SmartHBA 2100 and SmartRAID 3100 Software/Firmware Release Notes

Released / February 2021

Revision History : February 2021

Revision	Revision Date	Details of Change
26	February 2021	SR 2.6 Production Release
25	October 2020	SR 2.5.4 Production Release
24	August 2020	SR 2.5.2.2 Production Release with Firmware 3.00
23	March 2020	SR 2.5.2 Production Release with Firmware 2.93
22	March 2020	SR 2.5 Production Release with Firmware 2.66
21	February 2020	SR 2.5.2 Production Release
20	October 2019	SR 2.5 Production Release
19	September 2019	Updated for SR 2.4.8.1 (fw v2.31 Build 0)
18	August 2019	Updated for SR 2.4.8
17	January 2019	SR2.4 Production Release
16	June 2018	SR2.3 Production Release
15	June 2018	Updated for RC Release
14	October 2017	Update supported OSs
13	October 13, 2017	First Production Release
1-12	June 2016-July 2017	Pre-Production Releases.

Table of Contents

1 About This Release

The development release described in this document includes firmware, OS drivers, tools, and host management software for the SmartHBA 2100/SmartRAID 3100 controller 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
Package Release Date	February 22, 2021
Firmware Version	3.53 B0[1,2] (basecode 06.06.001.001)
UEFI Version	1.3.12.2
Legacy BIOS	1.3.12.3
Driver Versions[3]	Windows SmartPQI: • Windows 2012/2016/2019: 106.278.0.1043 Linux SmartPQI: • RHEL 6/7/8: 2.1.8-040 • SLES 12/15: 2.1.8-040 • Ubuntu 16/18/20: 2.1.8-040 • Debian 9/10: 2.1.8-040 • CentOS 6/7/8: 2.1.8-040 • Oracle Linux 7/8: 2.1.8-040 • Citrix XenServer 7/8: 2.1.8-040 VMware SmartPQI: • VMware 6.5/6.7/7.0: 4054.0.131 FreeBSD/Solaris SmartPQI: • FreeBSD 11/12: 4054.0.1005 • Solaris 11: 4044.0.1011
Management Software (arcconf, maxView™, Event Monitor, BootUSB)	B23971

Note:

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 " *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 volumes 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 volumes will need to be recreated.
3. Only run the driver on firmware 0.01 build 500 or later.

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 SmartHBA 2100/SmartRAID 3100 controller and SmartRAID 3100 and SmartRAID 3100 controller solutions from the Microsemi Web site at *https://storage.microsemi.com/en-us/support/start/*

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

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 • Host Management Utilities* on page 4
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 • Host Management Utilities* on page 4

Driver Files

Table 1-4 • Windows Storport Miniport SmartPQI Drivers

Package	Drivers	Binary	Version
2012	Server 2019 Server 2016 and Windows 10 Server 2012 SP1, R2 SP1 and Windows 8.1, 8	SmartPqi.sys	x64
		SmartPqi.inf	x64
		Smartpqi.cat	x64

Table 1-5 • Linux SmartPQI Drivers

Drivers	Version
Red Hat Enterprise Linux 8.3, 8.2, 8.1, 8.0, 7.9, 7.8, 7.7, 7.6, 6.10[1]	x64
CentOS 8.2, 8.1, 8.0, 7.9, 7.8, 7.7, 7.6, 6.10	x64
SuSE Linux Enterprise Server 12[1], SP5, SP4, SP3	x64
SuSE Linux Enterprise Server 15 SP2, SP1[1]	x64
Oracle Linux 7.6 with UEK5u2 (4.14.35)	x64
Oracle Linux 7.7 with UEK5u2 (4.14.35)	x64
Oracle Linux 7.9, 7.8 UEK 5u2	x64

Drivers	Version
Oracle Linux 8.1 UEK6	x64
Oracle Linux 8.3, 8.2 UEK6 U1	x64
Ubuntu 20.04.1, 20.04	x64
Ubuntu 18.04.5, 18.04.4, 18.04	x64
Ubuntu 16.04.5	x64
Debian 10.04	x64
Debian 9.12	x64
Citrix xenServer 8.2, 8.1, 8.0, 7.6	x64
Fedora 30 (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 12.2, 11.4	x64
Solaris 11.4, 11.3	x64
VMware 6.7 U3/U2/U1, 6.5 U3/U2/U1	x64
VMware 7.0 U1	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.
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.

Description	OS	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 ARCC-ONF and maxView Storage Manager	Linux x64	See the maxView BootUSB download package for the .iso file.

2 What is New?

2.1 Features

The following table lists features supported for this release.

Table 2-8 • 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	

2.2 Fixes

2.2.1 Firmware Fixes

2.2.1.1 Fixes and Enhancements for Firmware Release 3.53 B0

This release includes the following fixes and enhancements:

- Added support for persistent firmware logs (Serial Output Buffer) across warm-resets.

- Added a new API to allow upper-layer firmware to retrieve channel type (I2C or VDM) of MCTP messages before dispatching.
- Added support for long device model/product ID for SATA drives.
- Added support for user configurable PCIe Relaxed ordering mode setting.
- Added support for retrieving backup power source status information through Out-of-Band MCTP transport.
- Added capability for firmware to advertise to host software that a reboot is required to complete a configuration change, and whether that reboot needs to be a cold boot; for example, situations such as CPLD update, power mode changes, OOB interface changes, backplane discovery policy.
- Fixed an issue where the controller becomes unresponsive while waiting for a non-existent pre-fetch queue entry.
  - Root Cause: Firmware is unable to retrieve an entry from one of the pre-fetched hardware queues and becomes unresponsive, resulting in a controller hang.
  - Fix: Disable pre-fetch on the hardware queue, as this is not a queue needed for performance.
  - Risk: Low
- Fixed an issue where host timeouts occur due to continuing firmware attempts to discover devices during an expander configuration.
  - Root Cause: When an expander configuration is in progress, firmware retries forever until the configuration is complete to discover devices attached to the expander. This leads to an IO pile up causing host timeouts.
  - Fix:
    1. Add a 15 second limit for lower level firmware to retry the discovery when an expander is detected in configuration.
    2. Report removal of this expander and the topology behind it to upper layer firmware.
    3. Poll the expander once every 2 seconds, since the expander may not send a broadcast change when it is done configuring.
  - Risk: Low
- Fixed a controller lockup issue (with code 0x1E10) when a bad drive with unsupported block size is connected into slot 0.
  - Root Cause: Firmware is incorrectly trying to allocate zero block of data for a request from host tools when a bad drive is connected. This request is initiated only for slot 0.
  - Fix: Fail bad drives with unsupported block sizes earlier during discovery itself.
  - Risk: Low
- Fixed a controller TLB exception (NULL access) after changing all connector mode to HBA.
  - Root Cause: A corner case scenario where all but one of the controller connectors are in HBA mode and changing the last connector from Mixed/RAID mode to HBA mode, firmware tries to access the logical drive meta-data structure for a periodic background task, but meta data was just cleared out when controller was transitioned into HBA mode.
  - Fix: Skip accessing the meta-data when it is already cleared.
  - Risk: Low
- Fixed a PBSI issue where the Predictive Failure (PF) Drive is not reported correctly when configured in a RAID 0 logical drive.
  - Root Cause: Firmware was referring to LED control bits to report Predictive Fail drive instead of the actual predictive errors reported and for RAID 0 logical drive, the LED control bits were not set.
  - Fix: Change PBSI code to refer to actual PF error status instead of LED control bits.
  - Risk: Low
- Fixed an issue where firmware was not providing the response data for a few commands sent through Out-of-Band MCTP transport.

    - Root Cause: Firmware did not set up and populate the response data for the commands, which has SCSI direction attribute as both IN and OUT; instead it provided only the status response.
    - Fix: Set up the response buffer appropriately for these Out-of-Band MCTP commands.
    - Risk: Low
- Fixed an issue where controller hangs when hot-plug and hot remove events were being processed.
    - Root Cause: While handling hot-plug and hot removal of different devices, two different tasks handling these events got into a dead lock situation, resulting in commands not being sent to targets and reset event completions not being handled.
    - Fix: Allow the corresponding thread to submit commands to devices even when there are unprocessed hot removal events in the queue.
    - Risk: Low
- Fixed an issue where controller hangs on DDR caching enabled logical volumes with sequential read workload of 1MiB.
    - Root Cause: When the read workload reaches towards the end of the volume, the read ahead logic in firmware does not unlock the cache lines. When the request is large enough to bypass the cache, the firmware is stuck in invalidating the cache line indefinitely due to locked lines.
    - Fix: Unlock affected cache lines.
    - Risk: Low
- Fixed a PBSI issue where incorrect details are provided when there are Data Set Address Pointer (DSAP) I2C writes without STOP bit.
    - Root Cause: While updating the DSAP write, firmware looks for START and STOP bit. Once the request fails with no STOP bit, the I2C client issues the writes with STOP bit again; but the firmware invalidates this request due to incorrect calculation of total bytes received.
    - Fix: Fix the bytes-received count so the new DSAP is written properly into internal Data Set Table (DST).
    - Risk: Low
- Fixed a controller hang issue with outstanding IOs on a degraded logical volume.
    - Root Cause: When a physical drive is hot removed, the outstanding IOs on the failed drive will be returned with error status. Due to a race condition between the thread which handles drive removal and the error handling thread, there are a few requests getting re-queued again after handling the failure of the drive. These re-queued requests are not handled further because the drive is in the failed state.
    - Fix: Check the drive failure status before re-queuing the requests.
    - Risk: Medium
- Fixed a potential controller hang if a host IO and background consistency check simultaneously encounter a RAID-1 ADM stripe in which all drives have URE's on the same LBA.
    - Root Cause: While the consistency checking logic is checking the stripe, it holds the stripe lock. When that check subsequently fails in error handling due to the URE's, it checks a flag on the request to determine whether this attempt was a first try or a retry, clears some of the error counter information associated with the request, because it was a first attempt, then returns status to the consistency checking function. This function was incorrectly holding the retry status in a local variable and then enters a loop in which the request is retried, which causes this activity to enter a never-ending loop in which the error recovery task always thinks it's dealing with a first attempt instead of a subsequent retry, which would have preserved the error counting information and broken the loop.

      Note: Requires 3-drive RAID-1 logical drive and a URE on the same LBA of the different drives.

    - Fix: After the initial first attempt fails and returns to the consistency check process, mark the flag in the request to indicate subsequent operations on this stripe are retries so that the first unsuccessful attempt and retry of error recovery properly breaks out of the loop.
    - Risk: Low

- Fixed a controller hang issue when a drive is failed from a RAID6/60 logical drive when the host issues a Clear Controller Configuration command or any other configuration change command.
  - Root Cause: When a physical drive is hot removed in a RAID6/60 logical drive, the background parity consistency checking code incorrectly puts the background thread into sleep mode; but the code responsible for processing configuration change commands is trying in parallel to suspend this thread, resulting in the command not getting completed.
  - Fix: Correct the RAID 6 background consistency checking code to not put background thread into sleep mode when it is not necessary.
  - Risk: Low
- Fixed an issue where the SSD data drive is set offline (with reason code 0x37) when hot-plugged during a spare rebuild in progress.
  - Root Cause: When SSDs supporting TRIM/UNMAP commands are used in logical drive, firmware initiates these commands before actual rebuild starts. While these TRIM/UNMAP commands are in progress and the data drive is re-inserted, instead of deactivating the spare, the firmware fails the data drive with reason code 0x37.
  - Fix: Deactivate the spare which, in turn, will start the rebuild process on the data drive.
  - Risk: Low
- Fixed an issue where sanitize erase completion status is not preserved across system reboot.
  - Root Cause: When the sanitize erase process is completed on a drive, the final status is not saved in the RAID meta-data at that time. If the system is rebooted just after the sanitize erase finishes, the drive was not exposed to the OS because firmware is attempting to restart the sanitize erase operation.
  - Fix: Save the status information in the drive RAID meta-data when the sanitize erase is complete.
  - Risk: Low
- Fixed an issue where controller hangs when processing idle time followed by short burst of IOs.
  - Root Cause: When the controller processes idle time followed by short burst of IOs, it allocates memory required to fetch the IO data and the transfer buffer memory pool becomes exhausted; as a result, further IO processing stalls, leading to a controller hang.
  - Fix: Tune the transfer buffer max memory pool size during controller boot to avoid memory pool exhaustion that caused the controller hang.
  - Risk: Medium
- Fixed an issue where the controller can occasionally return previous drive firmware version, after a drive firmware update on SATA drives.
  - Root Cause: After SATA drive firmware update, a host command query to fetch the updated drive firmware version can get processed early, before the controller firmware reads the new firmware version from the drive.
  - Fix: Modify the host command query path to send appropriate commands to the drive to fetch the latest drive firmware information.
  - Risk: Low
- Fixed an issue where the controller could fail drives (with reason code 0x49, IO freeze timeout) during expander firmware upgrade on a multi-expander enclosure configuration.
  - Root Cause: During expander firmware upgrade, a race condition creates redundant IO freeze timeout timers for a single drive. After expander firmware upgrade is completed, the IO freeze timeout timer is not deactivated, causing the drive to be failed incorrectly.
  - Fix: Add new handshake mechanism between firmware threads to prevent creating redundant IO freeze timeout timers.
  - Risk: Low
- Fixed an issue where a hot-added drive LED control fails on specific fan-out expander type external enclosure models.

- Root Cause: During hot insertion of drive on a fan out type multi-expander enclosure configuration, device slot/index is not enumerated correctly, leading to failure of LED control on those enclosures.
- Fix: Changes were added to properly enumerate the device slot/index, including SEP devices, so LED control works as expected in those enclosure models.
- Risk: Low

- Fixed an issue where CSMI SSP passthrough commands do not display any data.
  - Root Cause: When processing CSMI SSP passthrough commands, the controller incorrectly fills the transferred length field value as 0 for under run cases; as a result, the application fails to fetch the exported data.
  - Fix: Modify CSMI SSP passthrough return response code to fill in the actual number of transferred bytes.
  - Risk: Low
- Fixed an issue where controllers without DDR controller cache hang when processing sequential IOs.
  - Root Cause: Controllers without DDR controller cache can have firmware wait indefinitely for a memory allocation call to complete. The firmware coalescing logic is waiting to get memory allocated from an internal memory pool, but there isn't enough memory available so the controller firmware remains in a loop trying to get the memory. This stalls all IOs from being processed that results in the LUN Reset from the OS which does not complete.
  - Fix: If the firmware coalescing logic can't get the memory, then the coalescing operation does not proceed and the requested IOs are sent individually to the logical drive.
  - Risk: Low
- Fixed an issue where maxCache configurations that encounter intermittent write IO errors to a primary logical drive may cause the file system or application to read old data.
  - Root Cause: Intermittent IO errors during a host write operation to the primary logical drive in a maxCache configuration will result in retries. Before the retries are completed, maxCache could load old data into a maxCache page. As a result, subsequent reads to that page will return old data.
  - Fix: maxCache will load the data after the host write and any retries are completed.
  - 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 1.3.12.2/Legacy BIOS 1.3.12.3

This release includes the following UEFI fixes and enhancements:

- Fixed an issue where HII allows setting drive undergoing erase as legacy bootable when controller is set to Mixed port mode.
  - Root Cause: Firmware masks erasing drives from the OS if the controller is in Mixed mode, but exposes them if the controller is in HBA mode. This causes the exposed drives to be presented as valid boot targets.
  - Fix: The driver was not checking if a physical drive is being erased before allowing it to be set as a boot target. Add a check to verify that the drive is not erasing before exposing it as a valid boot target.
  - Risk: Low
- Fixed an issue where full-length drive model is not shown in HII Drive Information.
  - Root Cause: Buffer used to retrieve drive model is not big enough to get complete model string for drives which have larger drive model information.
  - Fix: Increase buffer size to accommodate complete drive model information.

- Risk: Low

- Fixed an issue in HII where Physical Drive is still shown in Disk Utilities even after a member drive of a Logical Drive is removed.
  - Root Cause: Drive presence not verified while listing drives in HII Disk Utilities.
  - Fix: Filter out missing drives while listing drives in HII Disk Utilities.
  - Risk: Low
- Fixed an issue where Volume Unique Identifier for Logical Drive is not displayed in HII menu for Logical drive details.
  - Root Cause: No field or data displayed for Volume Unique Identifier as part of Logical Drive details.
  - Fix: Added new field named Volume Unique Identifier under Logical drive details.
  - Risk: Low
- Fixed an issue in HII Disk Utilities where Physical Drive Device information is not listing all the array details if the drive is part of multiple arrays.
  - Root Cause: Even if Physical drive is part of multiple arrays, the information is retrieved only for first associated array.
  - Fix: Consider all array associations for the drive.
  - Risk: Low
- Fixed an issue where SCSI pass thru Test Unit Ready (TUR) commands fail.
  - Root Cause: Wrong transfer direction set for SCSI TUR command.
  - Fix: Transfer direction set as None for TUR command.
  - Risk: Low

2.2.3 Driver Fixes

2.2.3.1 Fixes and Enhancements for Linux Driver Build 2.1.8-040

This release provides the following fixes and enhancements.

- Fixed a firmware ASSERT issue when scsi-mid-layer sends requests that exceeded the exposed host queue depth.
  - Symptom: scsi-mid-layer sends requests that exceeded the exposed host queue depth, resulting in a firmware ASSERT.
  - Root Cause: Before submitting the IO request to the low-level driver, scsi-mid-layer used to check host queue depth for each IO. Due to a recent change in the kernel, that check has been removed.
  - Fix: Added host queue depth counter. Driver will return back IOs to OS if it exceeds the exposed host queue depth limit.
  - Risk: Low
- Fixed an issue where I/O requests to the disk were blocked before Synchronize Cache requests are issued.
  - Symptom: Unloading driver with drive write cache enabled for HBA SAS/SATA disks results in the following error message: 'Synchronize Cache(10) failed: Result: hostbyte=DID_NOT_CONNECT driverbyte=DRIVER_OK'.
  - Root Cause: The function pqi_device_remove_start is called early in pqi_remove_device. This blocks I/O requests to the disk before Synchronize Cache requests are issued when the device is actually removed from the OS. The driver returns DID_NOT_CONNECT after pqi_device_remove_start is called.
  - Fix: Call pqi_device_remove_start at the end of the device removal chain.
  - Risk: Low
- Fixed an issue where IOBypass read I/O requests were failing.

- Symptom: IOBypass read I/O requests that hit UREs fail.
- Root Cause: The driver used the retry count in SCSI request packets received from the OS to determine when an I/O request was a failed IOBypass request that should be retried. The count was not incremented by the OS on retries, as expected.
- Fix: The driver now uses a count variable, maintained by the driver, instead of relying on the retry count field maintained by the OS.
- Risk: Low

2.2.3.2 Fixes and Enhancements for FreeBSD Driver Build 4054.0.1005

This release provides the following enhancements and fixes:

- Fixed an issue causing a kernel panic during array creation and deletion in a loop.
  - Symptom: Observed system crash while creating and deleting logical drives in a loop.
  - Root Cause: Accessing device null pointer in pqisrc_show_sense_data_simple function.
  - Fix: Added Null check to prevent accessing empty device pointer.
  - Risk: Low
- Fixed an issue where the first logical drive is not shown in format command after creating the logical drive.
  - Symptom: After creating the logical drive, the first logical drive is not shown in format command.
  - Root Cause: The controller and first logical drive entries have the same Target and LUN numbers (T:0 L:1) but different bus number. The code considers only Target and LUN for its device node creation per adapter. A duplicate Target and LUN are created, causing inconsistent behavior.
  - Fix: Create unique Target and LUN numbers.
  - Risk: Medium

2.2.3.3 Fixes and Enhancements for Solaris Driver Build 4044.0.1011

This release provides the following enhancements and fixes:

- Fixed an issue where an incorrect RAID request is submitted to firmware, causing none of the drives to get discovered.
  - Symptom: Extra padding byte gets added in pqisrc_raid_request, resulting in passing an incorrect RAID request to firmware. As a result, none of the drives get discovered.
  - Root Cause: pqisrc_raid_request structure has 16-byte cdb at offset 32. It is defined as union of cbd[16] and 16 byte bmic_cdb structure. It is not packed, and extra padding byte gets added with Solaris sun studio gcc compiler.
  - Fix: Added packed attribute to avoid padding.
  - Risk: Low
- Fixed an issue where the first logical drive is not shown in format command after creating the logical drive.
  - Symptom: After logical drive creation, the first logical drive is not shown in format command.
  - Root Cause: The controller and first logical drive device entries have the same Target and LUN numbers (T:0 L:1) but different bus number. The driver considers only Target and LUN for its device node creation per adapter. A duplicate Target and LUN are created and cause inconsistent behaviour.
  - Fix: Create unique Target and LUN numbers.
  - Risk: Medium

2.2.3.4 Fixes and Enhancements for Windows Build 106.278.0.1043

This release provides the following enhancements and fixes:

- Updated the CSMI specification to version 1.13.

  Implementation Details: Added Segment Number to the CSMI IOCTL CC_CSMI_SAS_GET_CNTLR_CONFIG data buffer. The DOMAIN/SEGMENT number is used in Configuration Space addressing and is primarily a PLATFORM level construct. DOMAIN and SEGMENT are used here interchangeably (Domain tends to be the Linux term, Segment is the Windows and PCISIG term). Logically, SEGMENT is the most significant selector (most significant address bits selector) in the DOMAIN:Bus:Device:Function:Offset addressing scheme of the PCI Family.
- Fixed an issue where disk has outstanding commands while deleting per-lun memory.
  - Symptom: BSOD observed while creating and deleting logical volume with more than 100+ drives.
  - Root Cause: A null pointer access occurred when completing outstanding commands from completion queue due to the per-lun memory already being deleted from the driver store during driver rescan.
  - Fix: If the drive has pending commands then the per-lun memory will be deleted in next bus rescan.
  - Risk: Low
- Fixed sleep wakeup stage adapter initialization issues.
  - Symptom: The timer callback based adapter initiation causes watchdog timeouts while waking up the controller from sleep states.
  - Root Cause: The driver code initialized the adapter from a timer callback function when the system was waking up from sleep states. Since the Storport calls the TimerCallback function in dispatch level, lengthy adapter initialization caused watchdog timeout based errors in the system.
  - Fix: Removed Storport timer and work-item based initialization from the driver code. The adapter initialization code has been moved under the 'ScsiSetRunningConfig' control code during sleep state power transitions.
  - Risk: Medium
- Fixed a BSOD watchdog when draining submission queue.
  - Symptom: Running large IO with SATA drives connected behind expander causes blue screen with message DPC_WATCHDOG_VIOLATION (133). The driver asks for Storport to provide 1600 max I/O at any given time even though the controllers max is 1000. This was a performance enhancement to ensure all submission queues remain full. In certain max configurations, doing 4M transfers to SATA drives causes the controller to fall behind driver submission rate. As a result, the driver constantly queues commands.
  - Root Cause: The SmartPqi driver spends too much time in the ISR routine when attempting to drain the driver`s overflow submission queues. This condition is caused by the controller failing to keep up with processing the inbound submission queues. The slower SATA drives triggered this "back pressure" causing the driver to queue the excess commands to the driver overflow queue.
  - Fix: Added routine to grow the number of submission queue group elements to prevent back pressure when excessive I/O is used. Increasing the submission queues allows more room for commands that don't need to be queued, thus preventing the constant queuing.
  - Risk: Low
- Fixed an issue that causes a system crash if all phys are disabled.
  - Symptom: If phy_control disable all phys is initiated, system crashes.
  - Root Cause: When all the controller SAS phys were disabled and an outstanding IO was completed to the driver, an invalid pointer was used to increment an IO counter that resulted in a system crash.
  - Fix: Remove usage of the IO counters to avoid the invalid pointer access.
  - Risk: Low
- Fixed an issue where BSOD is observed after running I/O and getting LUN resets.

- Symptom: BSOD observed after running I/O and getting LUN resets.
- Root Cause: Tag was not released for the LUN reset command that was not successfully submitted due to a queue full condition.
- Fix: Driver will correctly release the command memory and tag for a failed LUN reset, due to queue full condition.
- Risk: Low

2.2.3.5 Fixes and Enhancements for VMware Driver Build 4054.0.131

This release provides the following enhancements and fixes:

- Fixed an issue where drive is identified as new device after reinserting drive in a different slot.
  - Symptom: Drive will be identified as new device after reinserting drive in a different slot.
  - Root cause: When drive is reinserted in a different slot, scsi3addr of drive changes with respect to slot. As a result, drive is detected as a new device.
  - Fix: Compare only the World Wide ID of the drive to find whether the hot-added device is present in the removed device list and, if present, update scsi3addr of the corresponding device entry in the device list with the new scsi3addr of the hot-added drive.

    Note: Detecting the same SATA drive moved between bays of the same controller only works if the SATA WWN Unique ID feature is enabled in the driver module parameters.
  - Risk: Medium
- Fixed an issue where system panics during array deletion.
  - Symptom: System panics during array deletion test.
  - Root cause: When ESXi destroys the device path, the driver sets a flag to indicate that the OS has already freed the scsi path, and driver can now safely free the device memory through the device discovery path. When the OS later re-added the path, the driver had not cleared the flag. When the same device has been removed from the system, the driver discovery path gets triggered and freed the path since the flag was set. After that, the OS destroys the path again, which resulted in double free and PSOD.
  - Fix: Clear the flag which indicates OS has freed the SCSI path.
  - Risk: Medium
- Fixed a PSOD while deleting a logical drive.
  - Symptom: Observed system crash while creating and deleting a logical drive in a loop.
  - Root cause: PSOD is due to the lack of heartbeat from the CPUs. There are two possible issues:
    1. Busy wait while waiting for pending IO completion.
    2. Busy wait is due to pending IO during device removal.
  - Fix: There are two steps required to fix this issue:
    1. Change busy wait to sleep.
    2. Decide if "scsi IO or not" is done for each response.
  - Risk: Medium

2.2.4 Management Software Fixes

2.2.4.1 Fixes and Enhancements for Arcconf/maxView Build B23971

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

- Add support to report the drive model with 40 character length.

- Fixed an issue where “SSD Wear Out level status” events are not registered to OS log, Event log, and SNMP Trap for the newly inserted drives.
  - Root Cause: The condition to generate the “SSD Wear Out level status” event was not handled for newly interested drives.
  - Fix: Added the condition to generate the “SSD Wear Out level status” event for newly interested drives.
  - Risk: Low
- Fixed an issue where Arcconf Move Array operation failed if device 0 is assigned as a member drive.
  - Root Cause: Duplicate hard drive check at the wrong place.
  - Fix: Changed where checks for duplicate hard drive entries are made.
  - Risk: Low
- Fixed an issue where the RAID configuration is different after saving and replaying configuration with Arcconf.
  - Root Cause: In the specific configuration, logical device was created on the wrong array due to the wrong array ID reference.
  - Fix: Added code to get proper array based on array ID.
  - Risk: Medium
- Fixed an issue in maxView where using the Domain Admin User results in non-Admin user rights.
  - Root Cause: The buffer used for authentication was not cleared after calling Windows API.
  - Fix: Cleared the buffer after calling Windows API.
  - Risk: High

2.3 Limitations

2.3.1 Firmware Limitations

2.3.1.1 Limitations for Firmware Release 3.53 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. Similarly, this buffer will not be persistent across cold reboots.
  - Workaround: None
- SATA drives attached to a non-Microsemi 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: Perform expander firmware update without I/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
- maxCache reconfiguration operations, such as changing from write back to write through or maxCache cache delete, may not complete successfully.
  - Workaround: Reboot controller and reconfigure maxCache cache, as required.
- While a rebuild occurs along with a heavy IO workload, LUN Resets may be observed. The IOs will be retried by the OS and the rebuild will continue to progress slowly and eventually complete.
  - Workaround: Change the rebuild priority setting to ‘low’.

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 volume operations.
  - Description: Refer to entry "Fixed an issue where firmware may become unresponsive while attempting to flash firmware or execute other RAID logical volume 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 *Microsemi 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 Microsemi Support.

2.3.2 UEFI Limitations

2.3.2.1 Limitations for UEFI Build 1.3.12.2/Legacy BIOS Build 1.3.12.3

There are no known limitations for this release.

2.3.3 Driver Limitations

2.3.3.1 Limitations for Linux Driver Build 2.1.8-040

This release includes the following limitations:

- On AMD/RHEL 7.9 systems, the system might panic due to an issue in the IOMMU module. For details, refer to *https://lore.kernel.org/linux-iommu/20191018093830.GA26328@suse.de/t*.
  - Workaround: Disable the IOMMU setting option in BIOS.
- 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 on RHEL 7.9/RHEL8.3 and SLES 15 SP2.
  - Workaround: None
- When performing a driver injection (DUD) install, some Linux distributions (RHEL7.9, RHEL8.2, SLES15 SP2) will hang if a drive in HBA mode has enabled the Drive Write Cache.
  - Workaround: Two workarounds are available for this issue:
    1. Make sure the Drive Write Cache is disabled for any drive in HBA mode.
    2. For RHEL7.9 or 8.2, add `rd.driver.blacklist=smartpqi` to the grub entry along with `inst.dd`.
- Loading the out-of-box driver fails under SLES or RHEL during OS install with secure boot enabled.
  - Workaround:
    1. Install system with inbox driver in secure boot mode.
    2. Enroll the Microchip public key for secure boot.
    3. Install Microchip out-of-box signed driver package.

- The smartpqi.expose_ld_first parameter does not work correctly consistently.
  - Workaround: None

2.3.3.2 Limitations for Windows Driver Build 106.278.0.1043

There are no known limitations for this release.

2.3.3.3 Limitations for FreeBSD Driver Build 4054.0.1005

There are no known limitations for this release.

2.3.3.4 Limitations for Solaris Driver Build 4044.0.1011

This release includes the following limitation:

- If the OS is installed on a logical volume created from an external RAID array, it might not work on the same session.
  - Workaround: Reboot the system.

2.3.3.5 Limitations for VMware Driver Build 4054.0.131

This release includes the following limitation:

- Error messages such as "smartpqi01: pqisrc_show_sense_data_simple:0125: [ERR INFO] BTL: 2:1088:1 op=0x1a path=Raid K:C:Q: 5:20:0" may be seen while running IO. These are debug messages and can be ignored.
  - Workaround: None

2.3.4 Management Software Limitations

2.3.4.1 Limitations for Arcconf/maxView Build B23971

There are no limitations for this release.

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 Microsemi 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.

Note:

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 *Limitations for Firmware Release 1.32 Build 0* on page 16.
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 "Microsemi Adaptec® SmartRAID 3100 Series and SmartHBA 2100 Series Host Bus Adapters Installation and User's Guide (ESC-2171547)" for complete driver installation instructions.

7 The Microchip Web Site

Microchip provides online support via our web site at *http://www.microchip.com/*. This web site is used as a means to make files and information easily available to customers. Accessible by using your favorite Internet browser, the web site contains the following information:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user guides and hardware support documents, latest software releases, and archived software
- **General Technical Support** – Frequently Asked Questions (FAQ), technical support requests, online discussion groups, and Microchip consultant 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

7.1 Customer Change Notification Service

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

To register, access the Microchip web site at *http://www.microchip.com/*. Under “Support”, click on “Customer Change Notification” and follow the registration instructions.

7.2 Customer Support

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Field Application Engineer (FAE)
- Technical Support

Customers should contact their distributor, representative or Field Application Engineer (FAE) for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in the back of this document.

Technical support is available through the web site at: *http://www.microchip.com/support*

7.3 Microchip Devices Code Protection Feature

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

- Microchip products meet the specification contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is one of the most secure families of its kind on the market today, when used in the intended manner and under normal conditions.
- There are dishonest and possibly illegal methods used to breach the code protection feature. All of these methods, to our knowledge, require using the Microchip products in a manner outside the operating specifications contained in Microchip’s Data Sheets. Most likely, the person doing so is engaged in theft of intellectual property.
- Microchip is willing to work with the customer who is concerned about the integrity of their code.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of their code. Code protection does not mean that we are guaranteeing the product as "unbreakable."

Code protection is constantly evolving. We at Microchip are committed to continuously improving the code protection features of our products. Attempts to break Microchip's code protection feature may be

a violation of the Digital Millennium Copyright Act. If such acts allow unauthorized access to your software or other copyrighted work, you may have a right to sue for relief under that Act.

7.4 Legal Notice

Information contained in this publication regarding device applications and the like 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. 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 ITS CONDITION, QUALITY, PERFORMANCE, MERCHANTABILITY OR FITNESS FOR PURPOSE. Microchip disclaims all liability arising from this information and its use. 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.

7.5 Trademarks

The Microchip name and logo, the Microchip logo, AnyRate, AVR, AVR logo, AVR Freaks, BitCloud, chipKIT, chipKIT logo, CryptoMemory, CryptoRF, dsPIC, FlashFlex, flexPWR, Heldo, JukeBlox, KeeLoq, Kleer, LANCheck, LINK MD, maXStylus, maXTouch, MediaLB, megaAVR, MOST, MOST logo, MPLAB, OptoLyzer, PIC, picoPower, PICSTART, PIC32 logo, Prochip Designer, QTouch, SAM-BA, SpyNIC, SST, SST Logo, SuperFlash, tinyAVR, UNI/O, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

ClockWorks, The Embedded Control Solutions Company, EtherSynch, Hyper Speed Control, HyperLight Load, IntelliMOS, mTouch, Precision Edge, and Quiet-Wire 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, BodyCom, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, EtherGREEN, In-Circuit Serial Programming, ICSP, INICnet, Inter-Chip Connectivity, JitterBlocker, KleerNet, KleerNet logo, memBrain, Mindi, MiWi, motorBench, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICkit, PICtail, PowerSmart, PureSilicon, QMatrix, REAL ICE, Ripple Blocker, SAM-ICE, Serial Quad I/O, SMART-I.S., SQI, SuperSwitcher, SuperSwitcher II, Total Endurance, TSHARC, USBCheck, VariSense, ViewSpan, WiperLock, Wireless DNA, 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.

Silicon Storage Technology is a registered trademark 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.

7.6 Quality Management System Certified by DNV

ISO/TS 16949

Microchip received ISO/TS-16949:2009 certification for its worldwide headquarters, design and wafer fabrication facilities in Chandler and Tempe, Arizona; Gresham, Oregon and design centers in California and India. The Company's quality system processes and procedures are for its PIC® MCUs and dsPIC® DSCs, KEELOQ® code hopping devices, Serial EEPROMs, microperipherals, nonvolatile memory and

analog products. In addition, Microchip's quality system for the design and manufacture of development systems is ISO 9001:2000 certified.

7.7 Worldwide Sales and Service

AMERICAS	ASIA/PACIFIC	ASIA/PACIFIC	EUROPE
Corporate Office	**Australia - Sydney**	**India - Bangalore**	**Austria - Wels**
2355 West Chandler Blvd.	Tel: 61-2-9868-6733	Tel: 91-80-3090-4444	Tel: 43-7242-2244-39
Chandler, AZ 85224-6199	**China - Beijing**	**India - New Delhi**	Fax: 43-7242-2244-393
Tel: 480-792-7200	Tel: 86-10-8569-7000	Tel: 91-11-4160-8631	**Denmark - Copenhagen**
Fax: 480-792-7277	**China - Chengdu**	**India - Pune**	Tel: 45-4450-2828
Technical Support:	Tel: 86-28-8665-5511	Tel: 91-20-4121-0141	Fax: 45-4485-2829
http://www.microchip.com/	**China - Chongqing**	**Japan - Osaka**	**Finland - Espoo**
support	Tel: 86-23-8980-9588	Tel: 81-6-6152-7160	Tel: 358-9-4520-820
Web Address:	**China - Dongguan**	**Japan - Tokyo**	**France - Paris**
www.microchip.com	Tel: 86-769-8702-9880	Tel: 81-3-6880- 3770	Tel: 33-1-69-53-63-20
Atlanta	**China - Guangzhou**	**Korea - Daegu**	Fax: 33-1-69-30-90-79
Duluth, GA	Tel: 86-20-8755-8029	Tel: 82-53-744-4301	**Germany - Garching**
Tel: 678-957-9614	**China - Hangzhou**	**Korea - Seoul**	Tel: 49-8931-9700
Fax: 678-957-1455	Tel: 86-571-8792-8115	Tel: 82-2-554-7200	**Germany - Haan**
Austin, TX	**China - Hong Kong SAR**	**Malaysia - Kuala Lumpur**	Tel: 49-2129-3766400
Tel: 512-257-3370	Tel: 852-2943-5100	Tel: 60-3-7651-7906	**Germany - Heilbronn**
Boston	**China - Nanjing**	**Malaysia - Penang**	Tel: 49-7131-67-3636
Westborough, MA	Tel: 86-25-8473-2460	Tel: 60-4-227-8870	**Germany - Karlsruhe**
Tel: 774-760-0087	**China - Qingdao**	**Philippines - Manila**	Tel: 49-721-625370
Fax: 774-760-0088	Tel: 86-532-8502-7355	Tel: 63-2-634-9065	**Germany - Munich**
Chicago	**China - Shanghai**	**Singapore**	Tel: 49-89-627-144-0
Itasca, IL	Tel: 86-21-3326-8000	Tel: 65-6334-8870	Fax: 49-89-627-144-44
Tel: 630-285-0071	**China - Shenyang**	**Taiwan - Hsin Chu**	**Germany - Rosenheim**
Fax: 630-285-0075	Tel: 86-24-2334-2829	Tel: 886-3-577-8366	Tel: 49-8031-354-560
Dallas	**China - Shenzhen**	**Taiwan - Kaohsiung**	**Israel - Ra'anana**
Addison, TX	Tel: 86-755-8864-2200	Tel: 886-7-213-7830	Tel: 972-9-744-7705
Tel: 972-818-7423	**China - Suzhou**	**Taiwan - Taipei**	**Italy - Milan**
Fax: 972-818-2924	Tel: 86-186-6233-1526	Tel: 886-2-2508-8600	Tel: 39-0331-742611
Detroit	**China - Wuhan**	**Thailand - Bangkok**	Fax: 39-0331-466781
Novi, MI	Tel: 86-27-5980-5300	Tel: 66-2-694-1351	**Italy - Padova**
Tel: 248-848-4000	**China - Xian**	**Vietnam - Ho Chi Minh**	Tel: 39-049-7625286
Houston, TX	Tel: 86-29-8833-7252	Tel: 84-28-5448-2100	**Netherlands - Drunen**
Tel: 281-894-5983	**China - Xiamen**		Tel: 31-416-690399
Indianapolis	Tel: 86-592-2388138		Fax: 31-416-690340
Noblesville, IN	**China - Zhuhai**		**Norway - Trondheim**
Tel: 317-773-8323	Tel: 86-756-3210040		Tel: 47-72884388
Fax: 317-773-5453			**Poland - Warsaw**
Tel: 317-536-2380			Tel: 48-22-3325737
Los Angeles			**Romania - Bucharest**
Mission Viejo, CA			Tel: 40-21-407-87-50
Tel: 949-462-9523			**Spain - Madrid**

AMERICAS	ASIA/PACIFIC	ASIA/PACIFIC	EUROPE
Fax: 949-462-9608			Tel: 34-91-708-08-90
Tel: 951-273-7800			Fax: 34-91-708-08-91
Raleigh, NC			**Sweden - Gothenberg**
Tel: 919-844-7510			Tel: 46-31-704-60-40
New York, NY			Sweden - Stockholm
Tel: 631-435-6000			Tel: 46-8-5090-4654
San Jose, CA			**UK - Wokingham**
Tel: 408-735-9110			Tel: 44-118-921-5800
Tel: 408-436-4270			Fax: 44-118-921-5820
Canada - Toronto			
Tel: 905-695-1980			
Fax: 905-695-2078			