value.
 - *Risk:* Low
- Fixed an issue where the Drive.Identifiers.DurableName value for NVMe drives did not conform to the standards regular expression.
 - *Root cause:* The code to separate DurableName for NVMe drives with a colon was not implemented.
 - *Fix:* Modified the DurableName string for NVMe drives to separate each pair of characters with a colon.
 - *Risk:* Low
- Fixed an issue where Volume CREATE requests succeeded with malformed data drive @odata.id strings.
 - *Root cause:* The validation of the @odata.id string did not enforce that the string contained at least one instance of the substring "/Drives/".
 - *Fix:* Modified the @odata.id string validation to enforce inclusion of the "/Drives/" substring prior to the drive collection index number.
 - *Risk:* Low
- Fixed an issue where the Volume.Identifiers.DurableName value would not conform to the standards regular expression.
 - *Root cause:* The code to appropriately format the DurableName string was not implemented.
 - *Fix:* Added hyphens to the DurableName string for Volumes as needed to meet the spec requirement.
 - *Risk:* Low

2.2 Limitations

This section shows the limitations for this release.

2.2.1 Firmware Limitations

This section shows the firmware limitations for this release.

2.2.1.1 Limitations for Firmware Release 3.01.14.62

There are no known limitations for this release.

2.2.2 UEFI/Legacy BIOS Limitations

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

2.2.2.1 Limitations for UEFI Build 2.2.4/Legacy BIOS Build 2.2.2

There are no known limitations for this release.

2.2.3 Driver Limitations

This section shows the driver limitations for this release.

2.2.3.1 Windows Driver Limitations

This section shows the Windows driver limitations for this release.

2.2.3.1.1 Limitations for Windows Driver Build 1010.42.0.1020

This release includes the following limitations:

- The following are the limitations of Multi-Actuator:
 - Supports only
 - HBA drive
 - Direct-Attached
 - Windows/Linux/VMware
 - Intel/AMD
 - UEFI mode (for multi-LUN display)
 - No Storage Manager support
 - No boot support

2.2.3.2 Linux Driver Limitations

This section shows the Linux driver limitations for this release.

2.2.3.2.1 Limitations for Linux Driver Build 2.1.18.045

This release includes the following limitations:

- On AMD/RHEL 7.9 systems, the system might panic due to a bug in the IOMMU module. For details, see <https://lore.kernel.org>
 - Workaround: Disable the IOMMU setting option in BIOS.
- The following are the limitations of Multi-Actuator:
 - Supports only
 - HBA drive
 - Direct-Attached
 - Windows/Linux/VMware
 - Intel/AMD
 - UEFI mode (for multi-LUN display)
 - No Storage Manager support
 - No boot support
- On AMD/UEK6 systems, the system might hang during kdump if IOMMU is enabled.
 - Workaround: Disable IOMMU setting option in BIOS.
- 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. Reported on RHEL 8.5, 8.6 and 9.0.
 - Workaround: Load OS from USB device instead of virtual media. Load OS from virtual media but initiation ISO verification (media test) during install followed by ESC to cancel media test.
- This release includes the following limitation when doing a driver injection (DUD) install. On some distributions (RHEL7.9, RHEL8.2, RHEL8.3, SLES15SP2, 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:
 - i. Make sure the Write Cache is disabled for any attached drive.
 - ii. For RHEL7.9/8.2/8.3, add `rd.driver.blacklist=smartpqi` to the grub entry along with `inst.dd`.

- Depending on hardware configurations, the smartpqi expose_id_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 usual indication of the timeout is the console log entry: "scsi_hostX: error handler thread failed to spawn, error = -4".
 - Workaround: Extend the udev(systemd) timeout during a kdump operation. The steps to increase the timeout for udev(systemd) are:
 - i. vi /etc/sysconfig/kdump
 - ii. add udev.event-timeout=300 to KDUMP_COMMANDLINE_APPEND
 - iii. systemctl restart kdump
 - iv. systemctl status kdump

2.2.3.3 VMware Driver Limitations

This section shows VMware driver limitations for this release.

2.2.3.3.1 Limitations for VMware Driver Build 4330.0.116

This release includes the following limitations:

- The following are the limitations of Multi-Actuator:
 - Supports only
 - HBA drive
 - Direct-Attached
 - Windows/Linux/VMware
 - Intel/AMD
 - UEFI mode (for multi-LUN display)
 - No Storage Manager support
 - No boot support

2.2.3.4 FreeBSD Driver Limitations

This section shows FreeBSD driver limitations for this release.

2.2.3.4.1 Limitations for FreeBSD Driver Build 4280.0.1007

This release includes the following limitations.

- FreeBSD OS installation on a NVMe physical device is not supported in this release.

2.2.3.4.2 Limitations for Solaris Driver Build 11.4120.0.1005

There are no known limitations for this release.

2.2.4 Management Software Limitations

This section shows management software limitations for this release.

2.2.4.1 maxView Storage Manager/ARCCONF Limitations

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

2.2.4.1.1 Limitations for maxView Storage Manager/ARCCONF Build 25335

This release includes the following limitations.

- With maxView B25335 build, user will not be able to modify the direct attached cable settings.
 - *Workaround:* Use the ARCCONF CLI tool as an alternate option or use the previous version of maxView.

2.2.4.2 PLDM Limitations

This section shows the PLDM limitations for this release.

2.2.4.2.1 Limitations for PLDM Release 6.10.14.0

This release includes the following PLDM limitations:

- Action `Storage.ResetToDefault` with a `ResetType` of 'ResetAll' is not supported when the controller has logical drives which are encrypted.
 - Workaround: None
- RDE update on `Drive.WriteCacheEnabled` is unsupported for physical drives that are part of a logical drive.
 - Workaround: None

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
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>