

Table of Contents

1. About This Release.....	3
1.1. Release Identification.....	3
1.2. Components and Documents Included in this Release.....	3
1.3. Files Included in this Release.....	3
2. What's New?.....	6
2.1. Features.....	6
2.2. Fixes.....	6
2.3. Limitations.....	14
3. Updating the Controller Firmware.....	19
3.1. Updating the Controller Firmware.....	19
4. Installing the Drivers.....	21
5. Revision History.....	22
Microchip Information.....	23
Trademarks.....	23
Legal Notice.....	23
Microchip Devices Code Protection Feature.....	23

1. About This Release

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

1.1. Release Identification

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

Table 1-1. Release Summary

Solutions Release	2.9.4
Package Release Date	May 10, 2025
Firmware Version	7.60 B0 ¹
UEFI Driver Version	2.18.4
Legacy BIOS	2.18.2
Driver Versions	Windows SmartPQI: <ul style="list-style-type: none"> Windows Server 2019/2022/2025: 1016.18.0.1014 Windows 10/11: 1016.18.0.1014 Linux SmartPQI: <ul style="list-style-type: none"> RHEL 7/8/9/10: 2.1.34-035 SLES 12/15: 2.1.34-035 Ubuntu 20/22/24: 2.1.34-035 Debian 11/12: 2.1.34-035 Oracle Linux 7/8/9: 2.1.34-035 Citrix XenServer 8: 2.1.34-035 BC Linux 7: 2.1.34-035 OpenEuler 22/24: 2.1.34-035 VMware SmartPQI: <ul style="list-style-type: none"> VMware 7.0/8.0/9.0: 4856.0.105 FreeBSD SmartPQI: <ul style="list-style-type: none"> FreeBSD 13/14: 4620.0.1010
arcconf/maxView™	4.26.00.27449
PLDM	6.50.11.0

Note:

- Downgrading to 1.04 B0 or older builds from this release or prior 1.29 releases may cause the board to not boot or have supercap errors due to an incompatibility in SEEPROMs between this release and prior releases. See section “[Updating the Controller Firmware](#)”.

1.2. Components and Documents Included in this Release

Download the firmware, drivers, host management software, and supporting documentation for your HBA1100 controller solution from the Microchip Web site at <https://start.adaptec.com>

1.3. Files Included in this Release

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

Firmware Files

Table 1-2. Firmware Files

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

Table 1-3. Firmware Programming Tools

Tool	Description	Executable
Arcconf romupdate	The command allows to upgrade/downgrade the firmware and BIOS image to the controller.	Refer to Table 1-8
maxView™ firmware upgrade wizard	The firmware upgrade wizard allows to upgrade/downgrade the firmware and BIOS image to one or more controller(s) of same model in the system.	Refer to Table 1-8

Driver Files

Table 1-4. Windows Storport Miniport SmartPQI Drivers

Drivers	Binary	Version
Server 2025, 2022 and 2019	SmartPqi.sys	x64
Windows 10 (version 22H2) and 11 (version 24H2)	SmartPqi.inf	x64
	smartpqi.cat	x64

Table 1-5. Linux SmartPQI Drivers for Arm

Drivers	Version
Red Hat Enterprise Linux 9.5, 8.10	Arm®
SuSE Linux Enterprise Server 12 SP5	Arm
SuSE Linux Enterprise Server 15 SP6, SP5	Arm
Ubuntu 24.04.2, 22.04.5, 20.04.5	Arm
BC Linux 7.7	Arm
OpenEuler 24.03 SP1 LTS, 22.03 SP4 LTS	Arm

Table 1-6. Linux SmartPQI Drivers for Intel/AMD x64

Drivers	Version
Red Hat Enterprise Linux 9.5, 9.4, 9.0, 8.10, 8.9, 8.0, 7.9	x86_64
SuSE Linux Enterprise Server 12, SP5	x86_64
SuSE Linux Enterprise Server 15 SP6, SP5	x86_64
Oracle Linux 7.9 UEK6U3	x86_64
Oracle Linux 9.5, 9.4, 8.10, 8.9, UEK7U3	x86_64
Ubuntu 24.04.2, 24.04.1, 24.04, 22.04.5, 22.04.4, 22.04	x86_64
Ubuntu 20.04	x86_64

Table 1-6. Linux SmartPQI Drivers for Intel/AMD x64 (continued)

Drivers	Version
Debian 12.8, 12.6, 11.11, 11.10	x86_64
Citrix xenServer 8.2.1	x86_64
Fedora 41 (inbox only)	x86_64
OpenEuler 24.03 SP1 LTS	x86_64
OpenEuler 22.03 SP4 LTS	x86_64
SLE-Micro 6.1, 6.0 (Inbox only)	x86_64

Table 1-7. FreeBSD and VMware SmartPQI Drivers

Drivers	Version
FreeBSD 14.2, 13.4	x64
VMware 8.0 U3/U2, 7.0 U3/U2	x64

Host Management Software

Table 1-8. Host Management Utilities

Description	OS	Executable
ARCCONF Command Line Utility	Windows® x64 Linux® x64 VMware 7.0 and above XenServer FreeBSD x64 Linux ARM	See the Arccnf download package for the OS-applicable installation executable.
ARCCONF for UEFI	—	Included as part of the firmware downloadable image.
maxView™ Storage Manager	Windows x64 VMware 7.0 and above Linux x64 XenServer	See the maxView Storage Manager download package for the OS-applicable installation executable.
maxView™ vSphere Plugin	VMware 7.0 and above	See the VMware maxView Storage Manager download package for the OS-applicable installation executable.
Boot USB (offline or pre-boot) for ARCCONF and maxView Storage Manager	Linux x64	See the maxView BootUSB download package for the .iso file.

2. What's New?

This section shows what's new in this release.

2.1. Features

The following table highlights major features supported by each Solutions Release.

Table 2-1. Feature Summary

Feature		Supported Release
Added support to reduce UEFI load time.		2.9.2
Arcconf command to check Nand and NOR Flash type		2.9.0
Redfish Resource to Publish SuperCap Properties Support		2.8.2
Arcconf and Redfish Support in Secureboot ESXi Environment		2.8.2
Remote Key Management of Managed SED		2.8.0
Multi-Actuator Drive Support Enhancements		2.7.4
Managed SED Adapter Password Support		2.7.2
Managed SED Local Mode Support		2.7.0
Multi-Actuator Drive Support		2.7.0
Persistent Event Logging Support		2.6.2
Out of Band Interface Selection Support of MCTP or PBSI		2.5.2
MCTP BMC Management		2.4.8
SMR Drive Support	Enumeration, Unrestricted Command Flow-Through	2.3.0
	SATL Translation for HA/HM SMR Management	
	Identify all Drive Types	
Driver OS Certification Where Applicable		2.3.0
SNMP Management Software Support		2.3.0
4Kn, 512e and 512n Support		2.3.0
Legacy Boot Support		2.3.0
UEFI Driver, Boot Support		2.3.0

2.2. Fixes

2.2.1. Firmware Fixes

2.2.1.1. Fixes and Enhancements for Firmware Release 7.60

This release includes the following fixes and enhancements:

- Added support to clear cache error using PLDM.
- Fixed an issue when taking ownership of OPAL Kingston SEDC600ME1920G SED fails.
 - Root Cause: MSID retrieved from Kingston SEDC600ME1920G drive is 6 bytes long. When authentication session gets opened for MSID authorization, 32 bytes of MSID (original 6 bytes followed by zeroes) is passed to the drive and drive fails the security out protocol command with failure code 01 (NOT_AUTHORIZED).
 - Fix: When session is getting opened using MSID key, fill out the context data of the key equal to the length of the MSID key retrieved from the drive.
 - Risk: Low
- Fixed an issue that otherwise owned SED changes back to foreign SED after a controller reboot
 - Root Cause: When deleting a foreign volume, the foreign SEDs in the volume are set to otherwise owned SEDs by the firmware. The otherwise owned flag is not saved to the SED's

- datastore. After a controller reset, the firmware reads the SED's datastore, gets the foreign SED flag, not the otherwise owned flag.
- Fix: When the firmware sets the otherwise owned flag for the SED, saves it to the datastore. When reboot happens, the firmware can read the otherwise owned flag and restore it to otherwise owned state.
 - Risk: Low
 - Fixed a firmware crash/lockup issue when the Non-Disruptive Software Reset (NDSR) operation is performed on the expanders during the logical drive rebuilding and a drive is failed by firmware.
 - Root Cause: During NDSR operation and the logical drive rebuilding, firmware detected the I/O timeout on the rebuilding drive and issued a device reset to this drive. After NDSR operation, firmware restarted the device discovery and reset some internal variables. At the same time right after NDSR operation, firmware failed the I/O timeout drive due to the number of write retries failure on this drive. During failing the I/O timeout drive, firmware tried to clean up all outstanding requests already delivered to this failed drive and accidentally referencing the NULL pointer variable was reset at the restart of the device discovery. This leads to firmware crash/lockup.
 - Fix: During clean up all outstanding requests for the failed drive, firmware will check for the NULL pointer variable and skip the cleanup process.
 - Risk: Low
 - Fixed an issue where 4Kn SED physical drive tries to rebuild into a 512B block logical drive.
 - Root Cause: During the device hot-plug, firmware validates the replacement drive for block/sector size. If it is a mismatch, firmware fails the drive with the failure reason code of the non-MSED type mismatch in hot plug. But in this case the hot-plugged drive is MCHP foreign SED and is in the locked SED state, it cannot be failed with the non-MSED failure reason code. As it is not in the allowable failure reason code list for MSED support. As the result, this 4Kn MCHP foreign SED is still exposed to the host and allows it to be imported. After the import, it leads to the volume rebuild and drive failure on this 4Kn MCHP foreign SED.
 - Fix: The 4Kn MCHP foreign SED should be failed with the failure reason code of the MSED type mismatch in hot plug. So, it won't be exposed to the host and allow the import. This 4Kn MCHP foreign SED must be replaced with a SED that matches the block size.
 - Risk: Low
 - Fixed an issue where the logical drive entered the NEEDS_REBUILD state upon creation.
 - Root Cause: The issue occurred when an array was deleted during a spare rebuild process, with a RAID0 logical drive as the last logical drive. The firmware failed to clear the mapping between the data drives and spare drives, leaving stale data. When a new logical drive was created using the same set of data drives, the leftover stale data caused the logical drive to transition into the NEEDS_REBUILD state.
 - Fix: The firmware has been updated to ensure all drive mappings associated with an array are cleared during its deletion, preventing stale data from causing issues in future logical drive creation.
 - Risk: Low
 - Fixed an issue where drive connector details were incorrectly logged in persistent events.
 - Root Cause: The firmware, while logging persistent events with drive details, attempted to convert the connector name provided by the hardware. If the connector name was limited to two characters, the conversion would fail, preventing the data from being captured and posted into persistent events.
 - Fix: Updated the firmware to include a condition that checks the string length of the connector name. If the name is too short, it is converted to the required format to ensure proper logging in persistent events.

- Risk: Low
- Fixed an issue where BMC unable to retrieve controller information due to timeout errors.
 - Root Cause: The firmware's Out-of-Band (OOB) message threshold timeout was set to 100 milliseconds. This timeout was insufficient in scenarios where the firmware required additional processing time due to workload, resulting in commands timing out after 100 milliseconds.
 - Fix: The firmware's OOB message threshold timeout was increased to 300 milliseconds to ensure adequate processing time. This adjustment aligns better with the BMC's request timeout of 500 milliseconds, preventing unnecessary timeout errors.
 - Risk: Low
- Fixed an issue where multiple logical drives incorrectly showed a REBUILDING state at the same time when both transformation and rebuild processes were queued.
 - Root Cause: When a transformation is queued, the firmware creates two RAID metadata copies (old and new configurations) for each logical drive. During a spare rebuild, the firmware correctly updates the rebuild status (NEEDS_REBUILD → REBUILDING → OK) in the old configuration but fails to update the new configuration. As a result, the new configuration remains stuck in the REBUILDING state. Once the transformation completes and switches to the new configuration, all logical drives incorrectly show the REBUILDING state simultaneously.
 - Fix: The firmware now updates the rebuilding status in both old and new configurations immediately after the rebuild is completed.
 - Risk: Low
- Fixed an issue where changing the master key of a ManagedSED logical drive failed when LU cache was active.
 - Root Cause: When the user initiates a master key change, the firmware loops through the managed SED drives to update the key. At the same time, LU cache flush requests are being executed, which interfere with the password change process. This conflict causes the operation to fail and marking the affected drive as foreign.
 - Fix: Added a condition in the firmware to temporarily pause all I/O requests when a password change request is initiated. Once the master key change is successfully completed, resume all I/O operations.
 - Risk: Low
- Fixed an issue where tools and the BMC did not report a physical drive's status as Predictive Failure (PF), even when the PF LED was lit.
 - Root Cause: When a physical drive sets the DEXCPT bit to 1 in the Informational Mode Page (0x1C), it indicates that the drive will not report any deferred exception conditions, including Predictive Failure warnings. However, some physical drives set the DEXCPT bit to 1 but still report Predictive Failure warnings back to the firmware. In such cases, the firmware updates the drive status and LEDs accordingly, but tools and the BMC do not fetch this data because the DEXCPT bit is set to 1. As a result, the tools and BMC incorrectly display the drive status as "Online" instead of "Predictive Failure."
 - Fix: The firmware has been updated to check whether the DEXCPT bit on the physical drive is configurable. If the drive supports configurability, the firmware will reset the DEXCPT bit to 0. This ensures that Predictive Failure status is properly reflected across tools, the BMC, and the PF LED, providing consistent and accurate reporting.
 - Risk: Low
- Fixed an issue where a healed array would fail after a reboot if a hot-removed physical drive was reconnected in an offline state.

- Root Cause: The firmware tracks RAID metadata using a counter to identify the latest RIS copy across drives in the array. When a physical drive is hot removed, it retains outdated RAID metadata. If a logical drive in the array is healed from a failed state, the metadata is updated, and the counter is reset to 0. On the next reboot, if the hot-removed drive is reconnected in an offline state, its outdated metadata can incorrectly take priority, causing the array to fail.
 - Fix: The firmware now checks if drives are part of other logical drives in the array before resetting the RAID metadata counter to 0 during a heal operation. If they are, the counter is not reset, preventing outdated metadata from causing array failure.
 - Risk: Low
- Fixed an issue where controller will Lockup During Surface Scan.
 - Root Cause: When Unrecoverable Read Errors (UREs) occurred in an unmapped region of a logical drive, the firmware attempted to fix the UREs to maintain fault tolerance on the partial stripe of the logical drive. During this process, the firmware generated several internal read requests. However, stale information left in the internal resources used for these requests caused subsequent request processing to result in a controller lockup.
 - Fix: The firmware has been updated to clear stale data from internal resources before using them to address UREs in unmapped regions of a logical drive.
 - Risk: Low
- Fixed an issue where Controller Reporting Uncorrectable DDR ECC Errors at Boot.
 - Root Cause: The DDR cache was accessed before initialization, resulting in uncorrectable ECC errors. In certain scenarios, the firmware attempted to access the cache prematurely, leading to these errors.
 - Fix: The firmware has been modified to ensure the cache is fully initialized before any access. This prevents uncorrectable DDR ECC errors and ensures stable system operation during boot.
 - Risk: Low
- Fixed an issue where firmware incorrectly reporting the queue depth for Zoned Block Devices (ZBD) during hot-plug event.
 - Root cause: During the handling of ZBD physical drive hot-plug, the firmware prematurely exited the setup process before updating the queue depth reported by the physical drive.
 - Fix: The firmware has been updated to ensure that the queue depth is correctly updated in the drive parameters of the ZBD physical drive during hot-plug handling.
 - Risk: Low
- Fixed controller Lockup (0x3120C) when enabling MCTP due to flooded requests.
 - Root Cause: During initial MCTP communication, BMC sends more MCTP requests within a span of short time, since the route is not created and due to the controller VDM hardware issue, to get physical address of the initiator, resolve EID for all the requests is sent. When the response for the first resolve EID is received, route is created and then the outstanding requests with the same EID which are put in the list are cleared. Later when the response for the second resolve EID is received, controller goes through the list of outstanding requests for which the response of resolve EID is required, to get the physical address and to create the route. Since the outstanding requests are cleared when the first resolve EID response is received, there is no element present in the list and so the route is created with invalid address. Now that for the same endpoint ID EID(BMC) two routes got created which is incorrect. When BMC requests data from one of the routes, firmware sends data in another route and then host acknowledges with different route and hence completion ID, tag mismatched in PCIe and PCIe core detected mismatch and PCIe core hardware locked up the controller through fatal error.

- Fix: During the second resolve EID response, check whether the route is already created for the EID and if it is created, then use the same information for updating the routing table. The outstanding requests would have been already cleared during processing of first resolve EID response.
- Risk: Low
- Fixed an issue to re-enable `wr_32` command in production image for Chiplink testing.
 - Root Cause: Write 32 bit command was disabled in production image which had caused chiplink testing fail as write 32 command is required by Chiplink in its test.
 - Fix: Enable write 32 bit command in production image so that Chiplink testing can go ahead.
 - Risk: Low
- Fixed an issue to change the drive dampen timer value for expander attached SATA drives.
 - Root Cause: Dampen timer is used to track the time, from which the drive has been queued up in SSU queue. In Expander, based on SSU algorithm drives queued up in SSU queue will be spun up. From controller side, dampening timer is used to wake-up the drive after the drive in SSU queue has expired, which is not in sync with the proposed TTR time by drive vendor in case of SATA drives. Here the SATA drive is getting hot removed within 10 secs, which is not expected behavior from BC side.
 - Fix: During boot-up the drive dampen timer value will be initialized to 10 secs, which is default and will be used only for SAS drives. But when the timer is getting started, if it is SATA drive, then the drive dampen timer value will be overwritten to 45 secs, which is the TTR value proposed by SATA drive vendors.
 - Risk: Low

2.2.2. UEFI Fixes

Note: Microsoft signed and secure boot is supported.

2.2.2.1. Fixes and Enhancements for UEFI Driver 2.18.4/Legacy BIOS 2.18.2

This release includes the following fixes and enhancements:

- Added enhancement to show separate fields for the drive vendor and model name in the disk information.
 - Implementation Details: The Disk Information menu will show separate fields for the drive vendor and drive model data.
- Added support to transfer encryption keys securely in remote key management mode.
 - Implementation Details: Encryption keys received from the EFI KMS protocol, which are required for Controller Based Encryption and Controller Managed SED remote mode encryption, are transferred securely by the UEFI driver.
- Added support to enhance Drive write cache status information.
 - Implementation Details: The drive write cache status information in HII has been updated to handle the case for NVMe drives that do not support the command to retrieve the cache status.
- Fixed an issue where disk utilities display an invalid box number for drives after failing the active path.
 - Root Cause: Incorrect box number shown for the alternative path.
 - Fix: Show the box number from the alternative path only if the controller marks it as valid.
 - Risk: Low

2.2.3. Driver Fixes

2.2.3.1. Fixes and Enhancements for Linux Driver Build 2.1.34-035

This release includes the following fixes and enhancements:

- Fixed an issue where the kernel call trace when calling `smp_processor_id()` in real-time kernel.
 - Root Cause: `smp_processor_id()` checks to see if preemption is disabled. If enabled, it will issue an error message followed by a call to `dump_stack()`. `smp_processor_id()` can potentially return an inaccurate CPU ID if a context switch happens during its execution, while `raw_smp_processor_id()` is designed to avoid this issue by holding a lock internally while retrieving the ID. This makes it more reliable for CPU identification during highly time-sensitive situations.
 - Fix: Switch to using `raw_smp_processor_id()`.
 - Risk: Low
- Fixed a rare race condition between our scan thread and offline handler.
 - Root Cause: The scan thread was removing a SCSI device before the offline handler could access the SCSI device pointer to set its state to OFFLINE. The offline handler has been updated to check for a NULL SCSI device pointer and the device list is now protected with a `device_list` lock.
 - Fix: Add check for null SDEV in `pqi_take_ctrl_devices_offline()` function.
 - Risk: Low

2.2.3.2. Fixes and Enhancements for FreeBSD Driver Build 4620.0.1010

There are no known fixes for this release.

2.2.3.3. Fixes and Enhancements for Windows Build 1016.18.0.1014

There are no known fixes for this release.

2.2.3.4. Fixes and Enhancements for VMware Driver Build 4856.0.105

This release includes the following fixes and enhancements:

- Fixed an issue where initialization performs a capabilities inquiry to get values for max elements, and maximum and minimum size of the elements.
 - Root Cause: PSOD indicates a divide-by-zero happened when computing a value from PQI capability structure member `max_iq_elem_len`. Driver checks for capability values that don't align with driver expectations, but only prints a warning rather than halting the initialization process.
 - Fix: Show error message and exit the initialization smoothly if capabilities data don't match driver's expected values.
 - Risk: Low

2.2.4. Management Software Fixes

2.2.4.1. Fixes and Enhancements for Arcconf/maxView™ Build 4.26.00.27449

This release includes the following fixes and enhancements:

- Added support in Arcconf GETSTATUS command to filter and display the task details per logical/physical device.

2.2.4.2. Fixes and Enhancements for PLDM Release 6.50.11.0

This release includes the following fixes and enhancements:

- Added support that require privileged operations token for all RDE ACTION operations. Certain controllers require a custom request header to be sent to the RDE device for data destructive RDE Write operations. The current implementation only applies this condition to the `Drive.SecureErase` and `Storage.ResetToDefaults` RDE ACTIONS and all RDE CREATE, DELETE, and UPDATE operations. Modified the helper function which determines whether or not an RDE operation is data destructive to consider all RDE ACTION operations to fall in that category. This will result in the `Drive.Reset` and `Storage.SetEncryptionKey` RDE ACTIONS now requiring a custom request header to be passed in with the RDE operation request.

- Added support to provide meaningful updateInterval value for Numeric Sensor PDRs.
 - Set the updateInterval to 5 seconds for the controller temperature Numeric Sensor PDR.
 - Set the updateInterval to 60 seconds for the 'hottest or normalized' drive temperature Numeric Sensor PDR.
 - Set the updateInterval to 60 seconds for individual drive temperature Numeric Sensor PDRs on controllers supporting these PDRs.
- Fixed an issue where the controller firmware was not synchronizing status between PLDM Type 2 (PDR) and Type 6 (RDE) data. The present state of the controller composite state sensor Health state can sometimes have a reading that does not match that for the Status.HealthRollup property value obtained from a RDE READ on the Storage resource.
 - Root Cause: A previous fix to the logic used to determine the value to publish for Storage.Status.HealthRollup did not apply that same fix to the controller composite state sensor.
 - Fix: Modified the handling of GetStateSensorReadings requests for the controller composite state sensor to ensure that the Health state reading will match the value of Storage.Status.HealthRollup obtained from an RDE READ operation on the Storage resource.
 - Risk: Low
- Fixed an issue where SSD firmware update is failing due to PLDM error. RequestFirmwareData command when performing firmware update on a SAS drive fails when the drive is in a dual path configuration.
 - Root Cause: RequestFirmwareData sends a SCSI WRITE BUFFER command with mode 0xE. If this SCSI command is sent to the inactive path of the drive, it can cause the drive to return a check condition. PLDM fails the command if a check condition is returned by the drive.
 - Fix: There are two changes:
 - Before performing the first SCSI WRITE BUFFER command, PLDM will now send a Test Unit Ready command to force the drive to clear any check conditions.
 - If the update agent performs a self contained activation via ACTIVATE FIRMWARE command, PLDM will send a Test Unit Ready command to clear any check conditions to ensure the drive's inactive path can accept SCSI passthrough commands.
 - Risk: Low
- Fixed an issue where continuous serial log prints were observed when Predictive Failure state PD is Failed in system. When a drive in predictive failure goes into failure, periodic SCSI error prints appear in the controller serial log.
 - Root Cause: The SCSI error prints are related to repeated attempts to fetch drive monitoring and performance statistics in response to a RDE READ request incoming for the drive's associated DriveMetrics resource. A bug recently introduced into the API requesting this data allowed the request to be sent to controller firmware in spite of the failed status of the drive.
 - Fix: Updated the API submodule to pick up a fix for the erroneous behavior of the API.
 - Risk: Low
- Fixed the incorrect OriginOfCondition in Redfish Message Events when battery is in failure state. Battery Redfish Alerts are generated with incorrect OriginOfCondition data. The OriginOfCondition field contains a link reference to a Redfish Storage Controller resource, rather than a Battery resource.
 - Root Cause: The OriginOfCondition field in the Battery Event is being populated with a StorageController URI. The OriginOfCondition field should contain a link reference to the Redfish resource associated with the Redfish event.
 - Fix: Modified the OriginOfCondition URI for Battery Redfish Alerts to set a link to the Battery resource.

- Risk: Low
- Fixed an issue where performing an RDE UPDATE operation on a Volume resource's ReadCachePolicy property was completing successfully while a CBE rekey was running on the Volume. However, a follow-up RDE READ on the Volume showed that no change to the Volume's ReadCachePolicy was made.
 - Root Cause: Much of the cache RDE UPDATE logic is based on the controller cache status being either unconfigured or okay. In the case of this issue, the controller cache status was temporarily disabled, leading to unexpected behavior in the cache RDE UPDATE logic. Although there were checks already in place to block the cache patch if the controller cache status is temporarily disabled, these checks were not being hit when the controller cache status was in that state.
 - Fix: Adjusted the cache RDE UPDATE logic to no longer depend on the cache status being unconfigured or okay. Accurate error messages will now be returned either from logic within the cache patching function or internal API's called by the function, regardless of controller cache status.
 - Risk: Low
- Fixed an issue where RDE READ for a Drive resource shows the Status.State is Enabled for drives which are ATA security locked and are blocking all I/O requests.
 - Root Cause: PLDM had no means of detecting the ATA security locked state for drives and thus had no logic in place to handle this test case.
 - Fix: Picked up API changes which expose the ATA security enabled / locked states for a drive. Updated the helper functions which determine drive state for both RDE READ and event generation purposes. RDE READ for a Drive resource will now have Status.State of StandbyOffline and Health of Ok for ATA security locked drives. A DriveOffline Redfish alert event will be generated at boot when a drive is powered on in the ATA security locked state, and a DriveOfflineCleared alert will be generated at some point after the ATA security password has been supplied. Note that the controller firmware does not support configuration of drives with ATA security enabled, i.e. with a password set, regardless of locked status.
 - Risk: Low
- Fixed an issue where the HotspareActivationPolicy was showing OnDriveFailure in an allowable value when configure only RAID-0 with Dedicated spare. RDE READ on the Storage resource shows "OnDriveFailure" in the HotspareActivationPolicy@Redfish.AllowableValues array when volume configuration constraints would make this setting invalid, for example, having a RAID 0 Volume with a dedicated spare.
 - Root Cause: The logic for populating the HotspareActivationPolicy@Redfish.AllowableValues array was only checking for controller support and not inspecting the current configuration of Volumes.
 - Fix: Added checks of Volume configuration to make a more accurate determination of when to add the "OnDriveFailure" value to the HotspareActivationPolicy@Redfish.AllowableValues array.
 - Risk: Low
- Fixed an issue where there is a mismatch in Overall Health State and Logical Volume State for PLDM_OTHER_STORAGE_DEVICE_ENTITY_TYPE during RPI. While a Volume is undergoing RPI, GetStateSensorReadings for the logical drive state sensor would show the presentState as Warning, but the Status.Health property for all Volume resources would remain at OK when inspected via a RDE READ operation.
 - Root Cause: Logic which sets Volume Status.Health to Ok while RPI is underway was not being applied when calculating the presentState of the logical drive state sensor.

- Fix: The logic to process a GetStateSensorReadings request on the logical drive state sensor has been updated to prevent RPI from resulting in a presentState of Warning.
- Risk: Low

2.3. Limitations

2.3.1. General Limitations

This release includes the following general limitation:

- The following are the limitations of Multi-Actuator:
 - Supports only
 - HBA drive
 - Windows/Linux/VMware
 - Intel/AMD
 - UEFI mode (for multi-LUN display)

2.3.2. Firmware Limitations

2.3.2.1. Limitations for Firmware Release 7.60

This release includes the following firmware limitations:

- There is a chance of controller lockup if configuration changes (`set_config` commands) are issued simultaneously from the out-of-band (OOB) management path during high I/O workload testing. While this scenario does not result in data loss, it can cause temporary disruption to system operations.
 - Workaround: Reboot the system.
- Downgrading firmware from version 7.11 or later to a version prior to 7.11 during the remapping of Unrecoverable Read Errors (URE) in unmapped regions will cause the controller to lock up.
 - Workaround: Before downgrading the firmware, ensure that the event log shows `DETAIL_SA_READ_ERR` followed by `DETAIL_SA_READ_ERR_FIXED`. This confirms that the remapping process is complete and can prevent the controller from locking up.
- If a boot volume is secured by Managed SED Remote Key Management (RKM) or Managed SED Adapter Password enabled Local Key Management (LKM), it will fail to write Windows memory dump file during Windows OS crash dump.
 - Workaround: Don't use secured volumes as described above as an OS boot logical drive.
- Persistent Event Logs (PEL) are getting cleared when:
 - Upgrading from firmware releases prior to 5.61 to 5.61 or later firmware releases.
 - Downgrading from firmware releases 5.61 or later to firmware releases prior to 5.61.
- Firmware downgrade is blocked if disk-based transformation is in-progress.
 - Workaround: Wait for the transformation to complete and retry the firmware downgrade.
- Transformation is blocked if rebooting after the firmware update is pending or the flashed new firmware version is older than 5.32 B0.
 - Workaround: Reboot the system.
- Logical drive is not detected when disk-based transformation is in-progress during logical drive movement to a different controller and the different controller has a firmware version older than 5.32 B0, or, the firmware downgrade occurred while internal-cache based transformation was in progress, but the Backup Power Source failed before firmware activation.
 - Workaround: Move the logical drive to a controller with firmware version 5.32 B0 or later.
- Firmware downgrade from firmware version 7.11 B0 and newer to any firmware version before 7.11 B0 is blocked if Managed SED is enabled.

- Workaround: Disable Managed SED and try firmware downgrade.
- Managed SED cannot be enabled on the controller, where reboot is pending after firmware downgrade from firmware version 6.22 B0 to any older firmware version.
 - Workaround: Reboot the controller and enable the Managed SED.

2.3.2.2. Limitations for Firmware Release 1.32 Build 0

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

2.3.3. UEFI Limitations

2.3.3.1. Limitations for UEFI Build 2.18.4/Legacy BIOS Build 2.18.2

There are no known limitations for this release.

2.3.4. Driver Limitations

2.3.4.1. Limitations for Linux Driver Build 2.1.34-035

This release includes the following limitations:

- A call trace may be observed when performing a drive hot removal/re-add while I/O is running. This is seen exclusively on RHEL9.4 and has been tracked down to a kernel bug in the blk-mq subsystem. It has been patched starting with `kernel.org kernels 6.14-rc1`.
 - Workaround:
 - i. Stop I/O before doing drive hot removals/additions.
 - ii. Use a non-affected Linux OS release.
 - iii. Update to a later version of RHEL9.
- OS installation hangs when attempting to load the OOB SmartPQI driver DUD.
 - Workaround: This failure occurs when using `inst.dd` even if no driver update is provided. This is an OpenEuler 24.03 installer issue.
- SL-Micro 6.0 fails to boot after installation on 4Kn drives.
 - Workaround: This is a SUSE issue and only workaround is to use non-4Kn drives.
- On some distributions (RHEL7.9, RHEL8.2, RHEL8.3, SLES15SP2, SLES15SP3, OpenEuler 20.03LTS, and 22.03LTS including SP releases), the driver injection (DUD) install will hang if an attached drive (either HBA mode or Logical Volume) has Write Cache enabled.
 - Workaround: There are two workarounds for this issue:
 - Ensure that the Write Cache is disabled for any attached drive.

- For RHEL7.9/8.2/8.3 and OpenEuler 20.03LTS, 22.03LTS, add `rd.driver.blacklist=smartpqi` to the grub entry along with `inst.dd`.
- RHEL driver injection (DUD) install where OS ISO is mounted as virtual media on BMC based servers (non-ILO). Installer will hang after driver injection. It is reported on RHEL 8.5, 8.6, 9.0 to 9.4.
 - 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.
- 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.

2.3.4.2. Limitations for Windows® Driver Build 1016.18.0.1014

This release includes the following limitation:

- A system crash may occur when hibernating a system installed on a Dual Actuator drive.
 - Workaround:
 - Avoid hibernating the system while running heavy I/Os to multiple Dual Actuator drives.
 - Stop running the I/Os to the drives and then hibernate the system.
 - Reboot the server to recover the system.
- A crash dump file will not be created if the system is configured with the OS system files loaded on a partition which is NOT the first partition. If the first partition is deleted and then the system happens to bug check, the crash dump file will not be written out. For example:
 - a. Disk 0 is Array A
 - b. Disk 1 is Array B with the OS on it
 - c. If Array A is deleted and a crash dump occurs without a reboot, the OS will NOT write out the crash dump file.
 - Workaround: This is only seen in the above configuration and if the deletion is done without doing a system reboot. To avoid the problem, make sure the OS is on the first partition or ensure that any time an array is deleted the system is rebooted.
- A Logical drive goes into an offline state after a new array migration.

- Workaround:
 - i. Perform logical disk migration.
 - ii. Run DiskPart.
 - iii. Run the command "List Disk" to identify all the physical disks that have a duplicate unique disk IDs.
 - iv. Run the command "Select Disk X", where X is the physical disk with the duplicate Unique disk ID to be cleaned.
 - v. Run the command "clean". This cleans the physical disk with the duplicate disk ID(aka partition ID).
 - vi. Run command "select disk Y" where Y is the newly migrated logical disk.
 - vii. Run the command "online disk", which will bring the migrated logical drive online.

2.3.4.3. Limitations for FreeBSD Driver Build 4620.0.1010

This release includes the following limitations:

- FreeBSD 13.2 and later OS Installations will fail with the out of box driver.
 - Workaround: Install with inbox driver then update to latest.

2.3.4.4. Limitations for VMware Driver Build 4856.0.105

This release includes the following limitations:

- A system may PSOD if attached JBOD enclosures are power cycled while the system is running.
 - Workaround: Avoid powering OFF JBOD enclosures while the system is running.
- 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.
- Customers may encounter failures when attempting to add new Logical Drives (LD), particularly in cases involving a dead path.
 - Workaround: To facilitate recovery of new LD, customers are required to clear the dead path initially. Following the clearance of the dead path, if the newly created LD is still not exposed, then it is required to initiate a driver level rescan using the appropriate management tool. If clearing the dead path fails, a host reboot is required.

2.3.5. Management Software Limitations

2.3.5.1. Limitations for Arcconf/maxView Build 4.26.00.27449

There are no known limitations for this release.

2.3.5.2. Limitations for PLDMC Release 6.50.11.0

There are no known limitations for this release.

2.3.6. Hardware Limitations

This release includes the following hardware limitations:

- Two Wire Interface (TWI) address conflicts can cause system DDR memory to not be discovered.
 - Description: The HBA1100 boards include two TWI targets on the host-facing SMBUS interface with the following slave addresses:
 - 0xA0 – Field Replaceable Unit (FRU) EEPROM
 - 0xDE – PBSI (default)

According to the JEDEC specification, the default TWI addresses for the DDR SPD is 0xA0-0xAE (the spec uses 7 bit addressing which is 0x50-0x57). On platform system board designs with SMBUS wiring that has both PCIe slots and DDR slots shared on the same TWI bus, the TWI devices for the DDR and Smart controller are exposed to address conflicts which can result in the system memory not being discovered. The Smart controller PBI interface defaults to a value of 0xDE (0x6F in 7-bit addressing) and is not a problem unless it is changed to an address that conflicts with the JEDEC defined values. The Smart controller FRU EEPROM is hardwired to 0xA0.

- Workaround: None available. If this issue is encountered, contact your Microchip support engineer to determine the next steps for your system.
- Performance with workaround: Not applicable
- Performance without workaround: Not applicable

3. Updating the Controller Firmware

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



Important:

- If Managed SED is enabled, do not downgrade firmware to version 5.00 or earlier because they do not support Managed SED capabilities. Disable Managed SED if downgrading to firmware versions 5.00 or earlier.
- When downgrading firmware, there may be cases when newer hardware is not supported by an older version of firmware. In these cases, attempting to downgrade firmware will be prevented (fail). It is recommended to regularly qualify newer firmware versions, to ensure that newer hardware is supported in your system(s).

3.1. Updating the Controller Firmware

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

Note:

1. Complete these procedures exactly as described for proper functionality. If you do not follow all of the steps correctly, you could encounter unusual runtime behavior.

Flashing the board to the latest firmware:

This section describes how to update all the firmware components on HBA 1100 Adapter boards to the latest release.

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

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

Note:

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

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

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

Note:

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

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

1. **Mandatory:** Flash the controller with the provided "SmartFWx100_v1.29_b314.bin" image with arcconf/maxView software.

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

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

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

4. Installing the Drivers

See the “*Microchip Adaptec® HBA 1100 Series Host Bus Adapters Installation and User's Guide* (DS00004281D, previously ESC-2161232)” for complete driver installation instructions.

5. Revision History

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

Revision	Date	Description
R	05/2025	SR 2.9.4 Production Release.
Q	02/2025	SR 2.9.0 Patch Release to update "Fixes and Enhancements for Arcconf/maxView™ Build 4.18.00.26842" section.
P	12/2024	SR 2.9.2 Production Release.
N	07/2024	SR 2.9.0 Production Release.
M	03/2024	SR 2.8.4 Production Release.
L	12/2023	SR 2.8.0 Patch Release with maxView version B26068
K	11/2023	SR 2.7.0 Patch Release with maxView version B25339
J	11/2023	SR 2.8.2 Production Release
H	07/2023	SR 2.8.0 Production Release
G	03/2023	SR 2.7.4 Production Release
F	11/2022	SR 2.7.2 Production Release
E	08/2022	SR 2.7.0 Production Release
D	03/2022	VMware driver version updated from 4250.0.120 to 4252.0.103
C	02/2022	SR 2.6.6 Production Release
B	12/2021	SR 2.6.4.1 Patch Release with maxView™ version B24713. Updated Fixes and Enhancements for maxView Storage Manager/ARCCONF section for log4j vulnerabilities.
A	11/2021	SR 2.6.4 with VMware driver version 4230.0.103 (previously ESC-2162192)
22	08/2021	SR 2.6.2 with VMware driver version 4150.0.119
21	04/2021	SR 2.6.1.1 with VMware driver version 4054.2.118
20	03/2021	SR 2.6.1 with VMware driver version 4054.1.103
19	02/2021	SR 2.6 Production Release
18	10/2020	SR 2.5.4 Production Release
17	08/2020	SR 2.5.2.2 Production Release with Firmware 3.00
16	02/2020	Update for SR 2.5.2
15	10/2019	Update for SR 2.5
14	08/2019	Update for SR 2.4.8 Release
13	03/2019	Update for SR 2.4.4 Release
12	01/2019	SR2.4 Production Release
11	10/2018	SR2.3 firmware update with Cavium/ARM support and Ubuntu driver.
10	06/2018	SR2.3 Production Release
8	10/2017	Update supported OSs
8	10/2017	First Production Release
1-7	10/2016 to 07/2017	Pre-Production Release.

Microchip Information

Trademarks

The “Microchip” name and logo, the “M” logo, and other names, logos, and brands are registered and unregistered trademarks of Microchip Technology Incorporated or its affiliates and/or subsidiaries in the United States and/or other countries (“Microchip Trademarks”). Information regarding Microchip Trademarks can be found at <https://www.microchip.com/en-us/about/legal-information/microchip-trademarks>.

ISBN: 979-8-3371-1176-6

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.

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 products are 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.