

SmartRAID 3200 and SmartHBA 2200 Software/Firmware Release Notes



Table of Contents

1. About This Release.....	3
1.1. Release Identification.....	3
1.2. Files Included in this Release.....	3
2. What's New?.....	5
2.1. Fixes and Enhancements.....	5
2.2. Limitations.....	15
3. Updating the Controller Firmware.....	19
3.1. Updating Controllers to Latest Firmware.....	19
4. Revision History.....	20
Microchip Information.....	21
The Microchip Website.....	21
Product Change Notification Service.....	21
Customer Support.....	21
Microchip Devices Code Protection Feature.....	21
Legal Notice.....	21
Trademarks.....	22
Quality Management System.....	23
Worldwide Sales and Service.....	24

1. About This Release

The release described in this document includes firmware, OS drivers, tools, and host management software for the SmartRAID 3200 and SmartHBA 2200 solutions from Microchip.

1.1 Release Identification

The firmware, software, and driver versions for this release are shown in the following table.

Table 1-1. Release Summary

Solutions release	3.3.2
Package release date	November 2, 2023
Firmware version	03.01.26.036
UEFI/Legacy BIOS	2.10.2/2.10.2
Driver versions	<p>Windows Drivers:</p> <ul style="list-style-type: none"> Windows 2022, 2019, 2016, Windows 11, 10: 1010.84.0.1012 <p>Linux SmartPQI:</p> <ul style="list-style-type: none"> Rocky Linux 9: 2.1.26-030 RHEL 7/8/9: 2.1.26-030 SLES 12/15: 2.1.26-030 Ubuntu 20/22: 2.1.26-030 Oracle Linux 7/8/9: 2.1.26-030 Citrix Xenserver 8: 2.1.26-030 Debian 10/11/12: 2.1.26-030 <p>VMware:</p> <ul style="list-style-type: none"> VMware ESX 7.0/8.0: 4600.0.115 <p>FreeBSD:</p> <ul style="list-style-type: none"> FreeBSD 12/13: 4460.0.1002
ARCCONF/maxView	4.16.0.26273
PLDM	6.30.6.0

1.2 Files Included in this Release

This section details the files included in this release.

Table 1-2. Firmware Files

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

Table 1-3. Firmware Programming Tools

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

Driver Files

Table 1-4. Windows Drivers

OS	Version
Server 2022, 2019, 2016, Windows 11, 10	x64

Table 1-5. Linux Drivers

OS	Version
RHEL 9.3 ¹ , 9.2, 9.1, 9.0 ² , 8.9 ¹ , 8.8, 8.7, 8.6, 7.9	x64
SLES 12 SP5	x64
SLES 15 SP5, SP4, SP3	x64
Ubuntu 20.04.6, 20.04.5, 20.04	x64
Ubuntu 22.04.3, 22.04.2, 22.04	x64
Oracle Linux 7.9 UEK6U3	x64
Oracle Linux 9.2, 9.1, 8.8, 8.7, UEK7U1	x64
Debian 12, 11.7, 10.13	x64
Fedora 38 (inbox)	x64
Citrix XenServer 8.2.1	x64
Rocky Linux 9.1	x64

Notes:

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

Table 1-6. FreeBSD and VMware Drivers

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

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

Description	OS	Executable
ARCCONF Command Line Utility	Windows x64 Linux x64 VMware 7.0 and above XenServer UEFI support	See the arccconf_B#####.zip for the installation executables for the relevant OS.
maxView™ Storage Manager	Windows x64 Linux x64 VMware 7.0 and above XenServer	See the maxview_linux_B#####.zip, maxview_win_B#####.zip, and the maxview_vmware_B#####.zip for the installation executables.
maxView™ vSphere Plugin	VMware 7.0 and above	See the maxview_vmware_B#####.zip for the installation executables.
Boot USB (offline or pre-boot) for ARCCONF and maxView Storage Manager	Linux x64	See the maxview_offline_bootusb_B#####.zip for the .iso file.

2. What's New?

This section shows what's new in this release.

2.1 Fixes and Enhancements

This section shows the fixes and enhancements for this release.

2.1.1 Firmware Fixes

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

2.1.1.1 Fixes and Enhancements for Firmware Release 03.01.26.036

This release includes the following fixes and enhancements:

- Added support to save controller logs in host memory in the event of a system crash.
- Added support for remote managed SED rekey support.
- Added support for Configurable Spindown Spares policy.
- Fixed an issue where data could be lost in MaxCache upon panic shutdown and 1931 POST message seen with controllers that support 8GiB DDR cache.
 - *Root cause:* When creating MaxCache logical drives on a controller with 8GiB DDR cache, firmware was unable to restore some of the dirty cache lines because they were not in a DDR region that was backed up.
 - *Fix:* Correct logic to make sure dirty cache is placed in backed up regions.
 - *Risk:* Low
- Fixed an unexpected rate upshift on an SXP24G expander port upon SATA hot-removal on a different port.
 - *Root cause:* If two ports of the controller are each connected to an SXP24G expander, removing SATA drives from one SXP24G expander resulted in rate upshift for other port of controller.
 - *Fix:* Fixed logic to make sure rate upshift on one port does not affect other port of controller.
 - *Risk:* Medium
- Expected rate downshift does not occur with DCM active expander attached to SATA topology.
 - *Root cause:* There was a problem in logic where the firmware detected a SATA drive, but incorrectly thought the rate had been downshifted.
 - *Fix:* Corrected logic to make sure the rate downshift is applied.
 - *Risk:* Low
- Fixed an issue where taking ownership of Enterprise or Opal SED was failing on boot after panic shutdown.
 - *Root cause:* Changing a master key causes several SED authorities to also change to the new key. The SED flow requires an open session, perform an SED task, and an end session. During this flow, if the controller encounters a panic shutdown, but the SED drives do not encounter a power cycle, then the SED drives are left in the middle of the flow waiting for the session to end. When the controller restarts and attempts to start a new session to validate the datastore on the SED, a start session failure occurs.
 - *Fix:* Error recovery is added to perform a protocol stack reset and retry the start session.
 - *Risk:* Low
- Fixed the following three issues related to expander firmware upgrade:
 - a. OS hang after OS command timeout during expander firmware upgrade.

- b. Firmware lockup due to heartbeat timeout during expander firmware upgrade.
- c. Firmware does not detect all the drives after expander firmware upgrade.
 - *Root cause:* Following are the root cause for the preceding issues:
 - i. If an OS command is routed to target device's internal firmware queue upon expander firmware upgrade, it is not going to be processed by the firmware during the expander firmware upgrade. All host commands should be blocked while expander firmware upgrade is under way, but this is not possible.
 - ii. Expander firmware images can be large which can take almost three minutes to download to the expander. During the download a firmware thread is suspended and does not update the heartbeat counter. The lack of an updated heartbeat counter results in firmware triggering a lockup.
 - iii. After expander is instructed to activate newly downloaded expander firmware, it may takes tens of seconds for the expander to reset. As a result, firmware may not detect all the devices.
 - *Fix:* Following are the fixes for the preceding issues:
 - i. Upon device LUN reset, firmware aborts any command pending in the target device's internal firmware queue to resolve the OS hang.
 - ii. When downloading expander firmware, update the heartbeat counter to resolve the firmware lockup.
 - iii. Delay running device discovery so that all devices can be detected.
 - *Risk:* Low
- Fixed an issue where the controller firmware is not returning the correct queue depth back to the host driver for the RBOD device or storage array controller device. This can cause a degradation in the performance on these devices.
 - *Root cause:* The firmware does not return the queue depth value for these RBOD/storage array controller devices back to the OS device driver.
 - *Fix:* The firmware will check for the request of reporting the queue depth for these devices then it will return the value back to OS device driver.
 - *Risk:* Low
- Fixed an issue that disk drives attached behind an enclosure are shown to have a duplicate bay number.
 - *Root Cause:* The SES additional status element page may contain a valid additional status element descriptor for an empty slot/bay with only a valid device slot number. This causes firmware to assign an incorrect slot number to multiple drives which are in slots after the empty slot.
 - *Fix:* Firmware recognizes a valid additional status element descriptor for an empty slot/bay and skip over it during device discovery.
 - *Risk:* Low
- Fixed an issue that system hung upon device LUN reset due to I/O timeout on RBOD LUNs.
 - *Root cause:* Upon device LUN reset for a multi-LUN device, all requests pending at various firmware queues are flushed and if there is OS partition installed then OS may hang.
 - *Fix:* Do not flush requests pending at various firmware queues upon device LUN reset.
 - *Risk:* Low
- Fixed an issue where a SED drive state will be set to OFS instead of MCHP owned if Master key change process is interrupted due to panic shutdown.
 - *Root cause:* During Master Key change process, several SED tasks are performed. If this process is interrupted due to panic shutdown, it can leave the drive in some intermediate

- state. On the next reboot, when the firmware attempts to open a new session to validate the datastore, a session failure occurs. Power cycling the drive or a reset can clear this condition, or performing a "Protocol Stack Reset" can clear this condition.
- *Fix:* Fixed by performing a "Protocol Stack Reset" if the SED drive is in locked state.
 - *Risk:* Low
- Fixed an issue of possible controller lockup while deleting a logical drive configuration on a controller.
 - *Root cause:* There was a small timing hole while doing a check if cache needed to persist where a null pointer was hit.
 - *Fix:* Make sure to check for null pointer.
 - *Risk:* Low
 - Fixed an issue where a spare drive failed to spin up.
 - *Root cause:* The drive does spin up, but the controller gives up too soon. The controller was not expecting the 2:04:1A (Not Ready: Logical Unit Not Ready, Start Stop Unit Command In Progress) response.
 - *Fix:* Added support to handle the not ready error condition the drive was returning.
 - *Risk:* Low
 - Fixed an issue where a MaxCache logical drive is migrated from one controller to another and a MaxCache pair failed error message could be seen at power-up.
 - *Root cause:* Caching configuration parameters were not being copied over properly to the controller in a new server.
 - *Fix:* Make sure to properly copy caching configuration parameters when destination controller does not have caching enabled.
 - *Risk:* Low
 - Fixed an issue where security enabled SED drives that are in expected qualification failed state are staying failed after deleting the logical drive.
 - *Root cause:* When clearing the configuration, the failure state of the drive was persisting even though the clear configuration was successful.
 - *Fix:* Fixed firmware to clear the error condition if clear configuration is successful.
 - *Risk:* Low
 - Fixed an issue where the rebuild is not starting on a degraded managed SED logical drive that has the first spare drive as foreign SED.
 - *Root cause:* When the firmware tries to activate the available spare drive for the degraded managed SED logical drive, it will choose the first available spare drive without checking the foreign SED status. But during spare activation, the firmware will check foreign SED status and will not activate the spare if it is a foreign SED.
 - *Fix:* Added foreign SED status check in firmware while selecting the spare drive for activation.
 - *Risk:* Low
 - Fixed an issue where the controller is giving the second oldest event as the oldest event.
 - *Root cause:* The controller can store a maximum of 128 events. Once it reaches the maximum limit, a new event will replace the oldest event. While doing so, firmware incremented the event pointer twice in different places.
 - *Fix:* Removed the second increment operation on the event pointer during a rollover.
 - *Risk:* Low
 - Fixed an issue where the rebuild is not starting on the foreign SED after changing the controller master key to the foreign drive key.

- *Root cause:* When the controller-managed SED master key is changed to match the foreign drive key, firmware tries to import the foreign SED and unlock it. During this rekey process, firmware used the old master key to unlock foreign SED and failed to unlock and manage the foreign SED.
 - *Fix:* Added a check in firmware to decide which master key to use for unlocking the foreign SED among the reset key and old master key.
 - *Risk:* Low
- Fixed an issue where a transforming fault-tolerant managed SED logical drive is marked FAILED after hotplugging a foreign SED and rebooting the system.
 - *Root cause:* During bootup, if there is a locked SED in the managed SED logical drive, the logical drive will be marked as SED_LOCKED. Later, firmware will check whether the locked SED is a foreign drive and fail the drive to remove the SED_LOCKED state on the logical drive. But firmware failed to update the controller metadata of the foreign SED. Due to this, the transformation is not able to resume after the reboot, and the logical drive is moved to the FAILED state.
 - *Fix:* Update controller metadata while the firmware is failing the foreign SED to unlock the managed SED logical drive.
 - *Risk:* Low
- Fixed an issue where the controller allowed importing a foreign SED in a degraded logical drive which has a different drive type.
 - *Root cause:* During the bootup, the firmware will check and qualify the newly detected replacement drives to get listed in the logical drive. The firmware will fail the replaced physical drive if it is of a different drive type compared to the existing physical drives in the logical drive. The firmware failed to do this check on a foreign SED and allowed the import of the foreign SED of different drive type to the managed SED logical drive.
 - *Fix:* Added a drive type check for foreign SED in the firmware bootup path.
 - *Risk:* Low
- Fixed an issue where Predictive Spare Rebuild (PSR) is not starting on the managed SED logical drive after replacing the foreign SED spare drive with a controller owned drive.
 - *Root cause:* If PSR is enabled, the firmware will try to activate the spare drive when it observes a predictive failed data drive. But if the spare drive is a foreign SED, it will not be activated. When this foreign SED spare drive is replaced with a controller-owned SED, the firmware does not have any check to activate the spare on predictive failed SED.
 - *Fix:* While the firmware handles the hot insertion of a drive, it will check the need for the spare activation.
 - *Risk:* Low
- Fixed an issue where Predictive Spare Rebuild (PSR) is not starting on the RAID 0 logical drive, if it is the first logical drive in the array followed by multiple fault-tolerant logical drives.
 - *Root cause:* When a spare drive is assigned to the array, where the RAID 0 logical drive is the first logical drive in the array followed by multiple fault-tolerant logical drives, the spare will be assigned to all fault-tolerant logical drives, except the RAID 0 logical drives in the array by the host management tools in a sequence. If one of the data drives in the array is moved to the predictive failed state, the user can enable PSR to activate rebuilds of the predictive failed drive using the spare drive. When PSR is enabled, the firmware will start the PSR immediately on all the logical drives in the array, which have predictive failed drives and assigned spare drives. As the spare drive was not yet assigned to the RAID 0 logical drive by the host management tool, the PSR did not occur. The host management tool will assign a spare drive to RAID 0 logical drives with a delay after enabling the PSR, and the RAID0 logical drive will never get a chance to go through PSR.

- *Fix:* While the firmware handles the configuration update for a logical drive, it will check the need for the spare activation and rebuild.
- *Risk:* Low
- Fixed an issue where the controller may lockup or fail to detect NVME drives after hot add/hot plug or encounter a RSOD boot failure.
 - *Root cause:* The controller firmware processes Configuration Read and Write requests on a higher priority thread, whereas VDM messages are routed to a lower priority thread for processing. The processing of these requests share a common buffer pool. There is a corner case where the host sends a stream of VDM messages, thereby exhausting the shared common buffer pool. As the buffer pool is empty, the firmware cannot process the Configuration Read/Write requests triggered by the host causing a NVME drive hot plug failure, RSOD boot failure, or controller lockup.
 - *Fix:* Combine the processing of VDM messages and Configuration requests into the one high priority thread. This ensures the processing of one type of request will not starve the other of the buffers needed for processing.
 - *Risk:* Low

2.1.2 UEFI/Legacy BIOS Fixes

This section shows the UEFI/Legacy BIOS fixes and enhancements for this release.

2.1.2.1 Fixes and Enhancements for UEFI Build 2.10.2/Legacy BIOS Build 2.10.2

This release includes the following fixes and enhancements:

- Added support to the Controller Information menu to display controller CPLD and SEEPROM versions
- Added support in the logical drive members information menu to show transient drive information for transforming logical drives.
- Added a configuration option to change Spindown Spares Policy that indicates if controller will spindown inactive spares to a state of lower power consumption.
- Added support for additional data and content consistency for Save Support Archive operation, so the output will match with the output of other tools.
- Added HII menu option to perform rekey operation for Remote mode controller managed SED encryption.
- Added new options to the controller firmware update menu to select Active and Backup ROM region for firmware updates as well as to toggle the controller active ROM image.
- Added new HII menu under Disk Utilities to enumerate UBM backplanes along with the option to update backplane firmware.
- Fixed an issue with the unreadable characters for some HII menu help strings.
 - *Root cause:* Incorrect translation of string from English to Chinese language.
 - *Fix:* Updated unicode string file with correct translated string for Chinese language.
 - *Risk:* Low
- Fixed an issue where the System HII browser freezes after entering controller information menu when there is a lockup on the controller.
 - *Root cause:* Command transactions were not getting timed due to incorrect counter usage.
 - *Fix:* Added appropriate timer event to track the command time out and to detect controller firmware readiness.
 - *Risk:* Low
- Fixed an issue where UEFI Self Certification Tests for Block I/O protocol fails.

- *Root cause:* UEFI driver fails to handle the SCSI response sense data that includes unit attention status returned with a non-standard descriptor format.
- *Fix:* Updated error handling for SCSI unit attention response with non-standard sense data descriptor format.
- *Risk:* Low

2.1.3 Driver Fixes

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

2.1.3.1 Linux Driver Fixes

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

2.1.3.1.1 Fixes and Enhancements for Linux Driver Build 2.1.26-030

This release includes the following fixes and enhancements.

- Fixed an OS crash issue that happens while creating/deleting a logical drive or adding/removing physical drives.
 - *Root cause:* The driver is rescanning a device which does not exist. There was a problem where a device rescan operation is failing because the device pointer being used is not valid. This results in a NULL pointer de-reference issue and causes an OS crash.
 - *Fix:* Multiple conditions will be evaluated before notifying the OS to do a rescan. Driver will skip re-scanning the device if any one of the following conditions are met:
 - Device was not added to the OS scsi mid-layer yet or the device was removed.
 - Devices which are marked for removal or in the process of removal.
 - *Risk:* Low
- Fixed an issue to eliminate race condition while rescanning a logical drive. Under very rare conditions, after a logical drive size expansion the OS still sees the size of the original logical drive.
 - *Root cause:* The rescan flag in the driver is used to signal the need for a logical drive rescan. A race condition occurs in the driver which leads to one thread overwriting the flag inadvertently. As a result, the driver is not notifying the OS scsi mid-layer to rescan the logical drive.
 - *Fix:* Acquire spinlock and set the rescan flag.
 - *Risk:* Low

2.1.3.2 Windows Driver Fixes

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

2.1.3.2.1 Fixes and Enhancements for Windows Driver Build 1010.84.0.1012

This release includes the following fixes and enhancements.

- Added support for the DMA V3 kernel API, which simplifies the management of scatter/gather lists and reduces the need for driver intervention during complex DMA transfers.
- Added support for the driver to use StorPortMaskMsixInterrupt() API to enable and disable interrupts when the driver is running on Server 2022 (Version 21H1) and later. The driver will continue to use the legacy method on older versions of the OS where the API is not supported.
- Fixed the I/O errors that were observed inconsistently for MPIO enabled physical drives. The I/O errors reported for multipath during cable plug/unplug.
 - *Root cause:* The MPIO driver detects a device error instead of path failure when Read Capacity, Test Unit Ready, or Inquiry command failed with "SCSI status = Check Condition with Sense Key = Illegal Request (KCQ=5:26:00)" and "SRB status = SRB_STATUS_ERROR".

- *Fix:* The SmartPQI driver returns "SRB status = SRB_STATUS_NO_DEVICE" for "Sense Key = Illegal Request (KCQ=5:26:00)" to indicate that one of the paths has failed.
- *Risk:* Low

2.1.3.3 FreeBSD Driver Fixes

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

2.1.3.3.1 Fixes and Enhancements for FreeBSD Driver Build 4460.0.1002

This release includes the following fix:

- Fixed an issue where under certain I/O conditions a program doing large block disk reads can cause a controller to crash.
 - *Root cause:* The SCSI read request and destination address in the DMA descriptor is incorrect, causing the DMA engine in the controller to assert.
 - *Fix:* Change the alignment for creating `bus_dma_tags` in the driver from `PAGE_SIZE` (4k) to 1, which allows the controller to manage its own address range for DMA transactions.
 - *Risk:* Medium

2.1.3.4 VMware Driver Fixes

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

2.1.3.4.1 Fixes and Enhancements for VMware Driver Build 4600.0.115

There are no known fixes and enhancements for this release.

2.1.4 Management Software Fixes

This section shows the management software fixes and enhancements for this release.

2.1.4.1 maxView Storage Manager/ARCCONF Fixes

This section shows the maxView Storage Manager/ARCCONF fixes and enhancements for this release.

2.1.4.1.1 Fixes and Enhancements for maxView Storage Manager/ARCCONF Build 26273

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

- Microchip strongly recommends that maxView users update to the latest version of the tools to avoid a security vulnerability that has since been resolved.
- Added support in ARCCONF and maxView to configure the "Spindown spares policy".
- Added support to install and run ARCCONF and redfish server in the secureboot ESXi environment.
- Added support in ARCCONF and maxView to prevent the firmware rollback when the given controller firmware update contains a hardware security update and the controller write cache is enabled.
- Enhanced the UEFI ARCCONF `savesupportarchive` command to capture the support logs in the same format as host ARCCONF.
- Fixed an issue where ARCCONF create command help was missing the SSDIOBYPASS option.
 - *Root cause:* SSDIOBYPASS option was not added in ARCCONF help when the user executes create command.
 - *Fix:* Implemented changes to add SSDIOBYPASS option in ARCCONF help when the user executes create command.
 - *Risk:* Low
- Fixed an issue where invalid drive temperature was displayed in maxView for dedicated spare drive.
 - *Root cause:* The temperature data was collected using a signed character variable resulting the variable to hold incorrect data for high temperature values.

- *Fix:* Replaced the signed character variable with unsigned character variable to collect temperature and display invalid temperatures as “Not available” from ARCCONF/maxView.
 - *Risk:* Low
- Fixed an issue where `arccconf SLOTCONFIG` command was not displaying the UBM backplane attached drives.
 - *Root cause:* Slot mapping for the drive connected to the UBM backplane was missing in ARCCONF. This resulted in not displaying the UBM backplane attached drives.
 - *Fix:* Implemented changes to list the slots of the UBM backplane for the `arccconf SLOTCONFIG` command.
 - *Risk:* Low
- Fixed an issue where `savesupportarchive` captures empty crash dump file when the crash dump buffer size was zero from firmware.
 - *Root cause:* An empty crash dump file was created while executing `savesupportarchive` when crash dump buffer size was zero from firmware.
 - *Fix:* Implemented changes only to generate crash dump file when available in firmware using `savesupportarchive` command.
 - *Risk:* Low
- Fixed an issue where usage remaining and estimated life remaining properties for NVMe devices were not displayed in maxView.
 - *Root cause:* Drive interface check of NVMe was not added for usage remaining and estimated life remaining properties. This resulted in not displaying these properties of NVMe devices in maxView.
 - *Fix:* Implemented changes to add NVMe interface check for usage remaining and estimated life remaining properties of NVMe devices which were not displayed in maxView.
 - *Risk:* Low
- Fixed an issue where maxView was creating logical device with reduced size when creating maxCache with 256 KB cache line size.
 - *Root cause:* The size was calculated with default cache line size (64 KB) instead of user specified cache line size (256 KB) which resulted in reduced size of logical device.
 - *Fix:* Implemented the changes to calculate size with user specified cache line size (256 KB).
 - *Risk:* Low
- Fixed an issue where UBM backplanes were enumerated incorrectly in ARCCONF and maxView.
 - *Root cause:* ARCCONF and maxView were displaying UBM controllers as UBM backplane devices which resulted in incorrect enumeration of UBM backplanes.
 - *Fix:* Added changes in the ARCCONF to display one UBM backplane with multiple UBM controllers under it. Added the changes in the maxView to display one UBM backplane in the Enterprise Tree View and display only the required information along with UBM controller ID in the Backplane Summary tab.
 - *Risk:* Low
- Fixed an issue where maxView Desktop application was displaying the warning messages in the startup dialog.
 - *Root cause:* When maxView Desktop application is launched, the warning messages which were part of initialization process were displayed.
 - *Fix:* Suppressed the warning messages which were part of initialization process and added a similar information for all the Operating systems as “*maxView Storage Manager is initializing and will launch once initialization is complete*”.

- *Risk:* Low
- Fixed an issue where maxView was sending email notification twice per event.
 - *Root cause:* When email is configured with port 587, the SMTP server sent the email and returned the error as 'AUTH LOGIN failed (535 Authentication failed. Restarting authentication process.)' and another email was sent through fallback mechanisms because of SMTP error 535.
 - *Fix:* Blocked sending an email through fallback mechanism when SMTP server returns the error as 'AUTH LOGIN failed (535 Authentication failed. Restarting authentication process.)'.
 - *Risk:* Low
- Fixed an issue in maxView firmware upgrade wizard where maxView was displaying two options 'Flash Active ROM' and 'Flash Active and Backup ROM' for the controller that doesn't support flashing active ROM.
 - *Root cause:* maxView displays two options 'Flash Active ROM' and 'Flash Active and Backup ROM' for controllers in firmware upgrade wizard where the controller supports flashing only backup ROM.
 - *Fix:* The maxView firmware upgrade wizard displays only "Flash on Backup ROM" option for those controllers.
 - *Risk:* Low
- Fixed an issue in maxView where maxCache Statistics tab was missing.
 - *Root cause:* The Advanced Statistics code clean-up caused the maxCache statistics tab not to be displayed in the maxView.
 - *Fix:* Added the logic to display the cache statistics tab for maxCache logical device and update the data.
 - *Risk:* Low
- Fixed an issue where the RAID 50/60 legs (parity group) migration operation was displaying success message, but the legs were not migrated to the valid user specified value.
 - *Root cause:* The user specified legs (parity group) value was not passed as an argument to the migrate logical device SDK API which resulted in the operation returning success, but the parity groups were not migrated.
 - *Fix:* Added the changes to input the user specified legs count as an argument to the migrate logical device SDK API for migrating the RAID 50/60 logs.
 - *Risk:* Medium

2.1.4.2 PLDM Fixes

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

2.1.4.2.1 Fixes and Enhancements for PLDM Release 6.30.6.0

This release includes the following fixes and enhancements:

- Added support for the RDE ACTION #Storage.ResetToDefaults to clear event logs and crash dumps stored on the controller in addition to its previous functionality. Logs and crash dumps will be deleted regardless of the requested ResetType.
- Added support for the PLDM Type 6 long-running task for certain RDE operations. When a new RDE operation request is received that is anticipated to exceed the 6 second timeout limit, the PLDM Type 6 state machine will transition to the TASK_RUNNING state. When in this state, the BMC can send the RDEOperationStatus command to query for task completion or failure. Additionally, a RedfishTaskExecuted event will be sent to any event listeners when the task is completed. Long-running tasks will only be supported if the BMC negotiates for Task support using the NegotiateRedfishParameters command. The following RDE operations will be executed through the long-running task:

- RDE DELETE for a Volume resource which is not the last created on a given array.
- RDE DELETE for the last remaining Volume on an array which has SED encryption enabled.
- RDE ACTION for `#Storage.ResetToDefaults` when `ResetType = ResetAll` when encrypted Volumes are present.
- RDE ACTION for `#Storage.ResetToDefaults` when a crash dump is present on the controller.
- The `AutoVolumeCreate` property is now published in RDE READ responses for the Storage resource. This property will have a value of "NonRAID" on controllers which support the HBA Volume feature. Additionally, the schema version of the Storage resource was updated to v1.15.0 to incorporate this new property in the dictionary.
- The following properties are now published in RDE READ responses for the StorageController resource:
 - `Links.Batteries`
 - `Links.Batteries@odata.count`
- Added support for the Battery Redfish resource to publish properties describing a supercap backup power source connected to the controller. The addition of this resource includes parent resources `BatteryCollection`, `PowerSubsystem`, and `Chassis`. The following schema versions are supported for these resources:
 - `Chassis` - v1.22.0
 - `PowerSubsystem` - v1.1.0
 - `BatteryCollection` - unversioned
 - `Battery` - v1.2.0
- Added support for `Storage.Status.HealthRollup` and `StorageController.Health` to publish a value of "Critical" in RDE READ responses when dirty cache flush failures occur during POST.
- Fixed an issue where `WriteCacheProtected` alert is sometimes not generated after clearing a cache temporarily disabled condition. After experiencing a condition that results in a `WriteCacheTemporarilyDegraded` event being generated, no corresponding `WriteCacheProtected` event is sent after rectifying the condition.
 - *Root cause:* A previous fix to suppress `WriteCacheProtected` events at boot time inadvertently prevented the event from ever being sent.
 - *Fix:* Revised the condition for generating `WriteCacheProtected` events to more narrowly define the specific sequence of conditions that denote a good cache status just after boot time.
 - *Risk:* Low
- Fixed an issue where `StorageController.CacheSummary.Status.Health` has "OK" value instead of "Warning" when the cache is in a temporarily degraded state.
 - *Root cause:* The logic to determine the cache health was incorrectly trying to consider the state of a backup power source in cases where such backup power sources are not required.
 - *Fix:* Corrected the logic related to the check of the backup power source state.
 - *Risk:* Low
- Fixed an issue where the response to a `GetDownstreamFirmwareParameters` command did not return "DC power cycle" as a supported component activation method for drives or expander SEPs.
 - *Root cause:* Only the "AC power cycle" and "System reboot" `ComponentActivationMethods` bits were tagged as valid activation methods for drives and expander SEPs.
 - *Fix:* Modified the `ComponentActivationMethods` field for drives and expander SEPs to include setting the "DC power cycle" bit as a supported activation method.

- Risk: Low

2.2 Limitations

This section shows the limitations for this release.

2.2.1 General Limitations

This release includes the following general limitations.

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

2.2.2 Firmware Limitations

This section shows the firmware limitations for this release.

2.2.2.1 Limitations for Firmware Release 03.01.26.036

This release includes the following firmware limitations:

- Power cycle to the enclosure may be needed if connected server goes through abnormal shutdown under following condition: SED operation on OPAL drives like taking ownership, reverting the ownership, or changing the master key where firmware internally performs open session, performs SED management, and ends session gets interrupted due to abnormal shutdown on the server. This condition causes firmware to restart on reboot while drives are left off in the middle of performing SED task and so drives needs to be power cycled also.
 - *Workaround:*
 - Allow the change master key operation to complete before shutting down the server.
 - If SEDs are in an external enclosure, power cycle the external enclosure and SEDs before powering up the server with the controller.

2.2.3 UEFI/Legacy BIOS Limitations

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

2.2.3.1 Limitations for UEFI Build 2.10.2/Legacy BIOS Build 2.10.2

There are no known limitations for this release.

2.2.4 Driver Limitations

This section shows the driver limitations for this release.

2.2.4.1 Linux Driver Limitations

This section shows the Linux driver limitations for this release.

2.2.4.1.1 Limitations for Linux Driver Build 2.1.26-030

This release includes the following limitations:

- On some distributions (RHEL7.9, RHEL8.2, RHEL8.3, SLES15SP2, SLES15SP3, OpenEuler 20.03LTS, and 22.03LTS including SP releases), the driver injection (DUD) install will hang if an attached drive (either HBA mode or Logical Volume) has Write Cache enabled.
 - *Workaround:* There are two workarounds for this issue:
 - Ensure that the Write Cache is disabled for any attached drive.

- For RHEL7.9/8.2/8.3 and OpenEuler 20.03LTS, 22.03LTS, add `rd.driver.blacklist=smartpqi` to the grub entry along with `inst.dd`.
- Unable to do a driver injection (DUD) install on RHEL 8.7 when NVMe drives are attached to the system. This is a multipath issue with the OS install process.
 - *Workaround:* Edit grub to include the boot argument "nompath". Replace "inst.dd" with "nompath inst.dd" for DUD install.
- When doing a driver injection (DUD) install where the OS ISO is mounted as virtual media on BMC based servers (non-ILO). The installer will hang after driver injection. It is reported on RHEL 8.5, 8.6, 9.0, and 9.1.
 - *Workaround:*
 - Load the OS from USB device instead of virtual media.
 - Load the OS from virtual media but initiate ISO verification (media test) during the installation followed by ESC to cancel the media test.
 - Edit grub to include the boot argument "nompath". Replace "inst.dd" with "nompath inst.dd" for DUD install.
- Oracle 9 UEK 7 kernel causes SmartPQI rpm dependency failures. This is an issue with how the kernel package was created by Oracle. Correct UEK7 kernel for Oracle 9, which is expected in the mid-October UEK7 release, version number is still pending.

Note: This does not affect Oracle 8 UEK 7.

 - *Workaround:* Install the rpm using "--nodeps" when dependency failures occur.
 - Update:
 - For SmartPQI driver versions > 2.1.20-020 and UEK7 kernels >= 5.15.0-3.60.2.el9uek.x86_64, the SmartPQI rpm will install normally.
 - For UEK7 kernels < 5.15.0-3.60.2.el9uek.x86_64, the SmartPQI rpm needs to be installed using the "--nodeps".
- On AMD systems, the system might crash or hang due to a bug in the IOMMU module. For details, see lore.kernel.org/linux-iommu/20191018093830.GA26328@suse.de/t/.
 - *Workaround:* Disable the IOMMU setting option in BIOS.
- Depending on hardware configurations, the SmartPQI `expose_ld_first` parameter may not always work consistently.
 - *Workaround:* None
- When multiple controllers are in a system, `udev(systemd)` can timeout during `kdump/kexec` resulting in an incomplete `kdump` operation. The usual indication of the timeout is the console log entry: "scsi_hostX: error handler thread failed to spawn, error = -4".
 - *Workaround:* The workaround for this issue involves extending the `udev(systemd)` timeout during a `kdump` operation. The steps to increase the timeout for `udev(systemd)` are as follows:
 - i. `vi /etc/sysconfig/kdump`
 - ii. Add `udev.event-timeout=300` to `KDUMP_COMMANDLINE_APPEND`
 - iii. `systemctl restart kdump`
 - iv. `systemctl status kdump`
- On some distributions (including RHEL 9.0/Oracle Linux 9.0), you are unable to inject the OOB driver (DUD) during install when a multi-actuator drive is attached.
 - *Workaround:* Install using the inbox driver, complete OS installation, then install the OOB driver.

2.2.4.2 Windows Driver Limitations

This section shows the Windows driver limitations for this release.

2.2.4.2.1 Limitations for Windows Driver Build 1010.84.0.1012

This release includes the following limitation:

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

2.2.4.3 FreeBSD Driver Limitations

This section shows FreeBSD driver limitations for this release.

2.2.4.3.1 Limitations for FreeBSD Driver Build 4460.0.1002

This release contains the following limitations:

- Customized kernels built with the INVARIANTS flag are not currently supported.

2.2.4.4 VMware Driver Limitations

This section shows VMware driver limitations for this release.

2.2.4.4.1 Limitations for VMware Driver Build 4600.0.115

This release includes the following limitations:

- If the controller SED Encryption feature is “On” and locked, Datastores created from secured logical drives on the controller are not automatically mounted even after unlocking the controller, they are not visible through the ESXi hypervisor client.
 - *Workaround:* Use the command `vmkfstool -V` or ESXCLI storage filesystem rescan. Alternatively, use the Rescan option from the Devices tab in the Hypervisor's Storage section. Any of these options solve the issue by forcing a rescan, causing the datastore to mount.
- A controller lockup may occur when using VMDirectPath on a single-processor AMD system. These lockups have been seen with VMs running Linux and Windows. No known workaround at the present time. If a lockup of a passed-through controller occurs, a reboot of the ESXi server may be required to clear the lockup condition and restore the virtual machine to working condition.
- Hot-removal and re-insertion of the same drive may not work. Currently being seen with specific model (KCD6XVUL800G) of Kioxia NVMe drive.

- *Workaround:* Doing a manual unclaim/rescan may help in most of the cases.

2.2.5 Management Software Limitations

This section shows management software limitations for this release.

2.2.5.1 maxView Storage Manager/ARCCONF Limitations

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

2.2.5.1.1 Limitations for maxView Storage Manager/ARCCONF Build 26273

This release includes the following limitations:

- Windows SNMP service is not working after installing maxView.
 - *Description:* In some versions of Windows operating system when Adaptec SNMP subagent is installed, SNMP service does not respond due to the registry configuration which blocks TCP IN/OUT bound traffic to the SNMP subagents even when the firewall is disabled.
 - *Workaround:* Follow the below steps to enable the TCP IN/OUT bound traffic to the SNMP subagents:
 - i. Log on to the system as Administrator and open Registry by issuing regedit in the command prompt.
 - ii. Navigate to [HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\RestrictedServices\Static\System].
 - iii. Find the "Name" string starts with "SNMP-3" and "SNMP-4".
 - iv. Change the "Action=Block" to "Action=Allow" of those entries.
 - v. Restart the "Windows Firewall" service.
 - vi. Restart the "SNMP Service" service.

2.2.5.2 PLDM Limitations

This section shows the PLDM limitations for this release.

2.2.5.2.1 Limitations for PLDM Release 6.30.6.0

There are no known limitations for this release.

3. Updating the Controller Firmware

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

3.1 Updating Controllers to Latest Firmware

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

3.1.1 Upgrading to 3.0X.XX.XXX Firmware

1. For controllers running 3.01.02.042 or higher firmware, flash with 3.0X.XX.XXX version of firmware "SmartFWx200.bin" provided in this package using maxview or ARCCONF utility.
2. Power cycle the server.

4. Revision History

Table 4-1. Revision History

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

Microchip Information

The Microchip Website

Microchip provides online support via our website at www.microchip.com/. This website is used to make files and information easily available to customers. Some of the content available includes:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user’s guides and hardware support documents, latest software releases and archived software
- **General Technical Support** – Frequently Asked Questions (FAQs), technical support requests, online discussion groups, Microchip design partner program member listing
- **Business of Microchip** – Product selector and ordering guides, latest Microchip press releases, listing of seminars and events, listings of Microchip sales offices, distributors and factory representatives

Product Change Notification Service

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

To register, go to www.microchip.com/pcn and follow the registration instructions.

Customer Support

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Embedded Solutions Engineer (ESE)
- Technical Support

Customers should contact their distributor, representative or ESE for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in this document.

Technical support is available through the website at: www.microchip.com/support

Microchip Devices Code Protection Feature

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

- Microchip products meet the specifications contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is secure when used in the intended manner, within operating specifications, and under normal conditions.
- Microchip values and aggressively protects its intellectual property rights. Attempts to breach the code protection features of Microchip product is strictly prohibited and may violate the Digital Millennium Copyright Act.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of its code. Code protection does not mean that we are guaranteeing the product is “unbreakable”. Code protection is constantly evolving. Microchip is committed to continuously improving the code protection features of our products.

Legal Notice

This publication and the information herein may be used only with Microchip products, including to design, test, and integrate Microchip products with your application. Use of this information in any other manner violates these terms. Information regarding device applications is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure

that your application meets with your specifications. Contact your local Microchip sales office for additional support or, obtain additional support at www.microchip.com/en-us/support/design-help/client-support-services.

THIS INFORMATION IS PROVIDED BY MICROCHIP "AS IS". MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE, OR WARRANTIES RELATED TO ITS CONDITION, QUALITY, OR PERFORMANCE.

IN NO EVENT WILL MICROCHIP BE LIABLE FOR ANY INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, OR CONSEQUENTIAL LOSS, DAMAGE, COST, OR EXPENSE OF ANY KIND WHATSOEVER RELATED TO THE INFORMATION OR ITS USE, HOWEVER CAUSED, EVEN IF MICROCHIP HAS BEEN ADVISED OF THE POSSIBILITY OR THE DAMAGES ARE FORESEEABLE. TO THE FULLEST EXTENT ALLOWED BY LAW, MICROCHIP'S TOTAL LIABILITY ON ALL CLAIMS IN ANY WAY RELATED TO THE INFORMATION OR ITS USE WILL NOT EXCEED THE AMOUNT OF FEES, IF ANY, THAT YOU HAVE PAID DIRECTLY TO MICROCHIP FOR THE INFORMATION.

Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

Trademarks

The Microchip name and logo, the Microchip logo, Adaptec, AVR, AVR logo, AVR Freaks, BesTime, BitCloud, CryptoMemory, CryptoRF, dsPIC, flexPWR, HELDO, IGLOO, JukeBlox, KeeLoq, Kleer, LANCheck, LinkMD, maXStylus, maXTouch, MediaLB, megaAVR, Microsemi, Microsemi logo, MOST, MOST logo, MPLAB, OptoLyzer, PIC, picoPower, PICSTART, PIC32 logo, PolarFire, Prochip Designer, QTouch, SAM-BA, SenGenuity, SpyNIC, SST, SST Logo, SuperFlash, Symmetricom, SyncServer, Tachyon, TimeSource, tinyAVR, UNI/O, Vectron, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

AgileSwitch, APT, ClockWorks, The Embedded Control Solutions Company, EtherSynch, Flashtec, Hyper Speed Control, HyperLight Load, Libero, motorBench, mTouch, Powermite 3, Precision Edge, ProASIC, ProASIC Plus, ProASIC Plus logo, Quiet-Wire, SmartFusion, SyncWorld, Temux, TimeCesium, TimeHub, TimePictra, TimeProvider, TrueTime, and ZL are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, Augmented Switching, BlueSky, BodyCom, Clockstudio, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, Espresso T1S, EtherGREEN, GridTime, IdealBridge, In-Circuit Serial Programming, ICSP, INICnet, Intelligent Paralleling, IntelliMOS, Inter-Chip Connectivity, JitterBlocker, Knob-on-Display, KoD, maxCrypto, maxView, memBrain, Mindi, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICkit, PICtail, PowerSmart, PureSilicon, QMatrix, REAL ICE, Ripple Blocker, RTAX, RTG4, SAM-ICE, Serial Quad I/O, simpleMAP, SimpliPHY, SmartBuffer, SmartHLS, SMART-I.S., storClad, SQL, SuperSwitcher, SuperSwitcher II, Switchtec, SynchroPHY, Total Endurance, Trusted Time, TSHARC, USBCheck, VariSense, VectorBlox, VeriPHY, ViewSpan, WiperLock, XpressConnect, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

The Adaptec logo, Frequency on Demand, Silicon Storage Technology, and Symmcom are registered trademarks of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.

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

ISBN: 978-1-6683-3415-7

Quality Management System

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

Worldwide Sales and Service

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