SmartHBA 2100 and SmartRAID 3100 Software/Firmware Release Notes

© 2022 Microchip Technology Inc.
and its subsidiaries

Table of Contents

1. About This Release....3
   1.1. Release Identification....3
   1.2. Components and Documents Included in this Release....4
   1.3. Files Included in this Release....4
2. What's New?....7
   2.1. Features....7
   2.2. Fixes....8
   2.3. Limitations....16
3. Updating the Controller Firmware....19
   3.1. Updating the Controller Firmware....19
4. Installing the Drivers....20
5. Revision History....21
The Microchip Website....22
Product Change Notification Service....22
Customer Support....22
Microchip Devices Code Protection Feature....22
Legal Notice....22
Trademarks....23
Quality Management System....24
Worldwide Sales and Service....25

1. About This Release

The development release described in this document includes firmware, OS drivers, tools, and host management software for the 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	2.6.6
Package Release Date	March 9, 2022
Firmware Version	5.00 B0[1,2] (basecode 06.06.018.000)
UEFI Version	1.3.17.1
Legacy BIOS	1.3.17.2
Driver Versions[3]	Windows SmartPQI: • Windows 2016/2019/2022: 1010.24.0.1005 • Windows 10/11: 1010.24.0.1005 Linux SmartPQI: • RHEL 7/8: 2.1.16-030 • SLES 12/15: 2.1.16-030 • Ubuntu 16/18/20/21: 2.1.16-030 • Debian 10/11: 2.1.16-030 • CentOS 7/8: 2.1.16-030 • Oracle Linux 7/8: 2.1.16-030 • Citrix XenServer 8: 2.1.16-030 VMware SmartPQI: • VMware 6.5/6.7/7.0: 4252.0.103 FreeBSD/Solaris SmartPQI: • FreeBSD 11/12/13: 4210.0.1004 • Solaris 11: 11.4120.0.1005
Management Software **(arcconf, maxView™, Event Monitor, BootUSB)**	B24763

Notes:

1. Downgrading to 1.04 B0 or older builds from this release or prior 1.29 releases may cause the board to not boot or have supercap errors due to an incompatibility in SEEPROMs between this release and prior releases. Refer to the section "3. Updating the Controller Firmware" to downgrade an existing board.
2. If the firmware running on the board is older than 0.01 B594, existing data in the logical drives must be backed up if it needs to be used after the upgrade. After the upgrade from firmware prior to 0.01 B594, the logical drives will need to be recreated.
3. Only run the driver on firmware 0.01 build 500 or later.
4. Only Windows 11 Inbox driver is supported.

1.2 Components and Documents Included in this Release

Download the firmware, drivers, host management software, and supporting documentation for your SmartHBA 2100/ SmartRAID 3100 controller and SmartRAID 3100 and SmartRAID 3100 controller solutions from the Microchip Web site at https://start.adaptec.com

1.3 Files Included in this Release

This release consists of the files listed in the following tables:

Firmware Files

Table 1-2. Firmware Files

Component	Description	Pre-Assembly Use	Post-Assembly Use
SmartFWx100.bin	Programmable NOR Flash File Use to program NOR Flash for boards that are already running firmware.	—	X
SmartFWx100.fup	Programmable NOR Flash File Used for PLDM type 5 firmware flashing for boards that are already running firmware.	—	X

Table 1-3. Firmware Programming Tools

Tool	Description	Executable
Arcconf romupdate	The command allows to upgrade/downgrade the firmware and BIOS image to the controller.	Refer to Table 1-7
maxView™ firmware upgrade wizard	The firmware upgrade wizard allows to upgrade/ downgrade the firmware and BIOS image to one or more controller(s) of same model in the system.	Refer to Table 1-7

Driver Files

Table 1-4. Windows Storport Miniport SmartPQI Drivers

Package	Drivers	Binary	Version
2012	Server 2022, Server 2019 and Server 2016 Windows 10 (version 21H2) and 11	SmartPqi.sys	x64
		SmartPqi.inf	x64
		Smartpqi.cat	x64

Table 1-5. Linux SmartPQI Drivers

Drivers	Version
Red Hat Enterprise Linux 8.5, 8.4, 8.3, 8.2, 8.1, 7.9, 7.8, 7.7, 7.6	x64
CentOS 8.4, 8.3, 8.2, 8.1, 8.0, 7.9, 7.8, 7.7	x64
SuSE Linux Enterprise Server 12[1], SP5, SP4, SP3	x64
SuSE Linux Enterprise Server 15 SP3, SP2, SP1[1]	x64
Oracle Linux 7.9, 7.8, UEK6U3	x64
Oracle Linux 8.4, 8.3 UEK6U3	x64
Ubuntu 21.04	x64
Ubuntu 20.04.3, 20.04.2, 20.04	x64
Ubuntu 18.04.5, 18.04.4, 18.04	x64
Ubuntu 16.04.5	x64
Debian 11.1, 10.10, 10.05	x64
Citrix xenServer 8.2, 8.1, 8.0	x64
Fedora 35 (inbox only)	x64

Note: 1. To mitigate against the Spectre Variant 2 vulnerability, the RHEL 6u9/RHEL7u4/RHEL7u5 and SLES11 SP3 and higher drivers have been compiled to avoid the usage of indirect jumps. This method is known as "Retpoline".

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

Drivers	Version
FreeBSD 13, 12.3, 11.4	x64
Solaris 11.4	x64
VMware 6.7 U3/U2, 6.5 U3/U2	x64
VMware 7.0 U3/U2	x64

Host Management Software

Table 1-7. Host Management Utilities

Description	OS	Executable
ARCCONF Command Line Utility	Windows x64 Linux x64 VMware 6.5 and above XenServer FreeBSD x64 Solaris x86	See the Arcconf download package for the OS-applicable installation executable.
ARCCONF for UEFI		Included as part of the firmware downloadable image.

.....continued

Description	OS	Executable
maxView™ Storage Manager	Windows x64 Linux x64 VMware EXSi 6.5 and above XenServer	See the maxView Storage Manager download package for the OS-applicable installation executable.
maxView™ vSphere Plugin	VMware 6.5 and above	See the VMware maxView Storage Manager download package for the OS-applicable installation executable.
Boot USB (offline or pre-boot) for ARCCONF and maxView Storage Manager	Linux x64	See the maxView BootUSB download package for the .iso file.

2. What's New?

This section shows what's new in this release.

Important: Updated maxView to address log4j vulnerabilities.

2.1 Features

The following table lists features supported for this release.

Table 2-1. Feature Summary

Feature		Supported in this Release	Future Release
UEFI Driver, Boot Support		X	
Legacy Boot Support		X	
Dynamic Power Management		X	
SMR Drive Support	Enumeration, Unrestricted Command Flow-Through	X	
	SATL Translation for HA/HM SMR Management	X	
	Identify All Drive Types	X	
Driver Support	Windows	X	
	Linux	X	
	VMware	X	
	FreeBSD	X	
	Solaris	X	
	OS certification	X	
Out of Band interface selection support of MCTP or PBSI		X	
Flash Support		X	
MCTP BMC Management		X	
Configurable Big Block Cache Bypass		X	
Green Backup Support for SmartRAID		X	
4Kn Support in RAID		X	

© 2022 Microchip Technology Inc.
and its subsidiaries

2.2 Fixes

2.2.1 Firmware Fixes

2.2.1.1 Fixes and Enhancements for Firmware Release 5.00 B0

This release includes the following fixes and enhancements:

- Added support for Firmware Update Package (FUP) file to support PLDM type 5 firmware flashing.
- Added support in firmware to enable logical drive transformation without backup power source on the controllers supporting data preservation. Firmware will use the disk-based transformation strategy similar to SmartIOC controllers with a limitation of transformation backup memory of 8192 sectors. Once the backup power source reaches the optimal state, transformation will use DDR as backup memory as usual. An event is also logged to show if the DDR cache or drive is being used as transformation backup memory.
- Fixed a problem where the incorrect backup power source status was reported by the host management software.
  - Root Cause: During a Backup-in-progress failure, which occurs when the supercap is removed while DDR cache is being backed-up, the backup power status is reported as charging instead of failed. Firmware is not updating the “not configured state” status correctly.
  - Fix: Update the backup power source status to the Failed state.
  - Risk: Low
- Fixed an issue where SCSI pass through command fails due to incorrect data direction.
  - Root Cause: Firmware uses the opcode sent in the CDBs to determine the SCSI data direction. If an obsolete command is sent, firmware assumes IN direction without honoring the host provided direction.
  - Fix: Use the host provided direction information if the direction cannot be determined using the opcode within firmware.
  - Risk: Low
- Fixed an issue where a transforming logical drive may be reported in the Failed state for cache-less controllers after an unsafe reboot.
  - Root Cause: During transformation on cache-less controllers, the transformation backup data and a checksum is saved on one of the physical drives in the logical drive. After an unsafe reboot of the system, the firmware used an incorrect physical drive index to check for existing transformation backup data and encountered a checksum mismatch because it read a different physical drive than where the transformation backup data was previously written. Due to the checksum mismatch the transformation backup data is not available and the logical drive is failed.
  - Fix: The firmware will use the correct physical drive index to check for existing transformation backup data and if needed, firmware will check all the physical drives in the logical drive to find existing transformation backup data.
  - Risk: Low
- Fixed a 0x1E30 lockup during second logical drive creation without backup power source and with existing maxCache pair.
  - Root Cause: When a second logical is being created without a backup power source present, the DDR cache configuration is changed from 100% write to 100% read (read-only). Due to a pre-existing maxCache logical drive, the DDR cache configuration should remain as 100% write. Since the DDR cache configuration changed to 100% read (read-only), READ AHEAD request gets set in cache module and when maxCache processes the request, the lockup occurs as READ AHEAD requests are not supported by maxCache.
  - Fix: During logical drive creation, firmware will not change the DDR cache configuration to 100% read (read-only) when the backup power is missing and maxCache is present.
  - Risk: Low
- Fixed an issue where a transforming logical drive failed when system rebooted with transforming logical drive and newly created logical drive.

  - Root Cause: When the backup power source is removed during transformation, a transformation shutdown happens. Transformation shutdown will suspend the transformation and save the progress information in the transformation progress area in the drive's RAID meta data. This progress info has the current configuration signature also. When a new logical drive is created the configuration signature is changed and it is now different than what was stored in the transformation progress information. During system reboot firmware sees that the configuration signature present in the drive RAID metadata and transformation progress area are different, so failed the transforming logical drive.
  - Fix: While creating a logical drive, always update the transformation progress information with updated global signature if the transformation was running and was shutdown.
  - Risk: Low

- Fixed a 0x17005 lockup while resuming transformation after a cold boot when a logical drive under-going a transformation is moved from a cache-less controller to a cache-supported controller.
  - Root Cause: While the system is booting up, transformation is trying to restore the transformation backup data by doing a DMA from DDR cache for calculating and validating transforamtion progress data checksum. This DMA executed an uninitialized callback function which causes the lockup.
  - Fix: Correctly initialize the callback function used to restore the transformation data.
  - Risk: Low

- Fixed a TLB exception issue observed during surface scan retry failure.
  - Root Cause: Firmware accessed a NULL address because the logical drive data structure was not initialized correctly in the surface scan retry code path.
  - Fix: Initialize the logical drive data structure before calling surface scan for a retry.
  - Risk: Low

- Fixed a problem where I/O latency may be higher for logical drive I/O with a predictive failure drive.
  - Root Cause: Firmware may decide to optimize a read to a Predictive Failure drive when it should not optimize the read and may decide not to optimize a read to a Predictive Failure drive when it should optimize the read. If the request was a retry, this read optimization should not be done, but the conditions which determines the retry type was incorrect.
  - Fix: Correctly determine the retry type.
  - Risk: Low

- Fixed an issue that disk name is not showing for drives after deletion of logical drive and RAID metadata present on the SEDs.
  - Root Cause: Firmware is unable to read the RAID metadata on boot if SED is locked.
  - Fix: Firmware will no longer look for locked to unlocked state transition. It will check if current status is unlocked and it will then read the RAID metadata and update drive flags allowing RAID metadata to be cleared.
  - Risk: Low

- Fixed an issue where the devices are not discovered when an active optical cable is connected to the controller.
  - Root Cause: Firmware did not use the correct sequence when writing to a CPLD register to support active optical cables.
  - Fix: Firmware uses the correct sequence to write to CPLD registers.
  - Risk: Low

- Fixed a lockup issue that can randomly occur with the SED qualification process.
  - Root Cause: When firmware checks to see if the SED qualification process has completed, a race condition may occur that causes memory to be freed twice resulting in a controller lockup.
  - Fix: The race condition is fixed so the memory is freed once.
  - Risk: Low

- Fixed an issue that drive LEDs of SAS enclosure, attached behind SAS switch, are being turned off after Non-Disruptive Software Reset (NDSR) of a SAS switch.

  - Root Cause: The only indication of a SAS switch NDSR is a broadcast change event due to link down. Firmware does not update any drive LED in this scenario. The HBA inside the SAS switch clears all drive LEDs of any SAS enclosure attached behind the switch.
  - Fix: After broadcast change due to link down, enable drive LED update to all enclosures, in the background, once per minute for three times.
  - Risk: Low

- Fixed an issue displaying wrong message of 'cold boot required' when no CPLD update occurred.
  - Root Cause: The logic in the firmware that checks if cold boot is required was relying on the CPLD major revision. The definition of major revision has been changed, which caused the issue.
  - Fix: Check a different CPLD data bit to decide if a cold boot is required.
  - Risk: Low

- Fixed the “power on hours” value reports zero for locked SED SSD under physical drive details in the host management software.
  - Root Cause: For a locked SED SSD, the drive present flag is not set so the "power on hours" value is reported as zero.
  - Fix: Check for either drive present flag or if the drive is a locked SED and return the "power on hours" correctly.
  - Risk: Low

- Fixed an issue where the “power on hours” reports 0 for Micron SATA SED SSD at run-time.
  - Root Cause: Firmware was incorrectly reading the vendor specific page that is not supported by the drive.
  - Fix: Firmware will read the vendor specific page, if it is supported by the drive.
  - Risk: Low

- Fixed an issue that SED is not recognized by SEDUTIL as a valid drive.
  - Root Cause: Using SEDUTIL, resulted in the following:
    - IDENTIFY_DEVICE ATA passthru commands were sent with flags set to data out so no data was returned.
    - SCSI_PROTOCOL_IN was incorrectly set to unknown direction with flags set to data out so no data was returned.

    Therefore, devices were not displayed for unlocking.
  - Fix: For IDENTIFY_DEVICE ATA passthru commands, firmware can determine the direction by looking at CDB[2]. Use READ bit instead of looking at certain internal interface flags. For SCSI_PROTOCOL_IN, this is always a data in command so no need to look at internal interface flags.
  - Risk: Low

- Fixed an issue that hot removal event occurs on range locked SED due to SED Qualification on logical drive during boot.
  - Root Cause: While saving RAID metadata on range locked SED a write failure may occur. The SED device is reset and some drives do not complete the reset within 10 seconds so firmware issues a device removal event.
  - Fix: Set a flag to indicate the logical drive is going through SED Qualification so that the drive will not be hot removed.
  - Risk: Low

- Fixed an issue that range locked SED failed with reason error erasing RAID metadata during reboot and remains FAILED after clear configuration.
  - Root Cause: SED_QUAL_PENDING flag was not being set or cleared for SED Qualification replacement.
  - Fix: Correctly set and clear SED_QUAL_PENDING flags for SED Qualification replacement.
  - Risk: Low

- Fixed possible controller lockup when an encrypted logical drive with transformation queued and the system is rebooted

  - Root Cause: When firmware tries to see if transformation can be started it checks the outstanding IOBypass request count to the logical drive. This is a blocking call that ends up never returning, so the controller hangs and locks up.
  - Fix: Remove the check for outstanding IOBypass requests to the logical drive since this I/O is disabled for that logical drive while a transformation is pending.
  - Risk: Low

- Fixed a possible lockup when a drive fails due to handle swaps
  - Root Cause: During a LUN reset, the drive was failed because maximum number of handle swaps was reached. The task management service did not recognize this failure and kept going as if the drive was responsive.
  - Fix: Complete TMF commands when drive is failed to force all TMFs to complete to host.
  - Risk: Low

- Loss of redundant path seen on logical drive failure with single domain configuration.
  - Root Cause: In a single domain topolgy, if all physical drives are removed causing the logical drive to be failed, a loss of redundant paths message is incorrectly reported by host management software. The message occurs because firwmare incorrectly defaults to reporting redundant paths when there are no physical drives present for the failed logical drive.
  - Fix: Report no redundant paths are present for a logical drive when there are no physical drives present.
  - Risk: Low

- Fixed an issue where IOBypass is not used after LUN Reset.
  - Root Cause: A LUN reset results in disabling IOBypass in order to quiesce outstanding I/O. After reset completes, IOBypass was enabled, but the driver was not notified.
  - Fix: Notify driver when IOBypass is re-enabled.
  - Risk: Low

- Improved reporting of temperature sensor descriptors for controller virtual SES target.
  - Root Cause: Firmware was reporting “vendor defined” for some of the sensors.
  - Fix: Changed to more meaningful “top” and “bottom” for descriptors.
  - Risk: Low

- Fixed a controller lockup when removing a SMC SC846-P enclosure.
  - Root Cause: SES control and status pages do not match in size. Controller allocated memory based on the status page size, but the control page was smaller in size resulting in firmware clearing select bits beyond the end of the control page that caused a lockup.
  - Fix: Authenticate size of SES control and status page. If they are not the same, do not support the SES functionality of enclosure.
  - Risk: Low

- Fixed a potential controller lockup 0x1BC0 during write then read I/O with URE's.
  - Root Cause: A race condition could occur between two firmware threads acting on the same clean, valid cache line simultaneously. The two threads take opposing actions which cause a lockup in one thread’s checks.
  - Fix: Added a critical section within firmware so two threads can’t access the cache line simultaneously.
  - Risk: Low

- Fixed an issue with logical drives created with SED drives result in the controller losing one logical request object for each physical SED drive in the array.
  - Root Cause: When a logical drive is created with SEDs, each SED in this new logical drive will go through the SED qualification process, and it allocates a logical request object. Upon successful completion, firmware was not freeing up the allocated object.
  - Fix: Firmware will free up the allocated object before exiting the SED qualification process.
  - Risk: Low

© 2022 Microchip Technology Inc.
and its subsidiaries

- Fixed a firmware hang issue that can occur with logical drive expansion or spare assignment process on logical drives created with SEDs.
  - Root Cause: When a SED is added to an array as part of an array expansion or spare assignment process, the newly added SED goes through the SED qualification process, and the host IOs for the drive are put in a retry queue. Due to incorrect logic, firmware can never process the retry queue and becomes stuck in a loop.
  - Fix: Change the firmware to process the pending commands in the retry queue after the SED qualification is complete.
  - Risk: Low
- Fixed a controller lockup when deleting logical drive via PLDM over MCTP.
  - Root Cause: The count for the outstanding command was being decremented too soon, causing an input validation check to fail.
  - Fix: Updated to only decrement out of band commands count if request is a non-host request. The host transport layer will decrement the count for host requests.
  - Risk: Low
- Fixed a reporting error on sense low temperature at -5 °C.
  - Root Cause: The firmware logic for temperature sensor EMC 1414 and EMC 1464 assume that a value of 0 or 0xFF is invalid. By default, the sensor value can be 0 to 127. If the actual value, is below 0 then it's reported as 0. If the actual value is above 127, it is reported as 127. In normal operation, the temperature sensors never see 0 or below.
  - Fix: Allows 0 to be a valid value.
  - Risk: Low

2.2.2 UEFI Fixes

Note: Microsoft signed and secure boot is supported**.**

2.2.2.1 Fixes and Enhancements for UEFI Driver 1.3.17.1/Legacy BIOS 1.3.17.2

This release includes the following UEFI fixes and enhancements:

- Fixed an issue where the 0x1789 driver health error code is displayed as healthy when the controller is in Abnormal Volume state where all logical drives including healthy ones are offline.
  - Root Cause: Currently, the 0x1789 post message is triggered as healthy when all logical drives are made offline.
  - Fix: After sending bypass volume command if any of the logical drive state is still loose cable then treat it as Abnormal Volume state and set Driver Health state as configuration required.
  - Risk: Low
- Fixed an issue where there is no information in HII when the controller is in Abnormal Volume state.
  - Root Cause: No field in the HII controller information indicates the status that the controller is in Abnormal Volume state when all logical drives are offline.
  - Fix: Added new display field **Controller Configuration Status** under **Controller Information** menu to represent errors in controller due to configuration.
  - Risk: Low
- Fixed an issue where `ExtendedSCSIPassThruProtocolTest` is failing in SCT.
  - Root Cause: Mismatch in comparison to Target ID. UEFI driver does not validate complete 16 byte Target ID input.
  - Fix: Updated protocol function to compare complete 16 byte Target ID.
  - Risk: Low
- Fixed an issue where the Erase operation on an SMR-HM physical drive fails.
  - Root Cause: SMR-HM drives are supported as HBA targets, but not supported for RAID configurations or erase operations by the controller firmware.
  - Fix: Software will not send erase commands to SMR-HM physical drives.
  - Risk: Low

2.2.3 Driver Fixes

2.2.3.1 Fixes and Enhancements for Linux Driver Build 2.1.16-030

This release includes the following fixes and enhancements.

- Fixed an issue where the removal of a drive from the OS could be delayed up to 30 seconds after being physically pulled.
  - Root Cause: The driver was retrying a LUN reset three times even though the return code indicated the LUN was no longer valid. There was a 10 seconds delay between each retry. Additionally, the rescan worker was scheduled to run 10 seconds after the driver received the event.
  - Fix: Check the response code returned from the LUN Reset Task Management function and if it indicates the LUN is not valid, do not retry. Also, reduced the delay of the rescan worker to 5 seconds for the event handler only.
  - Risk: Low
- Fixed an issue where the outstanding requests and accessing the SCSI attributes for a device post linkdown leads to a hang.
  - Root Cause: Post linkdown, driver does not fail the outstanding requests leading to long wait time before all the I/Os eventually fail. Also, access to the SCSI attributes by the host applications lead to a system hang.
  - Fix: Sanity check added in various functions to block the host applications from accessing the SCSI attributes of a device when the controller goes offline.
  - Risk: Low
- Fixed an issue of disabling a controller by writing to this node in

  ```
  sysfs: /sys/bus/pci/devices/<Domain>:<Bus>:<Device>.<Function>/remove
  ```

  can cause processes with outstanding I/O to the controller or that attempt to access device attributes via sysfs for devices exposed by the SmartPQI driver to appear hung.
  - Root Cause: Disabling the controller using this technique can cause outstanding I/O requests to take a long time to fail, which can cause processes accessing the controller to appear hung.
  - Fix: The driver now promptly fails outstanding I/O requests and the device attribute accesses for a controller is removed using this technique.
  - Risk: Low
- Fixed an issue where the Linux kernel was not creating symbolic links in sysfs between SATA drives and their enclosure.
  - Root Cause: The driver was enabling the UNIQUE_WWID_IN_REPORT_PHYS_LUN PQI feature, which causes the firmware to return a WWID for SATA drives that is derived from a unique ID read from the SATA drive itself. The driver was exposing this WWID as the drive's SAS address. However, because this SAS address does not match the SAS address returned by an enclosure's SES Page 0xA data, the Linux kernel was unable to match a SATA drive with its enclosure.
  - Fix: The driver no longer enables the UNIQUE_WWID_IN_REPORT_PHYS_LUN PQI feature as it is not required.
  - Risk: Low
- Fixed an issue of inconsistent performance in RAID 10 logical drives when performing 256K sequential reads.
  - Root Cause: The driver was only using a single tracker to determine which physical drive to send a request to for IOBypass requests.
  - Fix: Allocate an array of trackers based on the number of data disks in a row of the RAID map.
  - Risk: Medium
- Fixed an issue where the controller boot timeout error message displays wrong number of seconds.
  - Root Cause: A recent change to the driver inadvertently changed the units displayed from seconds to Linux kernel “jiffies”.
  - Fix: Changed the units displayed by the error message back to seconds.
  - Risk: Low

2.2.3.2 Fixes and Enhancements for FreeBSD Driver Build 4210.0.1004

This release includes the following enhancements and fixes:

- Added support for 64 bit address. Changed 32 bit DMA address to 64 bit.
- Fixed an issue of inconsistent performance on RAID 10 logical drives when performing 256K sequential reads.
  - Root cause: The driver was only using a single tracker to determine which physical drive to send a request to for IOBypass requests.
  - Fix: Allocate an array of trackers based on the number of data disks in a row of the RAID map.
  - Risk: Medium

2.2.3.3 Fixes and Enhancements for Solaris Driver Build 11.4120.0.1005

There are no known fixes for this release.

2.2.3.4 Fixes and Enhancements for Windows Build 1010.24.0.1005

- Added support for the registry value “PostTimeOut”. Adding this registry value allows a user to override the default controller post timeout value of 3 minutes. The valid range is 30–1800 seconds.
  - Description: In some rare configurations, the controller can take more than 5 minutes to post. For these system, the registry value “PostTimeOut” can be set in seconds to something greater than 5 minutes to prevent the driver from timing out during the controller post.

    Registry location:

    Per Controller—`HKLM\CurrentControlSet\Enum\PCI\<Instance path>\DeviceParameters\SmartPqi`
- Fixed an issue where driver crashes when firmware encounters a lockup.
  - Root Cause: The driver is accessing the device records that are not valid while handling IOCTL, because the resources are freed while the lockup is detected.
  - Fix: The driver will not allow the IOCTLs directed to the controller or the drives after the HBA becomes failed. Clear the target offset from the “trie” node if the controller is failed.
  - Risk: Low
- Fixed a race condition in the driver LUN Reset handler that can result in accessing a bad pointer under heavy I/O Reset stress testing.
  - Root Cause: A changing driver resource race condition in the LUN Reset handler can lead to the driver referencing a bad pointer.
  - Fix: Removed the logic in the LUN Reset handler that resulted in accessing the bad pointer.
  - Risk: Low
- Fixed an issue of inconsistent performance in RAID 10 logical drives when performing 256K sequential reads.
  - Root Cause: The driver was only using a single tracker to determine which physical drive to send a request to for IOBypass requests.
  - Fix: Allocate an array of trackers based on the number of data disks in a row of the RAID map.
  - Risk: Medium

2.2.3.5 Fixes and Enhancements for VMware Driver Build 4252.0.103

This release includes the following enhancements and fixes:

- Fixed an issue where driver load/unload test was failing due to multiple issues with the driver error handling during initialization and topology rescan. Issues include management command timeout, memory leak, and PSOD.
  - Root Cause:
    - Management command time-out: MSIX interrupts were not getting enabled when driver initialization fails after enabling legacy interrupt.
    - Memory leak: Memory allocated was not getting freed when the following step fails.
    - PSOD: Due to a page fault, memory allocation is getting failed and the return value was not getting checked correctly.
  - Fix: Corrected the error handling cases during driver initialization and rescan paths.

  - Risk: Medium
- Fixed an issue of inconsistent performance in RAID 10 logical drives when performing 256K sequential reads.
  - Root Cause: The driver was only using a single tracker to determine which physical drive to send a request to for IOBypass requests.
  - Fix: Allocate an array of trackers based on the number of data disks in a row of the RAID map.
  - Risk: Medium

2.2.4 Management Software Fixes

2.2.4.1 Fixes and Enhancements for Arcconf/maxView Build B24763

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

- Support for logical device expansion without backup power source.
- Upgraded the log4j version to the latest version.
  Note: Microchip strongly recommends users of maxView update to the latest version of the tool to avoid the security vulnerabilities with the previous releases.
- Fixed an issue where the domain authentication is not working for certain admin users in maxView.
  - Root Cause: Windows APIs failed to identify the user as admin resulting in not allowing the user to perform admin level operation in maxView.
  - Fix: Added changes to call a new Windows API that identifies the user to be part of the admin group which allows them to operate in maxView.
  - Risk: Low
- Fixed an issue where arcconf/maxView displayed negative value in some error counters for certain hard drives.
  - Root Cause: arcconf/maxView displayed the negative value for certain error counter value as it is held by a smaller data type.
  - Fix: Added changes to hold the error counters in relevant data types to display it correctly.
  - Risk: Low
- Fixed an issue where the user could not stop the maxView Webserver service.
  - Root Cause: Conditions to stop the maxView webserver failed to allow stopping of service.
  - Fix: Added changes to allow the stopping of maxView Webserver service when relevant conditions are met.
  - Risk: Low
- Fixed an issue where maxView displays “XML Error intermitted” error.
  - Root Cause: Protected view exception was not handled in maxView which resulted in the XML error.
  - Fix: Added changes to handle protected view exception in maxView.
  - Risk: Low
- Fixed an issue where arcconf does not display the disk name information for the hard drive.
  - Root Cause: arcconf failed to display the disk name information for a RAW hard drive connected on a controller in Mixed mode.
  - Fix: Added changes to display the disk name information for a RAW hard drive on a controller in Mixed mode.
  - Risk: Low
- Fixed an issue where arcconf SETCONFIG command does not set all the controller parameters to factory default setting.
  - Root Cause: arcconf SETCONFIG command does not set the controller write cache policy settings to default values.
  - Fix: Added changes for arcconf SETCONFIG command to set the controller parameters to factory default settings.
  - Risk: Low
- Fixed an issue where firmware logs cannot be differentiated in a multi-controller environment.
  - Root Cause: In a multi-controller environment where controller’s physical slot is UNKNOWN, the firmware logs naming convention can differentiate where the log is collected from.

  - Fix: Added changes to add a controller ID along with the slot ID in firmware log names to identify the logs correctly.
  - Risk: Low
- Fixed an issue where arcconf failed to display the PHY error log information for an expander.
  - Root Cause: Displaying expander PHY errors is disabled in arcconf.
  - Fix: Added changes to enable expander PHY error logs from arcconf.
  - Risk: Low

2.3 Limitations

2.3.1 Firmware Limitations

2.3.1.1 Limitations for Firmware Release 5.00 B0

This release includes the following firmware limitations:

- A firmware update causes the UART log buffer (Serial Output Buffer) to be reinitialized, since the DDR gets reinitialized.
  - Workaround: None
- SATA drives attached to a non-Microchip expander may get into a failed state when upgrading the controller firmware from previous releases to this release due to the expander not clearing STP affiliation.
  - Workaround: Power cycle the expanders to clear the STP affiliation.
- A rare corner-case scenario where controller may hang during expander firmware update on multi-level expander/SEP device topology along with I/Os.
  - Workaround: After the enclosure firmware update, avoid enclosure Reset. It is recommended to download the new firmware and perform manual power cycle. This issue is intermittent and can cause a hang that requires a system reboot.
    Note:
    This issue was most frequently observed when using Linux based OS.
- Controller cache will not be converted into 100% read cache, if any backup power source cable error, charge or charge timeout error occurs when expansion or transformation task is active.
  - Workaround: None

2.3.1.2 Limitations for Firmware Release 1.32 Build 0

- Firmware release 1.32b0 may become unresponsive while attempting to flash firmware or execute other RAID logical drive operations.
  - Description: Refer to entry "Fixed an issue where firmware may become unresponsive while attempting to flash firmware or execute other RAID logical drive operations" in the Firmware fixes section.
  - A fix for this issue is available in the 1.60 B0 firmware release. If a firmware flash failure is occurring, try the following workarounds:
    - Workaround: If there are no target devices (expanders or drives) attached to the controller, attach a target device to the controller and try the host management operation again.
    - Workaround: If the system is operating using UEFI, the HII tool can be used to flash the firmware to this release as outlined in the *Microchip SmartIOC 2100/SmartROC 3100 Installation and User's Guide (ESC-2170577),* appendix entry “Updating the SmartIOC 2100/SmartROC 3100 Controller Firmware”.
    - Workaround: If there are target devices attached to the controller and this issue occurs or none of the workarounds can be used, contact Microchip Support.

2.3.2 UEFI Limitations

2.3.2.1 Limitations for UEFI Build 1.3.17.1/Legacy BIOS Build 1.3.17.2

There are no known limitations for this release.

2.3.3 Driver Limitations

2.3.3.1 Limitations for Linux Driver Build 2.1.16-030

This release has the following Linux limitations:

- The controller might lockup with the Inbox driver version 1.2.16-012 and RHEL 8.4 OS. The SCSI mid-layer might send requests beyond the exposed host queue depth resulting in a firmware assert.
  - Workaround:
    - The fix for this issue has been submitted to the open source community and it is part of the latest driver version.
- When doing a driver injection (DUD) install. On some distributions (RHEL7.9, RHEL8.2, RHEL8.3, SLES15SP2, and SLES15SP3), the DUD install hangs if a drive in HBA mode has the Drive Write Cache enabled.
  - Workaround: There are two workarounds for this issue:
    - Ensure the Drive Write Cache is disabled for any drive in HBA mode.
    - For RHEL7.9, RHEL8.2, and RHEL8.3, add `rd.driver.blacklist=smartpqi` to the grub entry along with `inst.dd`.
- Due to a change in the SCSI mid-layer, some Linux distributions may take a long time to come up if the system is rebooted while a hard disk(s) is being sanitized. This has currently been observed with inbox smartPQI drivers on RHEL 7.9, RHEL8.3, and SLES 15SP2.
  - Workaround: Do not reboot the system while a hard disk(s) is being sanitized or update to the SmartPQI 2.1.12-055 or later driver release.
- On AMD/RHEL 7.9 systems, the system might panic due to an issue in the IOMMU module. For more information, see lore.kernel.org/linux-iommu/20191018093830.GA26328@suse.de/t/
  - Workaround: Disable the IOMMU setting option in BIOS.
- On AMD/UEK6 systems, the system might hang during kdump if IOMMU is enabled.
  - Workaround: Disable IOMMU setting option in BIOS.
- The `smartpqi.expose_ld_first` parameter does not work correctly 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. Use the following to increase the timeout for udev(systemd):

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

2.3.3.2 Limitations for Windows Driver Build 1010.24.0.1005

This release includes the following limitations:

- If the controller encounters a lockup and if the windows host system is restarted within 30 seconds, a BSOD might happen.
  - Workaround: If the controller encounters a lockup, restart the windows host system after 30 seconds.
- 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, such as speed of drives, size of cache, type of data in cache, etc. 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.
  - Workaround: 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.

2.3.3.3 Limitations for FreeBSD Driver Build 4210.0.1004

This release includes the following FreeBSD driver limitations:

- Under heavy I/O with transfer size more than 128k , controller may go offline. This happens in FreeBSD 13.

– Workaround: Reduce the I/O transfer size of the application to less than 128k.

2.3.3.4 Limitations for Solaris Driver Build 11.4120.0.1005

There are no known limitations for this release.

2.3.3.5 Limitations for VMware Driver Build 4252.0.103

There are no known limitations for this release.

2.3.4 Management Software Limitations

2.3.4.1 Limitations for Arcconf/maxView Build B24763

This release includes the following Arcconf/maxView limitations:

- Logical drive creation fails when maxCache exists and supercap is removed.

2.3.5 Hardware Limitations

This release includes the following hardware limitations:

- Two Wire Interface (TWI) address conflicts can cause system DDR memory to not be discovered.
  - Description: The SmartRAID 3100 and SmartHBA 2100 boards include two TWI targets on the host-facing SMBUS interface with the following slave addresses:
    - 0xA0 – Field Replaceable Unit (FRU) SEEPROM
    - 0xDE – PBSI (default)

      According to the JEDEC specification, the default TWI addresses for the DDR SPD is 0xA0-0xAE (the spec uses 7 bit addressing which is 0x50-0x57). On platform system board designs with SMBUS wiring that has both PCIe slots and DDR slots shared on the same TWI bus, the TWI devices for the DDR and Smart controller are exposed to address conflicts which can result in the system memory not being discovered. The Smart controller PBSI interface defaults to a value of 0xDE (0x6F in 7-bit addressing) and is not a problem unless it is changed to an address that conflicts with the JEDEC defined values. The Smart controller FRU SEEPROM is hardwired to 0xA0.
  - Workaround: None available. If this issue is encountered, contact your Microchip support engineer to determine the next steps for your system.
  - Performance with workaround: Not applicable
  - Performance without workaround: Not applicable

3. Updating the Controller Firmware

This section describes how to update the board's firmware components to the latest release.

3.1 Updating the Controller Firmware

This procedure describes how to prepare your board to be programmed with the latest firmware.

Notes:

1. If the running firmware is older than 1.98 and a transformation is in progress, complete the transformation before proceeding with the following steps to upgrade the firmware.
2. Complete these procedures exactly as described for proper functionality. If you do not follow all of the steps correctly, you could encounter unusual runtime behavior.

Flashing the board to the latest firmware:
This section describes how to update all the firmware components on Adaptec controller boards to the latest release.

If the controller is currently running 1.60 b0 firmware or newer, follow these steps:

1. **Mandatory:** Flash the target with the provided " SmartFWx100.bin" image with arcconf/maxView software.
2. **Mandatory:** Use the OS shutdown/restart operation to gracefully reboot the system to complete the firmware update process.

Note:
After completing the firmware update, if the firmware version is still showing the prior version, retry the firmware update steps.

If the controller is currently running 1.32 b0 firmware, follow these steps:

1. **Mandatory:** Flash the target with the provided “SmartFWx100.bin" image with arcconf/maxView software.
   – If the arcconf/maxView software becomes unresponsive or hangs then power cycle the system to recover and refer to firmware limitation section 2.3.1.2. Limitations for Firmware Release 1.32 Build 0.
2. **Mandatory:** If flashing completes, use the OS shutdown/restart operation to gracefully reboot the system to complete the firmware update process.

Note:
After completing the firmware update, if the firmware version is still showing the prior version, retry the firmware update steps.

If the controller is currently running 1.04 b0 firmware, follow these steps:

1. **Mandatory:** Flash the controller with the provided “SmartFWx100_ v1.29_b314.bin" image with arcconf/maxView software.
2. **Mandatory:** Reboot the system to refresh all components**.**
3. **Mandatory**: Flash the target with the provided " SmartFWx100.bin" image with arcconf/maxView software.
4. **Mandatory**: Use the OS shutdown/restart operation to gracefully reboot the system to complete the firmware update process.

At this point, the controller would be updated and would be ready to use. Install the SmartPQI driver and the latest version of the Arcconf/maxView management utility to monitor and configure the controller.

Note: Downgrading firmware could lead to unexpected behavior due to an incompatibility in SEEPROMs between this release and the prior release.

4. Installing the Drivers

See the “Microchip Adaptec® SmartRAID 3100 Series and SmartHBA 2100 Series Host Bus Adapters Installation and User's Guide (DS00004258A, previously ESC-2171547)” for complete driver installation instructions.

5. Revision History

The revision history describes the changes that were implemented in the document. The changes are listed by revision, starting with the most current publication.

Revision	Date	Description
D	03/2022	VMware driver version updated from 4250.0.120 to 4252.0.103
C	02/2022	SR 2.6.6 Production Release
B	12/2021	SR 2.6.4.1 Patch Release with maxView™ version B24713. Updated Fixes and Enhancements for maxView Storage Manager/ARCCONF section for log4j vulnerabilities.
A	11/2021	SR 2.6.4 Production Release with firmware version 4.72 B0 (Previously ESC-2161026)
29	04/2021	SR 2.6.2 with firmware version 4.11 B0
28	04/2021	SR 2.6.1.1 with VMware driver version 4054.2.118.
27	03/2021	SR 2.6.1 with VMware driver version 4054.1.103.
26	02/2021	SR 2.6 Production Release
25	10/2020	SR 2.5.4 Production Release
24	08/2020	SR 2.5.2.2 Production Release with Firmware 3.00
23	03/2020	SR 2.5.2 Production Release with Firmware 2.93
22	03/2020	SR 2.5 Production Release with Firmware 2.66
21	02/2020	SR 2.5.2 Production Release
20	10/2019	SR 2.5 Production Release
19	09/2019	Updated for SR 2.4.8.1 (fw v2.31 Build 0)
18	08/2019	Updated for SR 2.4.8
17	01/2019	SR2.4 Production Release
16	06/2018	SR2.3 Production Release
15	06/2018	Updated for RC Release
14	10/2017	Update supported OSs
13	10/2017	First Production Release
1-12	06/2016 to 07/2017	Pre-Production Release.

© 2022 Microchip Technology Inc.
and its subsidiaries

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, AnyRate, 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, IntelliMOS, Libero, motorBench, mTouch, Powermite 3, Precision Edge, ProASIC, ProASIC Plus, ProASIC Plus logo, Quiet- Wire, SmartFusion, SyncWorld, Temux, TimeCesium, TimeHub, TimePictra, TimeProvider, TrueTime, WinPath, 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, 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, Inter-Chip Connectivity, JitterBlocker, Knob-on-Display, maxCrypto, maxView, memBrain, Mindi, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, NVM Express, NVMe, 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, SQI, SuperSwitcher, SuperSwitcher II, Switchtec, SynchroPHY, Total Endurance, 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, Symmcom, and Trusted Time 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.

© 2022, Microchip Technology Incorporated and its subsidiaries. All Rights Reserved.

ISBN: 978-1-5224-9847-6

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

© 2022 Microchip Technology Inc.
and its subsidiaries