

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.....	17
3. Updating the Controller Firmware.....	21
3.1. Updating Controllers to Latest Firmware.....	21
4. Revision History.....	22
Microchip Information.....	23
The Microchip Website.....	23
Product Change Notification Service.....	23
Customer Support.....	23
Microchip Devices Code Protection Feature.....	23
Legal Notice.....	23
Trademarks.....	24
Quality Management System.....	25
Worldwide Sales and Service.....	26

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	3.3.4
Package release date	February 29, 2024
Firmware version	03.01.28.82
UEFI/Legacy BIOS	2.12.1/2.12.3
Driver versions	<p>Windows Drivers:</p> <ul style="list-style-type: none"> Windows 2022, 2019, Windows 11, 10: 1010.96.0.1007 <p>Linux SmartPQI:</p> <ul style="list-style-type: none"> Rocky Linux 9: 2.1.28-025 RHEL 7/8/9: 2.1.28-025 SLES 12/15: 2.1.28-025 Ubuntu 20/22: 2.1.28-025 Oracle Linux 7/8/9: 2.1.28-025 Citrix Xenserver 8: 2.1.28-025 Debian 10/11/12: 2.1.28-025 <p>VMware:</p> <ul style="list-style-type: none"> VMware ESX 7.0/8.0: 4662.0.112 <p>FreeBSD:</p> <ul style="list-style-type: none"> FreeBSD 14/13: 4500.0.1024
ARCCONF/maxView	4.17.00.26540
PLDM	6.35.8.0

1.2 Files Included in this Release

This section details the files included in this release.

Table 1-2. Firmware Files

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

Table 1-3. Firmware Programming Tools

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

Driver Files

Table 1-4. Windows Drivers

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

Table 1-5. Linux Drivers

OS	Version
RHEL 9.3, 9.2, 9.1, 8.9, 8.8, 8.7, 7.9	x64
SLES 12 SP5	x64
SLES 15 SP5, SP4	x64
Ubuntu 20.04.6, 20.04.5, 20.04	x64
Ubuntu 22.04.3, 22.04.2, 22.04	x64
Oracle Linux 7.9 UEK6U3	x64
Oracle Linux 9.3, 9.2, 8.9, 8.8, 8.7, UEK7U2	x64
Debian 12.2, 11.8, 10.13	x64
Fedora 39 (inbox)	x64
Citrix XenServer 8.2.1	x64
Rocky Linux 9.3, 9.2	x64

Table 1-6. FreeBSD and VMware Drivers

OS	Version
ESX 8.0 U2/U1, 7.0 U3/U2	x64
FreeBSD 14.0, 13.2	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 arccnf_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 03.01.28.82

This release includes the following fixes and enhancements:

- Added support for ATA NCQ Priority which allows host application the option of sending ATA NCQ IO with priority.
- Added support for saving controller logs in host memory in the event of a system crash.
- Fixed an issue of host tools hang deleting a logical volume in an array having multiple secured volumes configured using foreign MCHP-owned SED drives.
 - *Root cause:* This is due to firmware changing the SED ownership of all the physical drives in the array to 'Otherwise Owned' when a logical volume is deleted in the array. This causes failure in the metadata update to the physical drives. Due to the SED ownership change, the physical drives are considered to be HBA access only and hence metadata is not saved. But, firmware waits for metadata update to complete in a loop and the tool's delete command never completes. This hangs the host application.
 - *Fix:* Fixed by changing the SED ownership to 'Otherwise Owned' only when the last logical volume in the array is deleted.
 - *Risk:* Low
- Fixed an issue where "Fatal Drive Error/IO Fatal Error" event unexpectedly observed in healthy configuration.
 - *Root cause:* When a logical request is sent to a drive and is completed with error status, firmware always posts "Fatal Drive Error/IO Fatal Error" event to the host. For example, the SCSI-ATA pass-through command which requests ATA information to be returned.
 - *Fix:* "Fatal Drive Error/IO Fatal Error" event should only be posted to the host only if the command is either a read or a write command.
 - *Risk:* Low
- Fixed an issue where controller may lockup when processing Out-Of-Band requests.
 - *Root cause:* When multiple Fragment Type MCTP requests are sent to the firmware with the same session ID, it may cause the wrong internal data structure to be allocated for the requests for processing. These internal data structures are indexed based on the session ID. Since two requests are sent with same session ID, this can cause the firmware to pick the wrong internal data structure and can result in a lockup.
 - *Fix:* If a new request is received with the session ID of an existing Out-Of-Band request, delete the internal data structure belonging to the existing request if it is stuck for more than five minutes and use that session ID for the newly received request. Otherwise return 'Session Already Exists' error for the newly received request. The host application should retry the request.
 - *Risk:* Low
- Fixed an issue where deleted foreign volume comes back after controller reboot.
 - *Root cause:* The foreign SEDs in the foreign volume have been set to otherwise owned before the volume metadata on the drive is cleaned up. That causes the SED's metadata not to be

- cleared. After a reboot, firmware could read the volume metadata again from the drives and the volume shows up again.
- *Fix:* Set the foreign SEDs in the foreign volume to otherwise owned after the metadata has been cleared.
 - *Risk:* Low
- Fixed a LED control failure upon hot insertion of the HGST H4060-J SAS external enclosure.
 - *Root cause:* The HGST H4060-J SAS external enclosure consists of two IO modules, and each IO module consists of multiple expanders, three in this case, with a single SEP. Upon hot insertion of an IO module (either primary or secondary path), firmware attempts to enumerate the "bay/slot index" upon detection of the SES device. However, this process is only successful, if and only if, all expanders, within the enclosure or IO module, are detected. Firmware fails to enumerate bay/slot index which leads to LED control failure.
 - *Fix:* Upon drive hot insertion and if drive is attached to a fan-out enclosure, ensure the bay/slot index is updated properly.
 - *Risk:* Low
 - Fixed an issue where the host is observing BSOD while clearing the controller configuration.
 - *Root cause:* Clearing the controller configuration with a large number of logical drives is a time-consuming process. During this long process, the host observed outstanding commands on a few logical drives and unconfigured physical drives and issued device resets. Firmware is completing the device resets but failed to send completion to the device driver for unconfigured physical drives as they don't have valid logical drive numbers associated with them, which resulted in the host OS triggering a BSOD.
 - *Fix:* Treat the unconfigured physical drive's max passthrough logical drive number as a valid logical drive number in device reset completion code path.
 - *Risk:* Medium
 - Fixed an issue where a foreign drive is getting listed in tools for logical drive creation.
 - *Root cause:* When a foreign MSED logical drive is deleted, the firmware will mark the physical SEDs that are part of the foreign logical drive as foreign SEDs. However, the firmware will take a max of one minute time to unregister these foreign SEDs from the tools. During this interval, the foreign SEDs will be listed in tools for logical drive creation.
 - *Fix:* The firmware will unregister the foreign SEDs immediately after the logical drive deletion.
 - *Risk:* Low
 - Fixed an issue where a LUN reset is waiting for completion on a failed physical drive.
 - *Root cause:* When the host issues a LUN reset to the physical drive, the firmware will mark the RESET BUSY flag on the physical drive. Due to persistent timeouts from the physical drive, the firmware will fail the physical drive but will fail to clear the RESET BUSY flag on the failed physical drive. Due to this flag, even though firmware completed the LUN reset, it will be waiting for this flag to clear before sending the completion to the host.
 - *Fix:* Clear the RESET BUSY flag while failing the physical drive.
 - *Risk:* Low
 - Fixed an issue where the logical drive under transformation is moving to a failed state after reboot.
 - *Root cause:* When the host initiates a transformation on a logical drive, the firmware may initiate the transformation on multiple logical drives to service the host request. While doing so, the firmware will maintain two copies of RAID metadata (Old and New) for the transforming logical drive. During the transformation, if the host is rebooted, during bootup, the firmware will start loading the RAID metadata for all logical drives. As part of this process, due to incorrect logic, firmware wrongly mapped the Old RAID metadata of the transforming

- logical drive to an unused logical drive index and discarded the RAID metadata. After this, when the firmware goes to resume the transformation, it observes the empty Old RAID metadata, fails the transformation, and the logical drive is moved to the failed state.
- *Fix:* Corrected the Old and New RAID metadata mapping during the firmware load configurations.
 - *Risk:* Low
- Fixed an issue where the managed SED logical drive moved to the `LOG_VOL_SED_DATA_LOCKED` state after inserting a foreign SED during Rapid Parity Initialization (RPI).
 - *Root cause:* When a foreign SED is inserted into the managed SED logical drive, the firmware will mark the foreign SED as `WRONG_REPLACED`. But when RPI is getting started, firmware marks all the physical drives as the `OK` state in the respective logical drive's RAID metadata without any checks. This resulted in a physical drive-state mismatch between global and logical drive metadata and firmware locked the managed SED logical drive and moved it to the `LOG_VOL_SED_DATA_LOCKED` state.
 - *Fix:* Prevented the initiation of RPI on the foreign SED. Added support to trigger RPI on foreign SED once it is imported.
 - *Risk:* Low
 - Fixed an issue where a possible controller lockup was observed after a factory reset.
 - *Root cause:* The controller firmware will allocate 3MiB of memory for storing Serial Output Buffer (SOB) logging and its metadata. During the factory reset, firmware wrongly updated the whole 3MiB memory as the size of SOB logging. Due to this, firmware tried accessing the memory beyond its region and ended up in lockup.
 - *Fix:* Update the correct SOB memory size after reducing the SOB metadata size from 3MiB in the factory reset code.
 - *Risk:* Low
 - Fixed a possible lockup when controller cache is being enabled and I/O is running.
 - *Root cause:* There is a small window during the cache enabling process where a firmware I/O completion process can access the cache before the cache enabling process completes, which may result in a lockup.
 - *Fix:* Fixed firmware logic to make sure the cache enabling process completes before the firmware I/O completion process can access the cache.
 - *Risk:* Low
 - Fixed an issue where an enclosure with drives powered off was unable to be powered on/ discovered by HBA.
 - *Root cause:* SES control page clearing device off bit was not sent to enclosure.
 - *Fix:* Fixed logic to make sure SES control page is sent to enclosure resulting in HBA being able to discover drives.
 - *Risk:* Low
 - Improved Raid 1/10 write sequential performance when DDR cache is on and writes are not aligned to stripe.
 - *Root cause:* Writes from cache were not stripe aligned, making them inefficient.
 - *Fix:* Ensure writes are stripe aligned.
 - *Risk:* Medium
 - Fixed an issue to remove encryption test failed event on boards that do not support encryption.
 - *Root cause:* Encryption test always runs on all boards and result was reported.
 - *Fix:* Do not log event if board does not support encryption.

- *Risk:* Low
- Fixed an issue where supercap is stuck in charging state indefinitely after charge timeout is exceeded.
 - *Root cause:* Status of backup power supply was not being updated when timeout event occurred.
 - *Fix:* When charge timeout event occurs, status of backup power supply is updated.
 - *Risk:* Low
- Fixed an issue where the controller may fail during SPDM authentication.
 - *Root cause:* SPDM authentication fails due to extra padded bytes caused by incorrect reading of CA certificate length.
 - *Fix:* Firmware will read the correct CA certificate length from the signed CA certificate.
 - *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 2.12.1/Legacy BIOS Build 2.12.3

This release includes the following fixes and enhancements:

- Added support to display UBM Backplane firmware information in both Decimal and Hexadecimal format.
- Added support of status menu for driver health ignore option.
- Fixed an issue where the driver health status is shown as failed even after selecting ignore driver health state.
 - *Root cause:* Driver health ignore state was considered only at the controller level not for the entire driver.
 - *Fix:* Flag to detect ignored driver health state considered for both controller and driver level.
 - *Risk:* Low.
- Fixed an issue where setting the value of Parallel Surface Scan Count to four is failing with error as no changes were detected.
 - *Root cause:* Eligibility for Parallel Surface Scan Count options was incorrectly compared.
 - *Fix:* Corrected condition to validate Parallel Surface Scan Count input options.
 - *Risk:* Low
- Fixed an issue where Managed SED controller password countdown timer was not getting displayed when failed unlock attempts are exceeded.
 - *Root cause:* Controller password countdown timer and remaining attempt information are shown only when attempts are remaining.
 - *Fix:* Show Controller password countdown timer and remaining attempt information in HII even when no attempts are available.
 - *Risk:* Low
- Fixed the incorrect Chinese translation for SED OPAL in HII.
 - *Root cause:* No translation required for technical term Opal.
 - *Fix:* Changed translations to keep the original technical term Opal.
 - *Risk:* Low

2.1.3 Driver Fixes

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

2.1.3.1 Linux Driver Fixes

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

2.1.3.1.1 Fixes and Enhancements for Linux Driver Build 2.1.28-025

This release includes the following fixes and enhancements.

- Fixed an issue to handle multi-path failover.
 - *Root cause:* Controller firmware does not return the proper error code for I/O errors caused by a multi-path path failure.
 - *Fix:* The driver maps I/O errors returned by the controller firmware into errors that cause the multi-path layers in the OS to detect the failure of a path.
 - *Risk:* Low
- Fixed an issue to correct RAID bypass counter. An OS crash issue occurs while updating the RAID bypass counter.
 - *Root cause:* The SmartPQI driver was using the RAID bypass counter pointer that was not allocated. This results in a NULL pointer de-reference issue which causes the OS to crash.
 - *Fix:* Driver now updates the RAID bypass counter pointer in the device structure when the driver detects that bypass has been enabled.
 - *Risk:* Low

2.1.3.2 Windows Driver Fixes

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

2.1.3.2.1 Fixes and Enhancements for Windows Driver Build 1010.96.0.1007

There are no known fixes for this release.

2.1.3.3 FreeBSD Driver Fixes

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

2.1.3.3.1 Fixes and Enhancements for FreeBSD Driver Build 4500.0.1024

This release includes the following fixes and enhancements:

- Added the ability to set the driver debugging levels in the `loader.conf` file.
- Fixed an issue where with INVARIANTS enabled kernel, panic observed while creating and deleting array.
 - *Root cause:* While creating and deleting array, freeing of memory was happening inside spinlock.
 - *Fix:* Move the memory freeing outside the spinlock.
 - *Risk:* Low
- Fixed an issue where a panic is observed while hot-removing a drive with an INVARIANTS-enabled kernel.
 - *Root cause:* `pqisrc_free_device` function frees the device memory inside the spinlock.
 - *Fix:* Move the memory freeing outside the spinlock.
 - *Risk:* Low

2.1.3.4 VMware Driver Fixes

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

2.1.3.4.1 Fixes and Enhancements for VMware Driver Build 4500.0.1024

This release includes the following fixes and enhancements:

- Fixed an issue where PSOD is observed due to double freeing of device memory.
 - *Root cause:* When a dead path is created due to the deletion of a logical drive, the corresponding device memory won't be freed until the dead path is eliminated. If another

logical drive is subsequently created, the firmware can assign the same SCSI-3 address that was associated with the previously deleted logical drive. Adding the new logical drive with the same SCSI-3 address may lead to a duplicate entry in the device list. This duplication can trigger PSOD due to double freeing of memory.

- *Fix:* If a new logical drive is assigned the same SCSI-3 address that was previously associated with a deleted logical drive, refrain from adding the new logical drive to the device list. To facilitate the recovery of the new logical drive, customers are required to first clear the dead path. Following the clearing of the dead path, it is necessary to initiate a driver rescan to identify the new logical drive. If clearing the dead path fails, a host reboot is required.
- *Risk:* This fix may cause failures when attempting to add new logical drives, particularly in cases involving a dead path.
- Fixed an issue where PSOD observed during array delete operation.
 - *Root cause:* The serial number received for an existing logical drive in the driver device list was reported as zero during a specific rescan. With the same SCSI address, the same WWID, but different serial number (all zeros), resulted in adding a new entry in device list with the same Bus:Target:Lun values creating a duplicate entry for the same device.
 - *Fix:* Avoid serial number checks for logical drives, as the serial number received for all logical drives is the same (assigned with the corresponding controller serial number).
 - *Risk:* Medium
- Fixed an issue where the firmware version shows only revision.
 - *Root cause:* The firmware version information currently used by driver comes from a 4-byte field returned as part of the 'ID controller' inquiry reply. This 4 byte field worked for previous products, but newer controllers have a more verbose versioning scheme and 4 bytes won't work. The 4-byte field is deprecated, and a 32-byte field is now available, containing the full ASCII character string for the firmware version.
 - *Fix:* Populate the VMware SAS Adapter structure's firmware field using content of the longer 32-byte firmware version field returned by 'ID controller' inquiry.
 - *Risk:* Low
- Fixed an issue where the failed Logical Volume takes too long to remove from OS level.
 - *Root cause:* When a Logical Volume fails, the SmartPQI driver does not detect the failure. Due to this, the upper storage layer might take some time to identify the failure which results in a delay in removing the failed volume.
 - *Fix:* Convert the information received for logical volume states to right Endian format before processing volume offline status.
 - *Risk:* Low
- Fixed an issue where the NVMe drive fails to be discovered after hot-plug.
 - *Root cause:* During hot-plug of the drive, the SmartPQI driver triggers driver rescans based on the events received from the firmware. In one of the rescans, the driver was not interpreting the information used to determine that the device is part of the logical drive.
 - *Fix:* The driver is correctly interpreting the that the physical device is part of a logical drive.
 - *Risk:* Low

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 26540

This release includes the following fixes and enhancements:

- Added support in maxView and ARCCONF to display the backplane firmware version in both Decimal and Hexadecimal format.
- Added GETCONFIGJSON arconf command to get the configuration in JSON format.
- Deprecated the MNPdelay property and settings from maxView and ARCCONF. Deprecated the arconf SETPERFORM command and moved the 'Degraded Performance Optimization (DPO)' property under SETCONTROLLERPARAM command.
- Fixed an issue where the manufacturing part number property was not displayed in the "arconf GETCONFIG" command output.
 - *Root cause:* Incorrect check was present to get manufacturing part number from manufacturing data.
 - *Fix:* Implemented changes to get the part number from the manufacturing data structure.
 - *Risk:* Low
- Fixed an issue where the alert message was not displayed in maxView for the cache backup failure.
 - *Root cause:* The alert message was not added in maxView for the cache backup failure.
 - *Fix:* Implemented changes to add warning device alert message in maxView for the cache backup failure.
 - *Risk:* Low
- Fixed an issue where the product codename "Flashlight" was incorrectly exposed in "arconf GETCONFIG" command output.
 - *Root cause:* Legacy product code name "Flashlight" was used while displaying cache state in arconf getconfig command.
 - *Fix:* Implemented changes to provide the appropriate message in the cache state in place of disclosing the product codename "flashlight" in the display.
- *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 6.35.8.0

This release includes the following fixes and enhancements:

- Added support for RDE ACTION operation on `#Storage.SetEncryptionKey`. This feature allows Redfish clients to enable SED encryption with local key management (LKM) and change the master key/master key identifier when MSED LKM is enabled. The GetPDR response for the Storage resource's RedfishAction PDR will now include the ActionName "Actions/Storage.SetEncryptionKey". The response for a RDE READ operation on a Storage resource will include `#Storage.SetEncryptionKey` as part of the Storage.Actions property only when EncryptionMode is UseLocalKey or disabled. To enable local SED encryption, the RDE ACTION operation payload should be as follows:

```
{
  "EncryptionKey": "The local encryption key to set on the storage subsystem.",
  "EncryptionKeyIdentifier": "The local encryption key identifier used by the storage subsystem."
}
```

To change the master key and/or master key identifier when MSED LKM is enabled, the RDE operation payload should be as follows:

```
{
  "EncryptionKey": "The local encryption key to set on the storage subsystem.",
  "EncryptionKeyIdentifier": "The local encryption key identifier used by the storage subsystem."
}
```

Note: In both of these cases, EncryptionKeyIdentifier is an optional parameter. If LKM MSED is being enabled initially and the EncryptionKeyIdentifier is not provided in the payload, then it will be given a default value. If LKM MSED is already enabled and the EncryptionKeyIdentifier is not provided in the payload, then it will be left at its current value.

- Added support to perform RDE UPDATE on Storage.EncryptionMode. This feature allows Redfish clients to enable remote mode SED encryption or to disable SED encryption.

Note: This feature does NOT allow Redfish clients to modify/enable controller-based encryption.

RDE operations targeting the Storage resource will have the update access bit (bit 1) of PermissionFlags set. The following changes have been made to RDE READ on a Storage resource:

- Updated the @Redfish.WriteableProperties property to indicate if "EncryptionMode" can be updated.
- Added EncryptionMode@Redfish.AllowableValues if "EncryptionMode" can be updated. The allowable values are:
 - If current SED encryption is Disabled, "UseExternalKey" to enable remote SED encryption will be published.
 - If current SED encryption is set to local mode, "Disabled" to disable SED encryption will be published.
 - If current SED encryption is set to remote mode, "Disabled" to disable SED encryption will be published.

The following steps can be followed to manage remote mode SED encryption via RDE UPDATE operations on the storage resource:

- To disable SED encryption, the following payload should be used with a RDE UPDATE operation on the Storage resource:

```
{
  "EncryptionMode": "Disabled"
}
```

Disabling SED encryption requires a long-running task. So, MCs are required to support long-running tasks.

- To enable remote mode SED encryption, the following payload should be used with a RDE UPDATE operation on the Storage resource:

```
{
  "EncryptionMode": "UseExternalKey"
}
```

- Added support for creating SED-encrypted volumes using a CREATE operation on the VolumeCollection resource. Updated the VolumeCapabilities resource RDE READ response to include the following properties if SED encryption is supported on the controller:
 - "Encrypted@Redfish.OptionalOnCreate": true
 - "EncryptionTypes@Redfish.OptionalOnCreate": true

- "EncryptionTypes@Redfish.AllowableValues": ["NativeDriveEncryption"]

Updated Volume creation as follows:

- SED encryption is enabled at the array level. Once the array has SED encryption enabled, all Volumes part of the same array will have SED encryption enabled.
- Users may not disable SED encryption on a Volume once enabled or enable SED encryption on existing Volumes that are not encrypted.
- Creating the first Volume on an array of SED drives will result in a long-running task being created to handle the RDE CREATE operation.
- Added support to perform RDE UPDATE on Volume.Encrypted. This feature allows Redfish clients to take ownership of an SED or revert a SED to OFS.

Note: This feature is only applicable to HBA volumes representing SEDs. An HBA volume representing a SED will have its "EncryptionTypes" set to "NativeDriveEncryption."

The following changes were made:

- RDE operations targeting a HBA Volume resource will have the update access bit (bit 1) of PermissionFlags set.
- Changes to RDE READ on a HBA Volume resource:
 - Added @Redfish.WritableProperties property to indicate if "Encrypted" can be updated.
- RDE UPDATE on HBA Volume resource:
 - To take ownership of a SED, the RDE UPDATE operation payload should be as follows:

```
{
  "Encrypted": true
}
```

- To revert a SED to OFS, the RDE UPDATE operation payload should be as follows:

```
{
  "Encrypted": false
}
```

- Added the following properties to the RDE READ responses for the associated resources:
 - Drive.Status.Conditions
 - Port.Status.Conditions
 - StorageController.Status.Conditions
 - StorageController.CacheSummary.Status.Conditions
 - Volume.Status.Conditions

These properties are of array type containing Condition objects whose MessageId child properties are taken from the DMTF Redfish StorageDevice v1.1.0 message registry. For a given resource, the Status.Conditions array will be empty, if the associated Status.Health value is Ok.

- Added support for PLDM Type 5 UBM PIC® firmware Flashing behind expander. PLDM Type 5 now updates the firmware of UBMs behind an SEP device when the update is being performed on the SEP.
- PLDM Type 5 VerifyComplete command will no longer return a VerifyResult of SECURITY_REVISION_ERROR (0x10). SECURITY_REVISION_ERROR was previously being returned on certain controllers after validating a firmware image that would result in a security revision update when write caching is enabled on the controller.
- Updated the PLDM base (Type 0) command GetPLDMTypes to indicate support for PLDM for File Transfer (PLDM Type 7).

- Updated the PLDM base (Type 0) command `GetPLDMVersion` to indicate new support for the following specification versions:
 - DSP0240 - PLDM Base Spec (Type 0) -> v1.1.0 and v1.2.0
 - DSP0248 - PLDM for Platform Monitoring and Control (Type 2) -> v1.3.0
 - DSP0242 - PLDM for File Transfer (Type 7) -> v1.0.0
 - Updated the PLDM base (Type 0) command `GetPLDMCommands` to indicate support for the following commands:
 - Type 0 command 0x07 - `NegotiateTransferParameters`
 - Type 0 command 0x09 - `MultipartReceive`
 - Type 7 command 0x01 - `DfOpen`
 - Type 7 command 0x02 - `DfClose`
 - Type 7 command 0x06 - `DfHeartbeat`
 - Added support for the PLDM Type 0 command `NegotiateTransferParameters` for PLDM Type 7. Implemented the PLDM Type 0 command `NegotiateTransferParameters` to allow an MC and device to negotiate the transfer part size to be used for multipart transfers using the PLDM Type 0 commands `MultipartSend` and `MultipartReceive`. This command will accept a transfer part size that is a power of two that is at least 256 bytes and will support negotiation for PLDM Type 7 (PLDM for file transfer). A successful response will report a responder part size of 512 bytes for PLDM Type 7.
 - Added support for PLDM Type 0 `MultipartReceive` command. Only PLDMType of 7 is supported for the `MultipartReceive` command. A successful `XFER_ABORT` will have the following response:
 - `CompletionCode` set to `SUCCESS`
 - `TransferFlag` set to `ZERO`
 - `NextDataTransferHandle` set to `ZERO`
 - `DataLengthBytes` set to `ZERO`
 - Updated the `GetPDR` response for the controller's `EntityAssociation` PDR to include a contained entity representing the crash dump `Device File`. This contained entity will be published at the tail end of the contained entities array and will be addressed with the following:
 - `EntityType` = `Device File (0x09, defined in DSP0249 v1.2.0)`
 - `EntityInstanceNumber` = 1
 - `ContainerId` = `0x9005`
 - Added support for reporting a PDR of type `FILE_DESCRIPTOR` for controller's crash dump log. The new PDR will be reported as part of PLDM Type 2 `GetPDR` command. The `FileClassification` field of the `GetPDR` response will be set to `CrashDump`.
 - Added a new `State Sensor` class PDR that will define state sensor readings for the crash dump device file. This PDR will report the following field values:
 - `StateSetId` - 68 (`Device File, defined in DSP0249 v1.2.0`)
 - `PossibleStates` - Bit 1 (Updated), Bit 4 (Max Size), and Bit 5 (Unchanged)
 - `PossibleStatesSize` - 5
- Updated `GetStateSensorReadings` to handle requests for the `sensorId` provided in the crash dump device file's `State Sensor` PDR. The reported state will be determined based on the current size of the crash dump relative to the last reading request.
- Added support for reporting a `Crash Dump Numeric Sensor` PDR of type `NUMERIC_SENSOR` to monitor the file size for a crash dump device file. The new PDR will be reported as part of PLDM Type 2 `GetPDR` command. Readings of numeric sensors associated with a given

FileDescriptor PDR will be supported using the existing Type 2 commands `GetSensorReading` and `GetSensorThresholds`.

- Added support to implement the PLDM Type 7 command `DfOpen` to allow opening a file for reading. This command will support opening a file with exclusive, regular (non-FIFO) read access. On success, a file descriptor will be provided that can be used as a handle for accessing and managing the opened file. Implemented the PLDM Type 7 command `DfClose` to close the file associated with a given file descriptor.

Implemented the PLDM Type 7 command `DfHeartbeat` that allows the MC and device to negotiate a timeout interval that will allow the device to close the file if no read or heartbeat refresh activity has occurred within that negotiated interval. On a successful negotiation, a response containing a `ResponderInterval` of 30000 milliseconds will be returned, and the smaller of that value and the requested `RequesterInterval` will govern as the `NegotiatedInterval` for the open file's heartbeat timer.

To successfully read the crash dump file, a file client MC should execute the following sequence:

- Send the `NegotiateTransferParameters` command to indicate a valid `RequesterPartSize` for PLDM Type 7.
- Send the `GetPDR` command to read the controller's `CrashDump` File Descriptor PDR and obtain the associated `FileIdentifier`.
- Send the `GetPDR` command to read the numeric sensor PDR associated with the crash dump file entity, and check that it has a non-zero size using `GetSensorReadings` with the appropriate `sensorId`.
- Send the `DfOpen` command with the crash dump file's `FileIdentifier` and requesting exclusive read access, obtaining a valid `FileDescriptor`.
- If desired, optionally send the `DfHeartbeat` command to negotiate a heartbeat interval for the open `FileDescriptor`.
- Send a sequence of `MultipartReceive` commands with `PLDMType = 7`, `TransferContext` equal to the open `FileDescriptor`, and the appropriate `TransferOperation` based on the progress through the file read as detailed in DSP0240.
- If a heartbeat interval has been negotiated and the file client is unable to handle another `MultipartReceive` request and response within that interval, send the `DfHeartbeat` command again to keep the open `FileDescriptor` alive.
- After concluding the full sequence of `MultipartReceive` requests and responses, send the `DfClose` command to close and free up the open `FileDescriptor`.
- Fixed an issue where an MC sends a `NegotiateRedfishParameters` request with the "BEJ v1.1 encoding and decoding supported" bit of `MCFeatureSupport` unset, that is, with the intent to negotiate for only BEJ v1.0 support, all RDE response bodies from the controller will be encoded using BEJ v1.1.
 - *Root cause:* The BEJ version v1.1 was hard-coded into all RDE READ responses and extended info messages as well as all Redfish alert messages.
 - *Fix:* Modified the response and message encoding to use the lowest common BEJ version supported by both the MC and the controller.
 - *Risk:* Low
- Fixed an issue where the `Drive.Manufacturer` property was not published during an RDE READ on SAS drives.
 - *Root cause:* The logic to publish `Drive.Manufacturer` property was being set to `False` on some controllers.
 - *Fix:* Corrected the logic to publish `Drive.Manufacturer` for SAS drives on all supported controllers.

- *Risk:* Low
- Fixed an issue where the `BlockSizeBytes` and `LogicalUnitNumber` are missing from the RDE READ response for all volume resources with `RAIDType = None`, that is, HBA Volumes.
 - *Root cause:* BEJ encoding for these two properties was missing from the function handling RDE READ operations for HBA volumes.
 - *Fix:* Modified the HBA Volume RDE READ operation handler function to publish `BlockSizeBytes` and `LogicalUnitNumber` for HBA volumes.
 - *Risk:* Low
- Fixed an issue where incorrect severity was observed in Redfish Eventing for cache status.
 - *Root cause:* The `CacheStatus` was hard-coded to only be sent with a severity of `Warning` when the cache state is `TemporarilyDegraded` and the power source is charging.
 - *Fix:* Modified the cache status event logic to send a `Severity` of `OK` for certain controllers.
 - *Risk:* Low
- Fixed an issue where incorrect state and health is observed when controller password is required. When the controller is waiting for boot password for CBE or SED encryption, the RDE READ on Storage Resource operation returns `Ok` for health and `Enabled` for state.
 - *Root cause:* The function which encodes the RDE READ response for the the `StorageController` resource didn't check the situation when the controller is waiting for boot password.
 - *Fix:* Modified the RDE READ response encoder function for the `StorageController` resource to check whether the controller is waiting for boot password on CBE and SED encryption. If so, set controller health as `Warning` and state as `StandbyOffline`.
 - *Risk:* Low
- Fixed an issue where the response data of RDE command `0x12` (`RetrieveCustomResponse`) is incorrect. In the response to a `RetrieveCustomResponseParameters` command, the size of the `NewResourceID` field is 2 bytes instead of the spec-defined size of 4 bytes.
 - *Root cause:* The internal structure used to build the command response assigned a type of `uint16` instead of `uint32` to the `NewResourceID` field.
 - *Fix:* Updated the command response structure to assign the `uint32` type to the `NewResourceID` field.
 - *Risk:* Low
- Fixed an issue to accommodate DMTF definition of `bejBoolean` for any non-`0x00` value as `TRUE`. BEJ boolean values other than numeric one are being set to `False`. However the DMTF specification indicates that a numeric value of zero indicates boolean `False` while any other value indicates `True`.
 - *Root cause:* The decoded boolean numeric value is being compared to one. If the numeric value is one, then the boolean value is being set to `True`. Otherwise, the boolean value is being set to `False`.
 - *Fix:* When the numeric value for the boolean has been decoded, it will be compared to zero. If the numeric value is zero, then it will stay zero. If the numeric value is not zero, then it will be set to one.
 - *Risk:* Low
- Fixed an issue where incorrect severity was observed in Redfish Eventing for Foreign SED Drive. When a Foreign owned Drive is inserted, `DriveOffline` alerts will have a severity of `OK` instead of `Warning`.
 - *Root cause:* There is no existing `DriveOffline` status with a severity of `Warning`. Default `Severity` is `OK` for this specific `DriveOffline` condition.

- *Fix*: Modified the drive offline status event logic at time of drive insertion to have a severity of Warning.
- *Risk*: Low
- Fixed an issue with improper capacity information in Redfish drive Name. The capacity in the drive name property does not include a two-digit remainder, when the whole part of the capacity is over one digit.
 - *Root cause*: When calculating capacity, the calculation of the remainder would be blocked if the whole part was larger than a single digit.
 - *Fix*: Capacity remainder is now calculated as long as the whole part of the capacity is fewer than four digits.
 - *Risk*: Low

2.2 Limitations

This section shows the limitations for this release.

2.2.1 General Limitations

This release includes the following general limitations.

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

2.2.2 Firmware Limitations

This section shows the firmware limitations for this release.

2.2.2.1 Limitations for Firmware Release 03.01.28.82

This release includes the following limitations:

- Persistent Event Logs (PEL) will be cleared under the following conditions:
 - Upgrading from firmware releases prior to 03.01.17.56 to 03.01.17.56 or later firmware releases.
 - Downgrading from firmware releases 03.01.17.56 or later to firmware releases prior to 03.01.17.56.
- Firmware downgrade is blocked if disk-based transformation is in-progress.
 - *Workaround*: Wait for the transformation to complete and retry the firmware downgrade.
- Transformation is blocked if a reboot is done after the firmware update is pending, and the flashed new firmware version is older than 03.01.17.56.
 - *Workaround*: Reboot the system.
- Logical drive is not detected when disk-based transformation is in-progress during logical drive movement to a different controller and the different controller has a firmware version older than 03.01.17.56, or, the firmware downgrade occurred while internal-cache based transformation was in progress, but the Backup Power Source failed before firmware activation.
 - *Workaround*: Move the logical drive to a controller with firmware version 03.01.17.56 or later.
- Firmware downgrade from firmware version 3.01.23.72 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.

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 2.12.1/Legacy BIOS Build 2.12.3

There are no known limitations for this release.

2.2.4 Driver Limitations

This section shows the driver limitations for this release.

2.2.4.1 Linux Driver Limitations

This section shows the Linux driver limitations for this release.

2.2.4.1.1 Limitations for Linux Driver Build 2.1.28-025

This release includes the following limitations:

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

Note: This does not affect Oracle 8 UEK 7.

 - *Workaround:* Install the rpm using "--nodeps" when dependency failures occur.
 - Update:
 - For SmartPQI driver versions > 2.1.20-020 and UEK7 kernels >= 5.15.0-3.60.2.el9uek.x86_64, the SmartPQI rpm will install normally.
 - For UEK7 kernels < 5.15.0-3.60.2.el9uek.x86_64, the SmartPQI rpm needs to be installed using the "--nodeps".
- On AMD systems, the system might crash or hang due to a bug in the IOMMU module. For details, see lore.kernel.org/linux-iommu/20191018093830.GA26328@suse.de/.
 - *Workaround:* Disable the IOMMU setting option in BIOS.
- Depending on hardware configurations, the SmartPQI `expose_ld_first` parameter may not always work consistently.

- *Workaround:* None
- On some distributions (including RHEL 9.0/Oracle Linux 9.0), you are unable to inject the OOB driver (DUD) during install when a multi-actuator drive is attached.
 - *Workaround:* Install using the inbox driver, complete OS installation, then install the OOB driver.

2.2.4.2 Windows Driver Limitations

This section shows the Windows driver limitations for this release.

2.2.4.2.1 Limitations for Windows Driver Build 1010.96.0.1007

This release includes the following limitations:

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

2.2.4.3 FreeBSD Driver Limitations

This section shows FreeBSD driver limitations for this release.

2.2.4.3.1 Limitations for FreeBSD Driver Build 4500.0.1024

This release includes the following limitations:

- If an ATA-locked drive is connected to the controller, the system may experience a hang during shutdown.
 - *Workaround:* There are two workarounds for this issue:
 - i. Perform a force reboot using the following command: `reboot -q -n -f.`
 - ii. If shutdown is hung, hot removing the ATA locked drive will allow the shutdown process to complete gracefully.
- FreeBSD 13.2 and later OS installations will fail with the out of box driver.
 - *Workaround:* Install with inbox driver then update to latest.

2.2.4.4 VMware Driver Limitations

This section shows VMware driver limitations for this release.

2.2.4.4.1 Limitations for VMware Driver Build 4662.0.112

This release includes the following limitations:

- Legacy interrupt mode is not supported in this release.
- If the controller SED Encryption feature is "On" and locked, Datastores created from secured logical drives on the controller are not automatically mounted even after unlocking the controller, they are not visible through the ESXi hypervisor client.
 - *Workaround:* Use the command `vmkfstool -V` or ESXCLI storage filesystem rescan. Alternatively, use the Rescan option from the Devices tab in the Hypervisor's Storage section. Any of these options solve the issue by forcing a rescan, causing the datastore to mount.
- A controller lockup may occur when using VMDirectPath on a single-processor AMD system. These lockups have been seen with VMs running Linux and Windows. No known workaround at the present time. If a lockup of a passed-through controller occurs, a reboot of the ESXi server may be required to clear the lockup condition and restore the virtual machine to working condition.
- Customers may encounter failures when attempting to add new Logical Drives (LD), particularly in cases involving a dead path.
 - *Workaround:* To facilitate recovery of new LD, customers are required to clear the dead path initially. Following the clearance of the dead path, if the newly created LD is still not exposed, then it is required to initiate a driver level rescan using the appropriate management tool. While clearing the dead path fails, a host reboot is required.

2.2.5 Management Software Limitations

This section shows management software limitations for this release.

2.2.5.1 maxView Storage Manager/ARCCONF Limitations

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

2.2.5.1.1 Limitations for maxView Storage Manager/ARCCONF Build 26540

There are no known limitations for this release.

2.2.5.2 PLDM Limitations

This section shows the PLDM limitations for this release.

2.2.5.2.1 Limitations for PLDM Release 6.35.8.0

There are no known limitations for this release.

3. Updating the Controller Firmware

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



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

3.1 Updating Controllers to Latest Firmware

If running firmware is 3.01.00.006 or lower, please contact Adaptec Apps team at ask.adaptec.com.

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

Microchip Information

The Microchip Website

Microchip provides online support via our website at www.microchip.com/. 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.

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

ISBN: 978-1-6683-4122-3

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>