



HBA 1200 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?.....	6
2.1. Features.....	6
2.2. Fixes and Enhancements.....	7
2.3. Limitations.....	12
3. Updating the Controller Firmware.....	15
3.1. Updating Controllers to Latest Firmware.....	15
4. Revision History.....	16
The Microchip Website.....	17
Product Change Notification Service.....	17
Customer Support.....	17
Microchip Devices Code Protection Feature.....	17
Legal Notice.....	18
Trademarks.....	18
Quality Management System.....	19
Worldwide Sales and Service.....	20

1. About This Release

The development release described in this document includes firmware, OS drivers, tools, and host management software for the HBA 1200 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	3.1.6
Package release date	November 18, 2021
Firmware version	3.01.07.046
UEFI/Legacy BIOS	1.4.4.1/1.4.4.2
Driver versions	<p>Windows Drivers:</p> <ul style="list-style-type: none"> Windows 2022, 2019, 2016, Windows 10: 1010.12.0.1007 <p>Linux SmartPQI:</p> <ul style="list-style-type: none"> RHEL 7/8: 2.1.14-035 SLES 12/15: 2.1.14-035 Ubuntu 18/20/21: 2.1.14-035 Oracle Linux 7/8: 2.1.14-035 Citrix Xenserver 8: 2.1.14-035 Debian 9/10: 2.1.14-035 CentOS 7/8: 2.1.14-035 <p>VMware:</p> <ul style="list-style-type: none"> VMware ESX 6/7: 4230.0.103 <p>FreeBSD/Solaris:</p> <ul style="list-style-type: none"> FreeBSD 11/12/13: 4170.0.1014 Solaris: 11: 11.4120.0.1005
ARCCONF/maxView	B24700

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 10	x64

Table 1-5. Linux Drivers

OS	Version
RHEL 8.4, 8.3, 8.2, 8.1, 7.9, 7.8, 7.7	x64
CentOS 8.4, 8.3, 8.2 ,8.1 ,8.0 ,7.9 ,7.8 ,7.7	x64
SLES 12 SP5, SP4	x64
SLES 15 SP3, SP2, SP1	x64
Ubuntu 20.04.2, 20.04.1, 20.04, 18.04.5, 18.04.4	x64
Ubuntu 21.04	x64
Oracle Linux 8.3, 8.2, 7.9, 7.8, UEK6U1 (5.4.17-2036)	x64
Oracle Linux 8.4 UEK6 (5.4.17-2102)	x64
Oracle Linux 8.2 UEK R6	x64
Debian 10.10, 9.13	x64
Fedora 34 (inbox)	x64
XenServer 8.2	x64

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

OS	Version
ESX6.5U3/U2	x64
ESX6.7U3/U2	x64
ESX7.0U3/U2	x64
FreeBSD 13, 12.2, 11.4	x64
Solaris 11.4	x64

Host Management Software

Table 1-7. maxView and ARCCONF Utilities

Description	OS	Executable
ARCCONF Command Line Utility	Windows x64 Linux x64 VMware 6.5 and above XenServer UEFI support	See the arconf_B#####.zip for the installation executables for the relevant OS.
maxView Storage Manager	Windows x64 Linux x64 VMware 6.5 and above XenServer UEFI support	See the maxview_linux_B#####.zip, maxview_win_B#####.zip, and the maxview_vmware_B#####.zip for the installation executables.
maxView vSphere Plugin	VMware 6.5 and above	See the maxview_vmware_B#####.zip for the installation executables.
Boot USB (offline or pre-boot) for ARCCONF and maxView Storage Manager	Linux x64	See the maxview_offline_bootusb_B#####.zip for the .iso file.

2. What's New?

This section shows what's new in this release.

2.1 Features

The following table lists the features supported for this release.

Table 2-1. Features Summary

Features		Supported in this Release	Future Release
UEFI driver, boot support		X	
Legacy boot support		X	
Dynamic power management		X	
Driver support	Windows	X	
	Linux	X	
	VMware	X	
	FreeBSD	X	
	Solaris	X	
	OS certification	X	
Flash support	ARCCONF utility	X	
maxView tool support		X	
ARCCONF tool support		X	
MCTP BMC management		X	
4Kn support in RAID and HBA		X	
Controller-based encryption (CBE) support ¹		X	
Out-of-band interface selection support of MCTP or PBSI		X	
VPP Backplane support		X	
PBSI support		X	
Configurable Expander SSU settings		X	

Note:

1. Only available for encryption-enabled products.

2.2 Fixes and Enhancements

This section shows the fixes and enhancements for this release.

2.2.1 Firmware Fixes

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

2.2.1.1 Fixes and Enhancements for Firmware Release 3.01.07.046

This release provides the following fixes and enhancements.

- Added workaround to work with AMI MG9100 UBM backplanes.
- Added workaround to rate-limit the attached SAS domain to 12 Gbps when a SAS-4 expander attached at 22.5 Gbps is present with at least one SATA drive attached.
- Increased the size of the host request object pool to account for a reduction in available request objects when the controller queue depth is at its maximum (1024). This should result in slightly better performance at that threshold.
- Improved request coalescing behaviors for non-parity HDD volumes as queue depth is scaled up.
- Added support for an off-board activity LED.
- Additional changes merged to modify various firmware error paths to avoid use of a predictive failure device.
- Added support for persistent log preservation through cold boot.
- Added support for SCSI_UNMAP for NVMe drives.
- Fixed an issue where background thread hangs while running device handle swap test.
 - *Root cause:* With frequent volume state transitions, it was possible for background thread to get stuck waiting for last physical request to wake it up, but all physical requests had been completed.
 - *Fix:* Adjust physical request count properly when a failure is seen reading drive block size
 - *Risk:* Low
- Fixed an issue where NVME drives would be reset internally on all drive firmware updates.
 - *Root cause:* NVME translation layer would reset all drive internally on drive firmware updates. Per the NVME spec, some drives support online activation and for those drives reset is not needed.
 - *Fix:* If drive supports online activation, do not reset the drive on firmware update.
 - *Risk:* Low
- Fixed an issue where both active and inactive image can be corrupted.
 - *Root cause:* Firmware was not checking status of redundant image and hence was allowing corruption of both active and inactive image components.
 - *Fix:* Changes were implemented in firmware to first check the redundant image status to disallow corrupting other image.
 - *Risk:* Low
- Fixed an issue where a controller crash dump is not recorded if the host issues PERST# during that process.
 - *Root cause:* Firmware reset was being invoked immediately on receipt of PERST#. If a crash dump was being generated at this time, then its contents would be lost.
 - *Fix:* Modified the crash handling routines to disable propagation of PERST# until the crash dump process is complete.
 - *Risk:* Low
- Fixed a potential controller lockup during NVMe drive firmware update.
 - *Root cause:* If the API call to "requery" a device for inventory purposes is sent at the same time the RAID metadata is being saved, there is a race condition exposure between the device info being cleared and the metadata update process in which a debug trap was being hit that is checking for valid drive parameters.
 - *Fix:* Instead of a lockup, modified the drive parameter checks to abort the current metadata update and retry later.
 - *Risk:* Low
- Fixed an issue where the controller can fail to discover devices after a cable is hot-added

- *Root cause:* The SFF-8449 specification lists a minimum setup time for the cable of 2 seconds before interrogating the cable to determine its interface type and setup details. The controller firmware was not providing this setup time and in some cases cables would not be ready for access and fail to be discovered.
 - *Fix:* Adjusted the cable insertion handler to provide the appropriate delay for standards compliance.
 - *Risk:* Low.
- Fixed a potential controller lockup when attached to a UBM backplane advertising bay count as zero
 - *Root cause:* A debug trap was encountered when a UBM backplane is attached advertising the bay count as zero that results in firmware attempting a memory allocation of size zero for the bay information
 - *Fix:* Added error handling for this case.
 - *Risk:* Low
- Fixed an issue where running 'abs_print_eye_capture' and 'cross_hair_enable' from ChipLink caused a firmware lockup.
 - *Root cause:* There was an issue in parsing the arguments for these commands that resulted in a firmware exception and lockup.
 - *Fix:* The logic that parses these commands is fixed to handle these commands with arguments.
 - *Risk:* Low
- Fixed an issue where the Fault LED is not turned ON when firmware fails a bad HBA drive connected to an expander during device discovery or hot-plug.
 - *Root cause:* When firmware fails a HBA drive during device discovery or hot-plug, it does not set the “select” bit to 1 in the SES control page to turn on the Fault LED. The issue is that firmware still sees this failed HBA drive is exposed to host so firmware does not control the drive LED and the host must control the drive LED. This leads to the firmware not setting the “select” bit to 1 in the SES control page for the failed HBA drive when firmware is performing the LED update operation.
 - *Fix:* Firmware sets the “select” bit to 1 in SES control page for the failed HBA drive which is not exposed to host and turns ON the drive's Fault LED correctly.
 - *Risk:* Low
- Fixed a problem where physical drive firmware update will not succeed when initiated through Out-Of-Band MCTP host transport
 - *Root Cause:* Firmware was not setting the correct response back to host for the SCSI pass through OUT direction commands.
 - *Fix:* To adjust the response buffer properly for these SCSI pass through commands.
 - *Risk:* Medium
- Fixed a problem where the I/O latency may be more for maxCache configured SATA drive volume, if there are non-remappable UREs/bad blocks.
 - *Root Cause:* When maxCache read request is failed due to bad block sense data, firmware is not logging bad blocks internally with required flags to invoke "short circuit logic" on bad blocks which sets the respective error status without having to send it to the drives. Due to this reason subsequent reads to known bad blocks are sent to the drive and leading to unwanted latency in logical volume read IOs.
 - *Fix:* Update the logical request type to trigger execution of the short circuit bad block logic.
 - *Risk:* Low
- Fixed an issue where failure of a logical drive in dual domain JBOD could also generate a warning about loss of a redundant path
 - *Root Cause:* The warning message is observed when SAS drives within the logical drive in the dual domain configuration are hot removed until the logical drive goes to a FAILED state. The warning message is due to the firmware logic that updates the “Redundant Cabling Flags”. These flags are used by host management tools for displaying the warning message. The flags were not updated when the logical drive is in FAILED state that resulted in an incorrect flag setting that generated the warning message by the host management tool.
 - *Fix:* Firmware will correctly set the "Redundant Cabling Flags" when the logical drive is in the FAILED state to avoid the incorrect warning message by the host management tool.
 - *Risk:* Low
- Fixed an issue where the device 'serial number' field was being overloaded with failure reason information

- *Root cause:* For a device which fails very early in target discovery, the serial number field was being overloaded with more detailed failure reason information to indicate the nature of the very early device fault. The incorrect serial number information was then being displayed in various inventory displays which causes confusion even though the target is known to be failed.
- *Fix:* Added additional device failure reason codes to describe each of these scenario's instead of lumping them all under a single generic reason.
- *Risk:* Low

2.2.2 UEFI/Legacy BIOS Fixes

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

2.2.2.1 Fixes and Enhancements for UEFI Build 1.4.4.1/Legacy BIOS Build 1.4.4.2

This release provides the following fixes and enhancements.

- Updated driver branding from Microsemi to Microchip.
- Added Support in Firmware Management Protocol to support fields Lowest Supported Image Version, Last Attempt Version and Last Attempt Status.
- Fixed an issue where eligible drives are not listing while performing Heal Array on a failed array which requires multiple drives to fix.
 - *Root cause:* Incorrect drive eligibility validation when heal array operation requires multiple drives.
 - *Fix:* Consider multiple drive requirement while performing eligibility validation for heal array operation.
 - *Risk:* Low
- Fixed an issue where the version field of Firmware Management Protocol is not populated with 32 bit version.
 - *Root cause:* Version field of Firmware Management Protocol is not populated with 32 bit version instead it is assigned with truncated long version.
 - *Fix:* Version field of Firmware Management Protocol assigned with 32 bit version format.
 - *Risk:* Low
- Fixed an issue where HII and health messages displayed incorrect translations for Chinese and Japanese strings.
 - *Root cause:* Incorrect translation for few HII options and driver health messages.
 - *Fix:* Corrected language translations for Unicode strings.
 - *Risk:* Low
- Fixed an issue where the driver health error code 0x1945 was not observed when parity initialization had not yet completed while logical drive rebuilding was in progress.
 - *Root cause:* Degraded logical drive case for error code 0x1945 does not consider logical drive state of rebuilding.
 - *Fix:* Add the logical drive with state rebuilding to degraded list for the case 0x1945 error.
 - *Risk:* Low
- Fixed an issue where creating a volume on an NVMe array used a default strip size of 256 KiB instead of 128 KiB.
 - *Root cause:* SA_GetEditableLogicalDriveStripSizeRange had no logic in place to set a default strip size specific to NVMe arrays.
 - *Fix:* Updated the default strip size calculation to be set to either 128 KiB or the max strip size supported by the controller for NVMe arrays or the desired RAID level, whichever is smallest.
 - *Risk:* Low
- Fixed an issue where the system would experience a Bootup Hang with Insyde Legacy BIOS on AMD ROME platform.
 - *Root cause:* The Legacy OpROM was not checking if the Keyboard buffer is empty before reading and writing to ports 60h and 64h.
 - *Fix:* Check if the keyboard buffer is clear before reading and writing to the Ports 60h and 64h.If not, clear the buffer by reading the data from Port 60h.
 - *Risk:* Medium

2.2.3 Driver Fixes

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

2.2.3.1 Windows Driver Fixes

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

2.2.3.1.1 Fixes and Enhancements for Windows Driver Build 1010.12.0.1007

There are no fixes and enhancements for this version.

2.2.3.2 Linux Driver Fixes

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

2.2.3.2.1 Fixes and Enhancements for Linux Driver Build 2.1.14-035

This release provides the following fixes and enhancements.

- Fixed an issue of driver spin down when system transitions to the Suspend (S3) state in certain systems.
 - *Root cause:* In certain system (based on PCI IDs), when the OS transitions the system into the Suspend (S3) state, the flush cache command indicates a system RESTART instead of SUSPEND. This avoids drive spin-down.
 - *Fix:* Avoid drive spin-down when system transitions to the Suspend state.
 - *Risk:* Medium
- Added enable SATA NCQ priority support to `sysfs`. The driver needed device attribute `sas_ncq_prio_enable` for I/O utility to enable SATA NCQ priority support and to recognize I/O priority in SCSI command and pass priority information to controller firmware. This device attribute works only when device has NCQ priority support and the controller firmware can handle I/O with NCQ priority attribute.
- Fixed an issue where logical drive size is not reflecting after expansion. After modifying the logical drive size, `lsblk` command still shows previous size of the logical volume.
 - *Root cause:* When the driver gets any event from firmware, the driver schedules a rescan worker with a delay of 10 seconds. If the array expansion completes too quickly (in a second), the driver does not catch the logical drive expansion due to worker delay. Since the driver doesn't detect logical drive expansion, it does not call rescan device to update the new size of the logical drive to the OS. This causes `lsblk` to report the original size.
 - *Fix:* For every logical device event notification, driver rescans the logical drive.
 - *Risk:* Low
- Fixed an issue where during `kdump` OS is dropping into a shell if the controller is in locked-up state.
 - *Root cause:* Driver issues SIS soft reset to restore the controller to SIS mode when OS boots into `kdump` mode. If the controller is in Locked-up state, the SIS soft reset does not work. Since the controller lockup code has not been cleared, the driver considers firmware is no longer up and running. In this case, the driver returns an error code to OS and `kdump` fails. After `kdump` failure, some OS distributions do not reboot cleanly which leads to the OS dropping into a recovery shell.
 - *Fix:* During `kdump`, driver will reboot the system if the controller is in Locked-up state.
 - *Risk:* Low
- Fixed an issue where the controller spins down drives during a warm boot on Linux.
 - *Root cause:* The Linux SmartPQI driver has a callback function that the OS calls when the system is being shut down or being rebooted. This callback function calls the Flush Cache command. The command has a parameter that allows the driver to indicate to the firmware the reason for the flush cache (shutdown, hibernate, suspend, or restart). The OS callback function does not indicate to the driver whether it is being called for shutdown or warm boot, so the driver indicates to the firmware that the reason for the flush cache is a system shutdown. The firmware always spins down drives in this case.
 - *Fix:* The SmartPQI driver uses a Linux kernel global variable to distinguish between a system shutdown and a warm boot and sets the Flush Cache command parameter accordingly.
 - *Risk:* Low
- Fixed an issue where duplicate device nodes for Ultrium tape drive and medium changer are being created.
 - *Root cause:* The Ultrium tape drive is a multi-LUN SCSI target. It presents a LUN for the tape drive and a second LUN for the medium changer. The controller firmware lists both LUNs in the report logical LUNS

results, so the SmartPQI driver exposes both devices to the OS. Then the OS does its normal device discovery through the SCSI REPORT LUNS command, which causes it to re-discover both devices a second time, resulting in duplicate device nodes.

- *Fix*: When the OS re-discovers the two LUNs for the tape drive and medium changer, the driver recognizes that they have already been reported and blocks the OS from adding them a second time.
- *Risk*: Low

2.2.3.3 VMware Driver Fixes

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

2.2.3.3.1 Fixes and Enhancements for VMware Driver Build 4230.0.103

This release provides the following fixes and enhancements.

- Fixed an issue where possibility of a null device pointer needs to be prevented in one of the functions where it waits for the outstanding commands to get completed.
 - *Root cause*: Device may have been removed.
 - *Fix*: Check for a null device pointer before starting the wait loop.
 - *Risk*: Low
- Fixed an issue where a failed lookup results in an array out of bounds condition.
 - *Root cause*: A device lookup function returns INVALID_ELEM (0xffff) when device is not found, but calling function does not check for error, and unconditionally uses lookup's return as index into the device list.
 - *Fix*: Print message and do not continue device addition or deletion if lookup function returns INVALID_ELEM.
 - *Risk*: Low

2.2.3.4 FreeBSD/Solaris Driver Fixes

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

2.2.3.4.1 Fixes and Enhancements for FreeBSD Driver Build 4170.0.1014

This release provides the following fixes and enhancements.

- Fixed an issue where debug log messages were flooding the kernel logs.
 - *Root cause*: There are a lot of DBG_INFO prints which are logged by SmartPQI and one DBG_ERR print causing log contention which should not be considered an error.
 - *Fix*: Disable the DBG_INFO prints from logging and change DBG_ERR to DBG_INFO for a message not considered an error.
 - *Risk*: Low

2.2.3.4.2 Fixes and Enhancements for Solaris Driver Build 11.4120.0.1005

There are no fixes and enhancements for this version.

2.2.4 Management Software Fixes

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

2.2.4.1 maxView Storage Manager/ARCCONF Fixes

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

2.2.4.1.1 Fixes and Enhancements for maxView Storage Manager/ARCCONF Version 2.0.0 Build 24700

This release provides the following fixes and enhancements.

- Added support for Redfish Server Daemon in ESXi 7.x.
- Added support to configure “unchanged” option for “drive write cache policy” at the controller level.
- Rebranded maxView applications from Microsemi to Microchip.
- Added support to set 128 kB as default stripe size for NVMe logical device creation.
- Fixed an issue where maxView doesn't work after upgrading build 23821 to 24308.
 - *Root cause*: maxView installer was not clearing the older files while upgrading, making maxView unusable.
 - *Fix*: Added changes to older files from installed directory while upgrading maxView.
 - *Risk*: Low

- Fixed an issue where secure erase task progress goes from 98% to 0% in ARCCONF.
 - *Root cause:* ARCCONF was displaying a secure erase task that is completed as still in progress with 0%.
 - *Fix:* Added changes to not display a secure erase task when task is completed.
 - *Risk:* Low
- Fixed an issue where ROMUPDATE command fails to open image in UEFI ARCCONF.
 - *Root cause:* ARCCONF was doing a redundant image file verification which resulted in failure.
 - *Fix:* Added changes for proper image file verification in ROMUPDATE command.
 - *Risk:* Low
- Fixed an issue where maxView does not display new firmware version after updating NVMe device with E+F mode.
 - *Root cause:* Incorrect mapping of E+F mode resulted firmware not flashing immediately.
 - *Fix:* Added changes for map proper E+F mode for NVMe devices to flash and reflect the new firmware image immediately.
 - *Risk:* Low
- Fixed an issue where maxView displays a warning message after selecting a firmware image file for flashing.
 - *Root cause:* Warning message of invalid message is displayed on the window of maxView after selecting the firmware image.
 - *Fix:* Added changes for move the warning message before uploading the firmware image in maxView.
 - *Risk:* Low

2.3 Limitations

This section shows the limitations for this release.

2.3.1 Firmware Limitations

This section shows the firmware limitations for this release.

2.3.1.1 Limitations for Firmware Release 3.01.07.046

This release includes the following limitation:

- This release has a limitation with SAS-4 (SXP24G) Expander-Attached SATA drives when the Controller-Expander operates at 24G. If a SATA device is connected to SXP24G, link between HBA 1200 and SXP24G will not operate at 24G rate. Please refer to SmartROC 3200/SmartIOC 2200 Device Errata for more details on root cause and impact due to this issue.
 - *Workaround:* Firmware will rate-limit the attached SAS domain to 12 Gbps when SAS-4 Expander attached at 22.5 Gbps is present with at least one SATA drive attached in the topology.

2.3.2 UEFI/Legacy BIOS Limitations

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

2.3.2.1 Limitations for UEFI Build 1.4.4.1/Legacy BIOS Build 1.4.4.2

There are no known limitations for this release.

2.3.3 Driver Limitations

This section shows the driver limitations for this release.

2.3.3.1 Windows Driver Limitations

This section shows the Windows driver limitations for this release.

2.3.3.1.1 Limitations for Windows Driver Build 1010.12.0.1007

There are no known limitations for this release.

2.3.3.2 Linux Driver Limitations

This section shows the Linux driver limitations for this release.

2.3.3.2.1 Limitations for Linux Driver Build 2.1.14-035

This release includes the following limitations.

- An issue can occur when doing a driver injection (DUD) install. On some distributions (RHEL7.9, RHEL8.2, RHEL8.3, SLES15SP2, SLES15SP3), the DUD install will hang if a drive in HBA mode has the Drive Write Cache enabled.
 - *Workaround:* There are two workarounds for this issue:
 - i. Make sure the Drive Write Cache is disabled for any drive in HBA mode.
 - ii. For RHEL7.9/8.2/8.3, add `rd.driver.blacklist=smartpqi` to the grub entry along with `inst.dd`.
- Due to a change in the SCSI mid-layer, some Linux distributions may take a long time to come up if the system is rebooted while a hard disk(s) is being sanitized. This has currently been observed with inbox smartpqi drivers on RHEL 7.9/RHEL8.3 and SLES 15SP2.
 - *Workaround:* Do not reboot the system while a hard disk is being sanitized, or update to the smartpqi 2.1.12-055 or later driver release.
- On AMD/RHEL 7.9 systems, the system might panic due to the a bug in the IOMMU module. For details, refer to lore.kernel.org/linux-iommu/20191018093830.GA26328@suse.de/
 - *Workaround:* Disable the IOMMU setting option in BIOS.
- On AMD/UEK6 systems, the system might hang during kdump if IOMMU is enabled.
 - *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
- Hibernating Linux system using `pm-hibernate` command causes system to hang.
 - *Workaround:* None
- When multiple controllers are in a system, `udev(systemd)` can timeout during kdump/kexec resulting in an incomplete kdump operation. The usual indication of the timeout is the console log entry: `"scsi_hostX: error handler thread failed to spawn, error = -4"`.
 - *Workaround:* Extend the `udev(systemd)` timeout during a kdump operation. Use the following to increase the timeout for `udev(systemd)`:

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

2.3.3.3 VMware Driver Limitations

This section shows VMware driver limitations for this release.

2.3.3.3.1 Limitations for VMware Driver Build 4230.0.103

There are no known limitations for this release.

2.3.3.4 FreeBSD/Solaris Driver Limitations

This section shows FreeBSD/Solaris driver limitations for this release.

2.3.3.4.1 Limitations for FreeBSD Driver Build 4170.0.1014

This release includes the following limitations.

- Under heavy I/O with transfer size more than 128k, controller may go offline. This happens in FreeBSD 13.
 - *Workaround:* Reduce the I/O transfer size of the application to less than 128k.

2.3.3.4.2 Limitations for Solaris Driver Build 11.4120.0.1005

There are no known limitations for this release.

2.3.4 Management Software Limitations

This section shows management software limitations for this release.

2.3.4.1 maxView Storage Manager/ARCCONF Limitations

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

2.3.4.1.1 Limitations for maxView Storage Manager/ARCCONF Version 2.0.0 Build 24700

This release includes the following limitations.

- Advanced statistics will no longer be available in maxView/ARCCONF.
- ADU report in support archive will no longer be available in zip format. The relevant logs are captured under Controller_X_Debug_Log.txt
- SSD report in support archive will not be available.
- OS partition information is not available in FreeBSD and Solaris OS in maxView/ARCCONF.
- Remote ARCCONF (CIM client) is not support for ESXi 7.x onwards.
- Due to data type mismatch between maxView and redfish server, eccRecoveredReadErrors and serviceHours properties in the drive error counter tab will not be reflecting the current value.
 - *Workaround:* User needs to use arccconf CLI GETCONFIG command to refer the current value for these error counter properties.
- In ESXi 7.x, maxView GUI may not update the latest configuration automatically when the operations are performed through the ESXi host ARCCONF.
 - *Workaround:* User needs to refresh the configuration using the refresh link provided in the top right corner in the maxView GUI before performing any operations.
- When user tries to access the maxView `main.xhtml` page directly when the previous session was still active, user may end up with a warning page mentioning “XML Parsing Error: no root element found”.
 - *Workaround:* User needs to use the login page to get authenticated and create a new session to access the `main.xhtml` page.
- When the SED drive is in locked state, the hard drive level refresh SED security status operation is not available in the maxView GUI.
 - *Workaround:* Use the controller level refresh SED security status operation or use ARCCONF command to refresh the SED security status.
- In Linux OS, the redfish server may get terminated when Delete Array operation is performed on an array with 64 logical devices.
 - *Workaround:* The user shall restart the redfish server or use ARCCONF CLI for configuration.

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 3.01.07.046 Firmware

1. For controllers running 3.01.02.042 or higher firmware, flash with 3.01.07.046 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
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.

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 devices:

- 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 and under normal conditions.
- There are dishonest and possibly illegal methods being used in attempts to breach the code protection features of the Microchip devices. We believe that these methods require using the Microchip products in a manner outside the operating specifications contained in Microchip's Data Sheets. Attempts to breach these code protection features, most likely, cannot be accomplished without violating Microchip's intellectual property rights.
- Microchip is willing to work with any customer who is concerned about the integrity of its code.
- 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. 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.

Legal Notice

Information contained in this publication is provided for the sole purpose of designing with and using Microchip products. Information 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.

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

IN NO EVENT WILL MICROCHIP BE LIABLE FOR ANY INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL OR CONSEQUENTIAL LOSS, DAMAGE, COST OR EXPENSE OF ANY KIND WHATSOEVER RELATED TO THE INFORMATION OR ITS USE, HOWEVER CAUSED, EVEN IF MICROCHIP HAS BEEN ADVISED OF THE POSSIBILITY OR THE DAMAGES ARE FORESEEABLE. TO THE FULLEST EXTENT ALLOWED BY LAW, MICROCHIP'S TOTAL LIABILITY ON ALL CLAIMS IN ANY WAY RELATED TO THE INFORMATION OR ITS USE WILL NOT EXCEED THE AMOUNT OF FEES, IF ANY, THAT YOU HAVE PAID DIRECTLY TO MICROCHIP FOR THE INFORMATION. Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

Trademarks

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

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

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, Augmented Switching, BlueSky, BodyCom, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, Espresso T1S, EtherGREEN, IdealBridge, In-Circuit Serial Programming, ICSP, INICnet, Intelligent Paralleling, Inter-Chip Connectivity, JitterBlocker, 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, SMART-I.S., storClad, SQL, SuperSwitcher, SuperSwitcher II, Switchtec, SynchroPHY, Total Endurance, TSHARC, USBCheck, VariSense, VectorBlox, VeriPHY, ViewSpan, WiperLock, XpressConnect, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

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

The Adaptec logo, Frequency on Demand, Silicon Storage Technology, 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.

© 2021, Microchip Technology Incorporated, Printed in the U.S.A., All Rights Reserved.

ISBN: 978-1-5224-8477-6

Quality Management System

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

Worldwide Sales and Service

AMERICAS	ASIA/PACIFIC	ASIA/PACIFIC	EUROPE
<p>Corporate Office 2355 West Chandler Blvd. Chandler, AZ 85224-6199 Tel: 480-792-7200 Tel: 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>