

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.....	12
3. Updating the Controller Firmware.....	17
3.1. Updating the Controller Firmware.....	17
4. Installing the Drivers.....	19
5. Revision History.....	20
Microchip Information.....	21
Trademarks.....	21
Legal Notice.....	21
Microchip Devices Code Protection Feature.....	21

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	2.10.0
Package Release Date	August 13, 2025
Firmware Version	7.81 B0'
UEFI Driver Version	2.20.1
Legacy BIOS	2.20.1
Driver Versions	Windows SmartPQI: <ul style="list-style-type: none"> Windows Server 2022/2025: 1016.24.0.1002 Windows 10/11: 1016.24.0.1002 Linux SmartPQI: <ul style="list-style-type: none"> RHEL 7/8/9/10: 2.1.36-026 SLES 12/15: 2.1.36-026 Ubuntu 22/24: 2.1.36-026 Debian 11/12: 2.1.36-026 Oracle Linux 7/8/9: 2.1.36-026 Citrix XenServer 8: 2.1.36-026 BC Linux 7: 2.1.36-026 OpenEuler 22/24: 2.1.36-026 VMware SmartPQI: <ul style="list-style-type: none"> VMware 9.0/8.0: 4862.0.104 FreeBSD SmartPQI: <ul style="list-style-type: none"> FreeBSD 13/14: 4660.0.1003
arcconf/maxView™	5.03.00.27960
PLDM	6.55.7.0

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
Arconformupdate	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 10, 9.6, 9.5, 8.10	Arm®
SuSE Linux Enterprise Server 12 SP5	Arm
SuSE Linux Enterprise Server 15 SP7	Arm
Ubuntu 24.04.2, 22.04.5	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 9.6, 9.5, 9.0, 8.10, 8.0	x86_64
SuSE Linux Enterprise Server 12, SP5	x86_64
SuSE Linux Enterprise Server 15 SP7, SP6	x86_64

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

Drivers	Version
Oracle Linux: <ul style="list-style-type: none"> 8.10 UEK7U3 9.4 UEK7U3 9.5 UEK8U1 9.6 UEK8U1 	x86_64
Ubuntu 24.04.2, 24.04.1, 22.04.5, 22.04.4	x86_64
Debian 12.10, 12.9, 11.11	x86_64
Citrix xenServer 8.4	x86_64
Fedora 42 (inbox only)	x86_64
OpenEuler 24.03 SP1 LTS	x86_64
OpenEuler 22.03 SP4 LTS	x86_64
SLE-Micro 6.1 (Inbox only)	x86_64

Table 1-7. FreeBSD and VMware SmartPQI Drivers

Drivers	Version
FreeBSD 14.3, 14.2, 13.4	x64
VMware 9.0, 8.0 U3/U2	x64

Host Management Software

Table 1-8. Host Management Utilities

Description	OS	Executable
ARCCONF Command Line Utility	Windows® x64 Linux® x64 VMware 8.0 and above XenServer FreeBSD x64 Linux ARM	See the Arccconf 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 8.0 and above Linux x64 XenServer	See the maxView Storage Manager download package for the OS-applicable installation executable.
maxView™ vSphere Plugin	VMware 8.0 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 7.81

This release includes the following fixes and enhancements:

- Added support for improved firmware logging related to lockup code 0x27006.
- Added support for SSD life expectancy logging.
- Fixed an issue where running high I/O workload together with simultaneous inband/outband management operations could cause controller lockup 0x1E30.
 - Root Cause: When hosts sent a LUN reset, the system attempted to complete all outstanding commands. If a configuration change was requested at the same time, the system could re-enter the command completion process, leading to an unexpected conflict and resulting in a controller lockup.
 - Fix: The system now checks if a command completion process due to a LUN reset is already in progress. If so, any new configuration request is safely aborted with an appropriate error status to prevent the lockup.
 - Risk: Low

- Fixed an issue where Persistent Log (PLOG) logs were not preserved on cards with NAND having block size smaller than PLOG size.
 - Root Cause: The process for backing up PLOG logs relied on the NAND block size being large enough for the entire PLOG to be copied in a single operation, with the DDR backup region aligned to the NAND block size. However, when NAND devices with smaller block sizes were used, this caused misalignment between the backup region and the PLOG area, resulting in only partial PLOG data being saved. This led to data integrity (CRC) failures during reload and loss of PLOG logs.
 - Fix: The backup process now aligns the PLOG area correctly for smaller NAND block sizes, ensuring the entire PLOG is reliably saved and restored.
 - Risk: Low
- Fixed an issue where an IO deadlock/hang occurs on RAID 0 during Predictive Spare Rebuild, resulting in a host IO timeout and LUN reset.
 - Root Cause: The host IO requests for the Predictive Fail (PF) drive are stuck in the RAID stack queues. These requests are queued, instead of being sent to a drive, because a rebuild is running at high priority. Concurrently, the RAID 0 rebuild waits in a loop, checking if all host IO requests have obtained a stripe lock. This situation leads to a deadlock, causing the IO timeout and LUN reset.
 - Fix: RAID Stack will cancel the Rapid Rebuild at the start of the RAID 0 rebuild while waiting for all host write requests to acquire the stripe lock. Then the host write requests can be sent to the basecode/target and are not stuck in the RAID Stack's internal queues.
 - Risk: Low
- Fixed an issue where the RAID Controller reports Online Firmware Activation even though support has been removed.
 - Root cause: A code change was not made in the PQI initialization of the configuration table to indicate that OFA is not supported. Conversely, the Linux SmartPQI driver still displays the feature in the message log "smartpqi 0000: 5c: 00.0: Online Firmware Activation enabled".
 - Fix: During initialization of the PQI configuration table and handling the PQI configuration table feature request for OFA from the host PQI driver, RAID Stack will check if the current running FW supports OFA. Then it will set or clear the PQI configuration table supported and enabled field accordingly.
 - Risk: Low
- Fixed an issue where the bay number is 255 for all drives connected to a "zoning disabled" enclosure, only when it's connected to a controller with a "zoning enabled" enclosure.
 - Root Cause: The combination of the expander's zoning support in the RAID Stack logic and the expander reporting connector type as 0x00 (that is, no information) and device slot number is all 0xFF. One enclosure reports zoning enabled, but the other enclosure does not. The "zoning disabled" enclosure reports the connector type as 0x00 (no information) and all device slot numbers as 0xFF. RAID Stack treats all enclosures/expanders as zoning enabled, even if only one expander has zoning enabled.
 - Fix: Added an additional check for zoning enabled on each expander when RAID Stack assigns the expander PHYs to device indexes.
 - Risk: Low
- Fixed an issue where the bay number is 255 for all drives connected to an enclosure after the enclosure was hot-plugged by powering on.
 - Root Cause: Some enclosures take time to become ready after being powered on. They might return a "not ready" status when the RAID Stack tries to read the SCSI enclosure services (SES) configuration pages from the SES device. The "not ready" SCSI status returned might be "Unit Attention - Logical Unit Is In Process Of Becoming Ready". As a result of this

status, the RAID Stack would not receive the SES drive bay information, causing the drive bay numbers to result to 255.

- Fix: The RAID Stack currently includes logic that waits up to 1 minute if the SES device is not ready. This logic was implemented to support enclosures returning a "not ready" status of "Not Ready - Enclosure Services Unavailable". Using the existing logic, an additional check has been added for the new SCSI "not ready" status code: "Unit Attention - Logical Unit Is In Process Of Becoming Ready".
- Risk: Low
- Fixed an issue with fixed format sense data for ATA passthrough command.
 - Root Cause: In the fixed format sense data returned by SAT layer for ATA pass-through commands, LBA (7: 0) and LBA (23: 16) bytes in command specific information are interchanged.
 - Fix: Corrected the issue in the fixed format sense data returned by SAT layer for ATA pass-through commands, according to T10 spec.
 - Risk: Low
- Fixed an issue where SCSI inquiry command was not able to get drive serial number.
 - Root Cause: SCSI inquiry command was unable to retrieve the drive serial number. To obtain the serial number of the drive, the Get VPD Serial Number command needs to be sent for SATA drives.
 - Fix: When the SCSI standard inquiry command is sent, copy the serial number of the SATA drive in swapped position to fit the ACS specification. The spec does not specify anything about the serial number in standard SCSI inquiry command output data, so the vendor-specific parameters field is used to fill the serial number of the SATA drive in the SCSI inquiry command output data.
 - Risk: Low
- Fixed an issue where device reset is due to ACK_NAK timeout in command phase.
 - Root Cause: Whenever there is an ACK_NAK timeout error sent during the command phase by OSSP for a device, the firmware enters into NCQ error handling mode and sends the Read Log Ext command to the device. However, the expander processing I/Os enters a different state machine as the Read Log Ext command is received in the middle of an I/O transaction, hence an error response to I/O is sent and the read log ext response is not sent. Consequently, the host, after timing out on the command, sends a LUN reset/device reset to the device.
 - Fix: Whenever OSSP sends an ACK_NAK timeout error during the command phase for a device, and if the device is a SATA device, return an error to the command to the host as the command has not reached the target.
 - 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 2.20.1/Legacy BIOS 2.20.1

There are no known fixes for this release.

2.2.3. Driver Fixes

2.2.3.1. Fixes and Enhancements for Linux® Driver Build 2.1.36-026

This release includes the following fixes and enhancements:

- Fixed a system crash issue caused by a divide-by-zero condition.
 - Root Cause: RAID map entries `blocks_per_row` and `strip_size` are used as divisors for AIO calculations. These entries can be zero, leading to a divide-by-zero issue.
 - Fix: Moved the check for 0 value higher in the code and added an additional check for `rmd->strip_size`.
 - Risk: Low
- Fixed an issue where after a device has been removed, it's possible for a previously scheduled work item for a LUN reset to be executed.
 - Root Cause: A race condition occurs between scheduling a work item for a LUN reset due to the abort handler being called and the device possibly being removed by a call to the `slave_destroy` function.
 - Fix: The solution includes the following changes:
 - In the device reset handler function, added a check to verify if the device is still in the controller's SCSI device list. If it is not, the handler exits without performing any actions.
 - In the slave destroy function, added code to cancel any scheduled TMF work that has not been executed.
 - In the slave destroy function, the device is now freed under the protection of the LUN reset mutex. This prevents the device resources from being freed while a LUN reset is occurring.
 - Risk: Medium
- Added a timeout value to RAID path requests to physical devices.
 - Root Cause: A driver change was requested to address a controller lockup condition that occurs if an I/O is never completed.
 - Fix: The driver now adds a timeout value to the requests sent to physical devices through the RAID path.
 - Risk: Low

2.2.3.2. Fixes and Enhancements for FreeBSD Driver Build 4660.0.1003

This release includes the following fixes and enhancements:

- Added documentation support for `device.hints` to the man page.
- Fixed an issue where setting `device.hints` for the SmartPQI driver caused a system crash.
 - Root Cause: Incorrect pointer casting was causing a crash when the `resource_long_value()` function was called.
 - Fix: Corrected the pointer issue introduced when fixing code related to changes made for FreeBSD.org repository updates. Additionally, corrected the size of `minor_version` in the `driver_info` structure and adjusted the `PQISRC_DRIVER_MINOR` value to fit in an unsigned char.
 - Risk: Medium

2.2.3.3. Fixes and Enhancements for Windows® Build 1016.24.0.1002

This release includes the following fixes and enhancements:

- Added support for NVMe admin passthru requests.
- Fixed an issue related to the upcoming "Driver isolation" requirement for Windows® Server 2025 security.
 - Root Cause: The "Driver isolation" is a new security requirement for Windows Server 2025. It requires the `SmartPqi.inf` section [DestinationDirs] to be set to `DefaultDestDir = 13`. This setting ensures that driver package files

are placed in the `Windows\System32\DriverStore\repository` instead of the `Windows\System32\Drivers` folder.

- Fix: The SmartPQI driver INF file is now compliant with "Driver isolation" requirements.
- Risk: Low

2.2.3.4. Fixes and Enhancements for VMware Driver Build 4862.0.104

This release includes the following fixes and enhancements:

- Fixed an issue where the driver's controller structure field was too small for the full ASCII firmware version.
 - Root Cause: The Identify Controller data buffer has two firmware version fields. The older field was smaller, on which the driver's controller structure field was based upon. Newer controllers use a 32-byte field for versioning.
 - Fix: Expanded the size of the firmware version field in the driver's controller structure.
 - Risk: Low
- Fixed an issue where a message from a periodic check on the controller heartbeat appeared as a system error instead of an informational message.
 - Root Cause: The message was incorrectly coded to print at a messaging level indicating an ERROR instead of a NOTE level.
 - Fix: The issue occurs only when the heartbeat check happens earlier than expected due to some flexibility in VMware's scheduling algorithm.
 - Risk: Low
- Fixed an issue where firmware versioning information was incorrect or blank on some of the controllers.
 - Root Cause: The driver uses a long firmware version ASCII field from the identify controller inquiry buffer, which is not supported on all controller types.
 - Fix: Check controller flags from the identify controller inquiry to verify if the long ASCII firmware version field is supported before using it. If not supported, use the original short firmware version field and build number from the same inquiry.
 - Risk: Low
- Fixed an issue where the driver maintains a list of removed devices to be processed later. When a previously removed device returns, it needs to be deleted from the list. In a specific scenario, the device removal handler attempts to remove a device that has already been removed by the normal device discovery flow.
 - Root Cause: The duplicate list removal attempt causes PSOD because it accesses a null pointer. Meanwhile, list insertions for the removed device list should be protected under spinlocks, but are not.
 - Fix: Removed the redundant list removal attempt in the device removal handler. Added spinlocks around list insertions.
 - Risk: Low

2.2.4. Management Software Fixes

2.2.4.1. Fixes and Enhancements for Arcconf/maxView™ Build 5.03.00.27960

This release includes the following fixes and enhancements:

- Added support for generic NVMe pass-through commands in Arcconf.
- Fixed an issue where the maxView web server lacked the default `-starttls` option, causing email alerts to fail on Windows.
 - Root Cause: The `starttls` option was disabled by default when sending emails, even though the SMTP server supported it.

- Fix: Enabled the `starttls` option for sending emails when the SMTP server supports it.
- Risk: Medium
- Fixed an issue where Backplane/Expander firmware upgrades failed when initiated from `maxView`.
 - Root Cause: The `upgradeMicrocode` call did not reach the webserver module, defaulting to an incorrect path and returning an improper response.
 - Fix: Added proper handling in the webserver module to correctly process the `upgradeMicrocode` call.
 - Risk: Medium

2.2.4.2. Fixes and Enhancements for PLDM Release 6.55.7.0

This release includes the following fixes and enhancements:

- For RDE READ on the VolumeCapabilities resource, the `CapacityBytes@Redfish.AllowableNumbers` annotation range maximum value will now be the larger of the following two possible values:
 - The largest free space range on any existing array of physical drives. This includes free space at any address location on the array, not just trailing space at the end of the array.
 - The sum of capacities of all unassigned drives with a common interface (NVMe/SAS/SATA) and media type (HDD/SSD). For controllers with drives of multiple interface/media type combinations, the largest free space capacity among those groupings will be considered.
- Updated the supported DriveMetrics schema version from 1.0.0 to 1.2.0. Added support for the following DriveMetrics schema properties:
 - `ReadIOPKiBytes`
 - `WriteIOPKiBytes`
- Fixed an issue where RDE READ on a volume resource with caching enabled could publish `ProtectedWriteBack` in the `WriteCachePolicy@Redfish.AllowableValues` array, even when a backup power source was not connected to the controller.
 - Root Cause: The check to add `ProtectedWriteBack` to the `AllowableValues` array only checked if the controller had a cache module with support for caching.
 - Fix: Added a check to ensure the backup power source is connected to the controller and is fully charged before publishing `ProtectedWriteBack` as an `AllowableValue` for the `WriteCachePolicy` in RDE PATCH operations.
 - Risk: Low
- Fixed an issue where RDE READ on a drive resource could sometimes cause a controller lockup.
 - Root Cause: An array out-of-bounds access could occur when RDE READ was received for a Drive resource deemed unsupported for configuration, leading to a lockup.
 - Fix: Updated the storagecore submodule to include a fix in the relevant API, adding bounds checking before accessing the array.
 - Risk: Low
- Fixed an issue where testing the ability to delete volumes through PLDM Type 6 could result in a controller lockup.
 - Root Cause: When deleting a volume with a long-running task thread, if another PLDM Type 6 command, such as `GetSchemaDictionary`, is sent to the controller firmware while it is still processing other commands related to the deletion, it can cause a lockup.
 - Fix: The following commands will now return a completion code of `NOT_READY` if an RDE operation requiring a long-running task is in progress:
 - `NegotiateRedfishParameters`

- NegotiateMediumParameters
- GetSchemaURI
- GetSchemaDictionary
- 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 7.81

This release includes the following firmware limitations:

- 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.
- Firmware will append spaces to serial numbers that are less than 15 characters.
- Firmware doesn't support moving array that is undergoing transformation.

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 2.20.1/Legacy BIOS Build 2.20.1

There are no known limitations for this release.

2.3.4. Driver Limitations

2.3.4.1. Limitations for Linux Driver Build 2.1.36-026

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:
 - Stop I/O before doing drive hot removals/additions.
 - Use a non-affected Linux OS release.
 - 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 1016.24.0.1002

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 4660.0.1003

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 4862.0.104

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.
- Customers may see a 'invalid device' in an environment with a high level of device resets.
 - Workaround: If this occurs, a system reboot will be needed to clear the problem.

2.3.5. Management Software Limitations

2.3.5.1. Limitations for Arcconf/maxView Build 5.03.00.27960

There are no known limitations for this release.

2.3.5.2. Limitations for PLDMC Release 6.55.7.0

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
T	08/2025	SR 2.10.0 Production Release.
S	05/2025	Removed VMware 9.0 support.
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: 979-8-3371-1796-6

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.