

This document does not apply to HPE Superdome servers. For information on HPE Superdome, see the following links:

[HPE Integrity Superdome X](#)  
[HPE Superdome Flex](#)

Information on HPE Synergy supported VMware ESXi OS releases, HPE ESXi Custom Images and HPE Synergy Custom SPPs is available at:

[VMware OS Support Tool for HPE Synergy](#)

Information on HPE Synergy Software Releases is available at:

[HPE Synergy Software Releases - Overview](#)

## Gen10 SPP v2022.09.01.00 Release Notes for VMware ESXi 8.0

[BIOS \(Login Required\) - System ROM](#)

[Driver - Lights-Out Management](#)

[Firmware - Network](#)

[Firmware - NVDIMM](#)

[Firmware - Storage Controller](#)

[Firmware - Storage Fibre Channel](#)

[Software - Management](#)

[Software - Storage Fibre Channel](#)

[Software - System Management](#)

### BIOS (Login Required) - System ROM

[Top](#)

ROM Flash Universal Firmware Package - HPE ProLiant DL360/DL380 Gen10 Plus (U46) Servers

Version: 1.64\_08-11-2022 (**Recommended**)

Filename: U46\_1.64\_08\_11\_2022.fwpkg

### **Important Note!**

#### **Important Notes:**

For customers who had already upgraded to v1.62 7/14/2022 ROM, it is highly recommended to upgrade to this version v1.64 if customers have Intel Optane Persistent Memory 200 Series installed in their system.

When installing with HPE Smart Update Manager (SUM) 8.9.5 or earlier, the warning message "Component Not Signed" might be displayed for SHA-384 signed components. This issue will be resolved in a future version of SUM. For details, see [https://support.hpe.com/hpesc/public/docDisplay?docLocale=en\\_US&docId=a00125612en\\_us](https://support.hpe.com/hpesc/public/docDisplay?docLocale=en_US&docId=a00125612en_us).

This version of the System ROM contains updates aligned with the Intel Product Update (IPU) version IPU.2022.2 guidance.

#### **Deliverable Name:**

HPE DL360 Gen10 Plus/DL380 Gen10 Plus System ROM - U46

#### **Release Version:**

1.64\_08-11-2022

#### **Last Recommended or Critical Revision:**

1.64\_08-11-2022

#### **Previous Revision:**

1.60\_06-01-2022

#### **Firmware Dependencies:**

None

#### **Enhancements/New Features:**

None

#### **Problems Fixed:**

Addressed an issue from previous published ROM release v1.62 7/14/2022 in which systems with Intel Optane Persistent Memory 200 Series may not appear under an OS.

This revision of the System ROM includes the latest revision of the Intel IPU 2022.2 microcode update which provides Intel's mitigation for BIOS advisory and security vulnerability documented as CVE-2022-21233. This security vulnerabilities are documented in Intel Security Advisory INTEL-SA-00657. This issue is not unique to HPE servers.

This revision of the System ROM includes the latest revision of the Intel Reference Code which provides mitigations for BIOS advisories and security vulnerability documented as CVE-2021-33060. This security vulnerabilities are documented in Intel Security Advisory INTEL-SA-00686. These issues are not unique to HPE servers.

This revision of the System ROM includes the latest revision of openssl which provides mitigation for BIOS security vulnerabilities documented as CVE-2022-0778 (<https://nvd.nist.gov/vuln/detail/CVE-2022-0778>) and CVE-2022-2068

(<https://nvd.nist.gov/vuln/detail/CVE-2022-2068>). This security vulnerability is documented in the CVE report site. This issue is not unique to HPE servers.

This revision of the System ROM includes the latest revision of FreeType library which provides mitigations for security vulnerability documented as CVE-2022-27405 (<https://nvd.nist.gov/vuln/detail/CVE-2022-27405>). This issue is not unique to HPE servers.

Addressed an issue where Microsoft Secured-core is not grayed-out when TPM is not present.

Addressed an issue where the QR code in Rom Based Set Up menu linked to a discontinued site.

Addressed an issue where a hot plug event occurring before event handler operates would cause ESXi kernel to call exception waiting for hot plug status from device and then triggered an ASSERT with PSOD (Purple Screen Of Death).

Addressed an issue with displaying FW versions for Intel's latest ATS-M GPU cards.

**Known Issues:**

None

**Fixes**

**Important Notes:**

For customers who had already upgraded to v1.62 7/14/2022 ROM, it is highly recommended to upgrade to this version v1.64 if customers have Intel Optane Persistent Memory 200 Series installed in their system.

When installing with HPE Smart Update Manager (SUM) 8.9.5 or earlier, the warning message "Component Not Signed" might be displayed for SHA-384 signed components. This issue will be resolved in a future version of SUM. For details, see [https://support.hpe.com/hpesc/public/docDisplay?docLocale=en\\_US&docId=a00125612en\\_us](https://support.hpe.com/hpesc/public/docDisplay?docLocale=en_US&docId=a00125612en_us).

This version of the System ROM contains updates aligned with the Intel Product Update (IPU) version IPU.2022.2 guidance.

**Firmware Dependencies:**

None

**Problems Fixed:**

Addressed an issue from previous published ROM release v1.62 7/14/2022 in which systems with Intel Optane Persistent Memory 200 Series may not appear under an OS.

This revision of the System ROM includes the latest revision of the Intel IPU 2022.2 microcode update which provides Intel's mitigation for BIOS advisory and security vulnerability documented as CVE-2022-21233. This security vulnerabilities are documented in Intel Security Advisory INTEL-SA-00657. This issue is not unique to HPE servers.

This revision of the System ROM includes the latest revision of the Intel Reference Code which provides mitigations for BIOS advisories and security vulnerability documented as CVE-2021-33060. This security vulnerabilities are documented in Intel Security Advisory INTEL-SA-00686. These issues are not unique to HPE servers.

This revision of the System ROM includes the latest revision of openssl which provides mitigation for BIOS security vulnerabilities documented as CVE-2022-0778 (<https://nvd.nist.gov/vuln/detail/CVE-2022-0778>) and CVE-2022-2068 (<https://nvd.nist.gov/vuln/detail/CVE-2022-2068>). This security vulnerability is documented in the CVE report site. This issue is not unique to HPE servers.

This revision of the System ROM includes the latest revision of FreeType library which provides mitigations for security vulnerability documented as CVE-2022-27405 (<https://nvd.nist.gov/vuln/detail/CVE-2022-27405>). This issue is not unique to HPE servers.

Addressed an issue where Microsoft Secured-core is not grayed-out when TPM is not present.

Addressed an issue where the QR code in Rom Based Set Up menu linked to a discontinued site.

Addressed an issue where a hot plug event occurring before event handler operates would cause ESXi kernel to call exception waiting for hot plug status from device and then triggered an ASSERT with PSOD (Purple Screen Of Death).

Addressed an issue with displaying FW versions for Intel's latest ATS-M GPU cards.

**Known Issues:**

None

---

**Driver - Lights-Out Management**

HPE iLO Native Driver for ESXi 7.0

Version: 10.8.0 **(Recommended)**

Filename: ilo-driver\_700.10.8.0.6-1OEM.700.1.0.15843807\_20300719.zip

[Top](#)

**Fixes**

- Fixed issue where ilo driver is failing to acquire contiguous physical memory below 4GB causing userworld apps like hponcfg to be unable to communicate with iLO.

## Enhancements

- Added support for vSphere 8.0

---

## **Firmware - Network**

[Top](#)

Mellanox Firmware Package (FWPKG) - Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE

Version: 26.33.1048 (**Recommended**)

Filename: 26\_33\_1048-MCX631102AS-ADA\_Ax.pldm.fwpkg

## **Important Note!**

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available

at: <https://docs.nvidia.com/networking/display/ConnectX6LxFirmwarev26331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

## Prerequisites

FWPKG will work only if the firmware version flashed on the adapter is 20.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

## Fixes

**The following issues have been fixed in version 26.33.1048:**

- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- A configuration issue which flipped the MSB of Partition Key field in CNP packets and led to P\_KEY mismatch between CNP packets and regular packets.
- A rare case where asserts and ext\_synd appeared in dmesg after performing driver restart.
- An issue that prevented the np\_cnp\_sent counter from increasing after it reached its maximum although there were CNPs sent upon receiving ECN-marked packets.

## Enhancements

**New Features and Changes in Version 26.33.1048:**

- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added support for VF migration. The hypervisor can now suspend its VF, meaning from that point the VF cannot perform action such as send/receive traffic or run any command. In this firmware version only the suspend resume mode is supported (on the same VM).
- Added support for VF migration.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

## Supported Devices and Features

HPE Part Number	Mellanox Ethernet Only Adapters	PSID
P42044-B21	Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE	MT_0000000575

Mellanox Firmware Package (FWPKG) - Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE  
Version: 26.33.1048 (**Recommended**)  
Filename: 26\_33\_1048-MCX631432AS-ADA\_Ax.pldm.fwpkg

### Important Note!

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available at: <https://docs.nvidia.com/networking/display/ConnectX6LxFirmwarev26331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

### Prerequisites

FWPKG will work only if the firmware version flashed on the adapter is 20.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

### Fixes

The following issues have been fixed in version 26.33.1048:

- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- A configuration issue which flipped the MSB of Partition Key field in CNP packets and led to P\_KEY mismatch between CNP packets and regular packets.
- A rare case where asserts and ext\_synd appeared in dmesg after performing driver restart.
- An issue that prevented the np\_cnp\_sent counter from increasing after it reached its maximum although there were CNPs sent upon receiving ECN-marked packets.

### Enhancements

New Features and Changes in Version 26.33.1048

- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added support for VF migration. The hypervisor can now suspend its VF, meaning from that point the VF cannot perform action such as send/receive traffic or run any command. In this firmware version only the suspend resume mode is supported (on the same VM).
- Added support for VF migration.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

## Supported Devices and Features

HPE Part Number	Mellanox Ethernet Only Adapters	PSID
P42041-B21	Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE	MT_0000000551

Mellanox Firmware Package (FWPKG) for Mellanox MCX562A-ACAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE  
Version: 16.33.1048 (**Recommended**)  
Filename: 16\_33\_1048-MCX562A-ACA\_Ax\_Bx.pldm.fwpkg

### **Important Note!**

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available at: <https://docs.nvidia.com/networking/display/ConnectX5Firmwarev16331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

### **Prerequisites**

FWPKG will work only if the firmware version flashed on the adapter is 16.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

### **Fixes**

**The following issues have been fixed in version 16.33.1048 :**

- A rare issue that caused the QP not to receive a completion.
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- The card occasionally masked some PCIe AER reporting.
- Incorrect flow of credits blockage that prevented booting during DC cycle test.
- An issue with BMC medium migration from SMBUS to PCIe, and increased FIFOs to pass large packets in case of the migration.

### **Enhancements**

**Important :** Security Hardening Enhancements - This release contains important reliability improvements and security hardening enhancements. HPE recommends upgrading your device firmware to this version to improve the firmware security and reliability of your device.

**New features and changes included in version 16.33.1048:**

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

### **Supported Devices and Features**

HPE Part Number	Mellanox Ethernet Only Adapters	PSID
P10112-B21	Mellanox MCX562A-ACAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE	MT_0000000241

Mellanox Firmware Package (FWPKG) for HPE Ethernet 100Gb 1-port QSFP28 PCIe3 x16 MCX515A-CCAT Adapter : HPE part numbers P31246-B21 and P31246-H21

Version: 16.33.1048 **(Recommended)**

Filename: 16\_33\_1048-MCX515A-CCA\_HPE\_Ax.pldm.fwpkg

### **Important Note!**

For PLDM enabled VPI (Virtual Protocol Interconnect) adapters supporting both InfiniBand mode and Ethernet modes, every firmware version is made available in two different formats at HPE.com:

1. HPE signed PLDM Firmware Package (.FWPKG filename extension) updatable via iLO.
2. Firmware binary (.bin filename extension) updatable via mstflint utility from the Operating System.

Choose the appropriate firmware file format based on your preference and what suits your environment.

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available

at: <https://docs.nvidia.com/networking/display/ConnectX5Firmwarev16331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

### **Prerequisites**

FWPKG will work only if the firmware version flashed on the adapter is 16.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

### **Fixes**

**The following issues have been fixed in version 16.33.1048 :**

- A rare issue that caused the QP not to receive a completion.
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- The card occasionally masked some PCIe AER reporting.
- Incorrect flow of credits blockage that prevented booting during DC cycle test.
- An issue with BMC medium migration from SMBUS to PCIe, and increased FIFOs to pass large packets in case of the migration.

### **Enhancements**

**Important :** Security Hardening Enhancements - This release contains important reliability improvements and security hardening enhancements. HPE recommends upgrading your device firmware to this version to improve the firmware security and reliability of your device.

**New features and changes included in version 16.33.1048:**

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRQ QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.

- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

## Supported Devices and Features

This software package contains the following firmware versions:

Mellanox Ethernet Only Adapters	Firmware Version	PSID
HPE Ethernet 100Gb 1-port QSFP28 PCIe3 x16 MCX515A-CCAT Adapter(P31246-B21 and P31246-H21)	16.33.1048	MT_0000000591

Mellanox Firmware Package (FWPKG) for HPE InfiniBand HDR/Ethernet 200Gb 1-port QSFP56 PCIe4 x16 MCX653105A-HDAT Adapter : HPE part numbers P23664-B21 and P23664-H21

Version: 20.33.1048 **(Recommended)**

Filename: 20\_33\_1048-MCX653105A-HDA\_HPE\_Ax.pldm.fwpkg

## Important Note!

For PLDM enabled VPI (Virtual Protocol Interconnect) adapters supporting both InfiniBand mode and Ethernet modes, every firmware version is made available in two different formats at HPE.com:

1. HPE signed PLDM Firmware Package (.FWPKG filename extension) updatable via iLO.
2. Firmware binary (.bin filename extension) updatable via mstflint utility from the Operating System.

Choose the appropriate firmware file format based on your preference and what suits your environment.

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available

at: <https://docs.nvidia.com/networking/display/ConnectX6Firmwarev20331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

## Prerequisites

FWPKG will work only if the firmware version flashed on the adapter is 20.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

## Fixes

**The following issues have been fixed in version 20.33.1048:**

- An issue where RDMA write may experience performance degradation when working with Adaptive Routing and DCT half-handshake mode.
- An issue that ignored the default value of TX\_SCHEDULER\_BURST when its value in the ini was different than "0".
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- RDMA\_WRITE traffic performance degradation that occurred when working with DC on Adaptive Routing network.
- Bad cache invalidations of destroyed QPs.
- A rare case where asserts and ext\_synd appeared in dmesg after performing driver restart.
- An issue related to host isolation on multi-host systems.

## Enhancements

Security Hardening Enhancements: This release contains important reliability improvements and security hardening enhancements. HPE recommends upgrading your device firmware to this release to improve the firmware security and reliability of your device.

**New features and changes included in version 20.33.1048:**

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Crypto features can be in either wrapped or unwrapped mode. Meaning, the key can be wrapped or in plaintext when running the CREATE\_DEK PRM command. To comply with the requirements specified in FIPS publication, all the created DEKs must be wrapped.  
This feature adds new NV\_CONFIG per device to control this mode, and enables the user to change all the crypto features to wrapped or cleartext.
- Implemented SLD detection code. Surprise Down Error Reporting Capable value was changed from 1 to 0 in boards where the downstream perst was not controlled thus causing SLD detection not to function properly.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

## Supported Devices and Features

This software package contains the following firmware versions:

Mellanox InfiniBand Adapter	Firmware Version	PSID
HPE InfiniBand HDR/Ethernet 200Gb 1-port QSFP56 PCIe4 x16 MCX653105A-HDAT Adapter (P23664-B21 and P23664-H21)	20.33.1048	MT_0000000451

Mellanox Firmware Package (FWPKG) for HPE InfiniBand HDR/Ethernet 200Gb 1-port QSFP56 PCIe4 x16 OCP3 MCX653435A-HDAI Adapter : HPE part numbers P31323-B21 and P31323-H21  
Version: 20.33.1048 (**Recommended**)  
Filename: 20\_33\_1048-MCX653435A-HDA\_HPE\_Ax.pldm.fwpkg

### Important Note!

For PLDM enabled VPI (Virtual Protocol Interconnect) adapters supporting both InfiniBand mode and Ethernet modes, every firmware version is made available in two different formats at HPE.com:

1. HPE signed PLDM Firmware Package (.FWPKG filename extension) updatable via iLO.
2. Firmware binary (.bin filename extension) updatable via mstflint utility from the Operating System.

Choose the appropriate firmware file format based on your preference and what suits your environment.

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available

at: <https://docs.nvidia.com/networking/display/ConnectX6Firmwarev20331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

### Prerequisites

FWPKG will work only if the firmware version flashed on the adapter is 20.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

## Fixes

### The following issues have been fixed in version 20.33.1048:

- An issue where RDMA write may experience performance degradation when working with Adaptive Routing and DCT half-handshake mode.
- An issue that ignored the default value of TX\_SCHEDULER\_BURST when its value in the ini was different than "0".
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- RDMA\_WRITE traffic performance degradation that occurred when working with DC on Adaptive Routing network.
- Bad cache invalidations of destroyed QPs.
- A rare case where asserts and ext\_synd appeared in dmesg after performing driver restart.
- An issue related to host isolation on multi-host systems.

## Enhancements

Security Hardening Enhancements: This release contains important reliability improvements and security hardening enhancements. HPE recommends upgrading your device firmware to this release to improve the firmware security and reliability of your device.

### New features and changes included in version 20.33.1048:

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Crypto features can be in either wrapped or unwrapped mode. Meaning, the key can be wrapped or in plaintext when running the CREATE\_DEK PRM command. To comply with the requirements specified in FIPS publication, all the created DEKs must be wrapped.  
This feature adds new NV\_CONFIG per device to control this mode, and enables the user to change all the crypto features to wrapped or cleartext.
- Implemented SLD detection code. Surprise Down Error Reporting Capable value was changed from 1 to 0 in boards where the downstream perst was not controlled thus causing SLD detection not to function properly.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

## Supported Devices and Features

This software package contains the following firmware versions:

Mellanox InfiniBand Adapter	Firmware Version	PSID
HPE InfiniBand HDR/Ethernet 200Gb 1-port QSFP56 PCIe4 x16 OCP3 MCX653435A-HDAI Adapter (P31323-B21 and P31323-H21)	20.33.1048	MT_0000000592

Mellanox Firmware Package (FWPKG) for HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 MCX653106A-HDAT Adapter : HPE part numbers P31324-B21 and P31324-H21

Version: 20.33.1048 (**Recommended**)

Filename: 20\_33\_1048-MCX653106A-HDA\_HPE\_Ax.pldm.fwpkg

### Important Note!

For PLDM enabled VPI (Virtual Protocol Interconnect) adapters supporting both InfiniBand mode and Ethernet modes, every firmware version is made available in two different formats at HPE.com:

1. HPE signed PLDM Firmware Package (.FWPKG filename extension) updatable via iLO.
2. Firmware binary (.bin filename extension) updatable via mstflint utility from the Operating System.

Choose the appropriate firmware file format based on your preference and what suits your environment.

ConnectX-6 VPI supports having one port as InfiniBand and the other port as Ethernet according to the following matrix of combinations.

Port #2 - InfiniBand				
Port #1 – Ethernet	HDR/HDR100	EDR	FDR	QDR
200GbE/50GbE	supported	not supported	not supported	supported
100GbE/25GbE	supported	not supported	not supported	supported
40GbE/10GbE	supported	not supported	not supported	supported
1GbE	supported	not supported	not supported	supported

Port #2 - Ethernet				
Port #1 - InfiniBand	200GbE/50GbE	100GbE/25GbE	40GbE/10GbE	1GbE
HDR / HDR100	supported	supported	not supported	supported
EDR	supported	supported	not supported	supported
FDR	not supported	not supported	not supported	not supported
QDR/SDR	supported	supported	not supported	supported

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available at: <https://docs.nvidia.com/networking/display/ConnectX6Firmwarev20331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

### Prerequisites

FWPKG will work only if the firmware version flashed on the adapter is 20.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

### Fixes

The following issues have been fixed in version 20.33.1048:

- An issue where RDMA write may experience performance degradation when working with Adaptive Routing and DCT half-handshake mode.
- An issue that ignored the default value of TX\_SCHEDULER\_BURST when its value in the ini was different than "0".
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- RDMA\_WRITE traffic performance degradation that occurred when working with DC on Adaptive Routing network.
- Bad cache invalidations of destroyed QPs.
- A rare case where asserts and ext\_synd appeared in dmesg after performing driver restart.
- An issue related to host isolation on multi-host systems.

### Enhancements

Security Hardening Enhancements: This release contains important reliability improvements and security hardening enhancements. HPE recommends upgrading your device firmware to this release to improve the firmware security and reliability of your device.

**New features and changes included in version 20.33.1048:**

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_OP will provide information on the syndrome type and which side caused the error.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Crypto features can be in either wrapped or unwrapped mode. Meaning, the key can be wrapped or in plaintext when running the CREATE\_DEK PRM command. To comply with the requirements specified in FIPS publication, all the created DEKs must be wrapped.  
This feature adds new NV\_CONFIG per device to control this mode, and enables the user to change all the crypto features to wrapped or cleartext.
- Implemented SLD detection code. Surprise Down Error Reporting Capable value was changed from 1 to 0 in boards where the downstream perst was not controlled thus causing SLD detection not to function properly.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

**Supported Devices and Features**

This software package contains the following firmware versions:

Mellanox InfiniBand Adapter	Firmware Version	PSID
HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 MCX653106A-HDAT Adapter(P31324-B21 and P31324-H21)	20.33.1048	MT_0000000594

Mellanox Firmware Package (FWPKG) for HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 OCP3 MCX653436A-HDAI Adapter : HPE part numbers P31348-B21 and P31348-H21  
Version: 20.33.1048 **(Recommended)**  
Filename: 20\_33\_1048-MCX653436A-HDA\_HPE\_Ax.pldm.fwpkg

**Important Note!**

For PLDM enabled VPI (Virtual Protocol Interconnect) adapters supporting both InfiniBand mode and Ethernet modes, every firmware version is made available in two different formats at HPE.com:

- HPE signed PLDM Firmware Package (.FWPKG filename extension) updatable via iLO.
- Firmware binary (.bin filename extension) updatable via mstflint utility from the Operating System.

Choose the appropriate firmware file format based on your preference and what suits your environment.

**ConnectX-6 VPI supports having one port as InfiniBand and the other port as Ethernet according to the following matrix of combinations.**

Port #2 - InfiniBand				
Port #1 – Ethernet	HDR/HDR100	EDR	FDR	QDR
200GbE/50GbE	supported	not supported	not supported	supported
100GbE/25GbE	supported	not supported	not supported	supported
40GbE/10GbE	supported	not supported	not supported	supported
1GbE	supported	not supported	not supported	supported

Port #2 - Ethernet				
Port #1 - InfiniBand	200GbE/50GbE	100GbE/25GbE	40GbE/10GbE	1GbE

<b>HDR / HDR100</b>	supported	supported	not supported	supported
<b>EDR</b>	supported	supported	not supported	supported
<b>FDR</b>	not supported	not supported	not supported	not supported
<b>QDR/SDR</b>	supported	supported	not supported	supported

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available

at: <https://docs.nvidia.com/networking/display/ConnectX6Firmwarev20331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

## **Prerequisites**

FWPKG will work only if the firmware version flashed on the adapter is 20.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

## **Fixes**

**The following issues have been fixed in version 20.33.1048:**

- An issue where RDMA write may experience performance degradation when working with Adaptive Routing and DCT half-handshake mode.
- An issue that ignored the default value of TX\_SCHEDULER\_BURST when its value in the ini was different than "0".
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- RDMA\_WRITE traffic performance degradation that occurred when working with DC on Adaptive Routing network.
- Bad cache invalidations of destroyed QPs.
- A rare case where asserts and ext\_synd appeared in dmesg after performing driver restart.
- An issue related to host isolation on multi-host systems.

## **Enhancements**

Security Hardening Enhancements: This release contains important reliability improvements and security hardening enhancements. HPE recommends upgrading your device firmware to this release to improve the firmware security and reliability of your device.

**New features and changes included in version 20.33.1048:**

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Crypto features can be in either wrapped or unwrapped mode. Meaning, the key can be wrapped or in plaintext when running

the CREATE\_DEK PRM command. To comply with the requirements specified in FIPS publication, all the created DEKs must be wrapped.

This feature adds new NV\_CONFIG per device to control this mode, and enables the user to change all the crypto features to wrapped or cleartext.

- Implemented SLD detection code. Surprise Down Error Reporting Capable value was changed from 1 to 0 in boards where the downstream perst was not controlled thus causing SLD detection not to function properly.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

## Supported Devices and Features

This software package contains the following firmware versions:

Mellanox InfiniBand Adapter	Firmware Version	PSID
HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 OCP3 MCX653436A-HDAI Adapter (P31348-B21 and P31348-H21)	20.33.1048	MT_0000000593

Mellanox Firmware Package (FWPKG) for HPE InfiniBand HDR100/Ethernet 100Gb 1-port QSFP56 PCIe4 x16 MCX653105A-ECAT Adapter : HPE part numbers P23665-B21 and P23665-H21

Version: 20.33.1048 (**Recommended**)

Filename: 20\_33\_1048-MCX653105A-ECA\_HPE\_Ax.pldm.fwpgk

### Important Note!

For PLDM enabled VPI (Virtual Protocol Interconnect) adapters supporting both InfiniBand mode and Ethernet modes, every firmware version is made available in two different formats at HPE.com:

1. HPE signed PLDM Firmware Package (.FWPKG filename extension) updatable via iLO.
2. Firmware binary (.bin filename extension) updatable via mstflint utility from the Operating System.

Choose the appropriate firmware file format based on your preference and what suits your environment.

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available

at: <https://docs.nvidia.com/networking/display/ConnectX6Firmwarev20331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

### Prerequisites

FWPKG will work only if the firmware version flashed on the adapter is 20.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

### Fixes

**The following issues have been fixed in version 20.33.1048:**

- An issue where RDMA write may experience performance degradation when working with Adaptive Routing and DCT half-handshake mode.
- An issue that ignored the default value of TX\_SCHEDULER\_BURST when its value in the ini was different than "0".
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- RDMA\_WRITE traffic performance degradation that occurred when working with DC on Adaptive Routing network.
- Bad cache invalidations of destroyed QPs.
- A rare case where asserts and ext\_synd appeared in dmesg after performing driver restart.
- An issue related to host isolation on multi-host systems.

### Enhancements

Security Hardening Enhancements: This release contains important reliability improvements and security hardening enhancements. HPE recommends upgrading your device firmware to this release to improve the firmware security and reliability of your device.

**New features and changes included in version 20.33.1048:**

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_ism\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Crypto features can be in either wrapped or unwrapped mode. Meaning, the key can be wrapped or in plaintext when running the CREATE\_DEK PRM command. To comply with the requirements specified in FIPS publication, all the created DEKs must be wrapped.  
This feature adds new NV\_CONFIG per device to control this mode, and enables the user to change all the crypto features to wrapped or cleartext.
- Implemented SLD detection code. Surprise Down Error Reporting Capable value was changed from 1 to 0 in boards where the downstream perst was not controlled thus causing SLD detection not to function properly.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

**Supported Devices and Features**

This software package contains the following firmware versions:

Mellanox InfiniBand Adapter	Firmware Version	PSID
HPE InfiniBand HDR100/Ethernet 100Gb 1-port QSFP56 PCIe4 x16 MCX653105A-ECAT Adapter (P23665-B21 and P23665-H21)	20.33.1048	MT_0000000452

Mellanox Firmware Package (FWPKG) for HPE InfiniBand HDR100/Ethernet 100Gb 2-port QSFP56 PCIe4 x16 MCX653106A-ECAT Adapter :  
HPE part numbers P23666-B21 and P23666-H21  
Version: 20.33.1048 **(Recommended)**  
Filename: 20\_33\_1048-MCX653106A-ECA\_HPE\_Ax.pldm.fwpkg

**Important Note!**

For PLDM enabled VPI (Virtual Protocol Interconnect) adapters supporting both InfiniBand mode and Ethernet modes, every firmware version is made available in two different formats at HPE.com:

1. HPE signed PLDM Firmware Package (.FWPKG filename extension) updatable via iLO.
2. Firmware binary (.bin filename extension) updatable via mstflint utility from the Operating System.

Choose the appropriate firmware file format based on your preference and what suits your environment.

**ConnectX-6 VPI supports having one port as InfiniBand and the other port as Ethernet according to the following matrix of combinations.**

Port #2 - InfiniBand				
Port #1 – Ethernet	HDR/HDR100	EDR	FDR	QDR
50GbE	supported	not supported	not supported	supported
100GbE/25GbE	supported	not supported	not supported	supported
40GbE/10GbE	supported	not supported	not supported	supported
1GbE	supported	not supported	not supported	supported

Port #2 - Ethernet				
Port #1 - InfiniBand	50GbE	100GbE/25GbE	40GbE/10GbE	1GbE

<b>HDR / HDR100</b>	supported	supported	not supported	supported
<b>EDR</b>	supported	supported	not supported	supported
<b>FDR</b>	not supported	not supported	not supported	not supported
<b>QDR/SDR</b>	supported	supported	not supported	supported

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available at: <https://docs.nvidia.com/networking/display/ConnectX6Firmwarev20331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

### **Prerequisites**

FWPKG will work only if the firmware version flashed on the adapter is 20.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

### **Fixes**

**The following issues have been fixed in version 20.33.1048:**

- An issue where RDMA write may experience performance degradation when working with Adaptive Routing and DCT half-handshake mode.
- An issue that ignored the default value of TX\_SCHEDULER\_BURST when its value in the ini was different than "0".
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- RDMA\_WRITE traffic performance degradation that occurred when working with DC on Adaptive Routing network.
- Bad cache invalidations of destroyed QPs.
- A rare case where asserts and ext\_synd appeared in dmesg after performing driver restart.
- An issue related to host isolation on multi-host systems.

### **Enhancements**

Security Hardening Enhancements: This release contains important reliability improvements and security hardening enhancements. HPE recommends upgrading your device firmware to this release to improve the firmware security and reliability of your device.

**New features and changes included in version 20.33.1048:**

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Crypto features can be in either wrapped or unwrapped mode. Meaning, the key can be wrapped or in plaintext when running

the CREATE\_DEK PRM command. To comply with the requirements specified in FIPS publication, all the created DEKs must be wrapped.

This feature adds new NV\_CONFIG per device to control this mode, and enables the user to change all the crypto features to wrapped or cleartext.

- Implemented SLD detection code. Surprise Down Error Reporting Capable value was changed from 1 to 0 in boards where the downstream perst was not controlled thus causing SLD detection not to function properly.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMI's, based on the delay drop status.

## Supported Devices and Features

This software package contains the following firmware versions:

Mellanox InfiniBand Adapter	Firmware Version	PSID
HPE InfiniBand HDR100/Ethernet 100Gb 2-port QSFP56 PCIe4 x16 MCX653106A-ECAT Adapter (P23666-B21 and P23666-H21)	20.33.1048	MT_0000000453

Mellanox Firmware Package (FWPKG) for Mellanox MCX623105AS-VDAT Ethernet 200Gb 1-port QSFP56 Adapter for HPE

Version: 22.33.1048 **(Recommended)**

Filename: 22\_33\_1048-MCX623105AS-VDA\_Ax.pldm.fwpkg

### Important Note!

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available

at: <https://docs.nvidia.com/networking/display/ConnectX6DxFirmwarev22331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

### Prerequisites

FWPKG will work only if the firmware version flashed on the adapter is 22.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

### Fixes

The following issues have been fixed in version 22.33.1048:

- vDPA traffic unbalance issue in active-backup VF LAG mode.
- An issue that occurred after powering off DC in Multi-Host system which resulted in OOB connection to the BMC getting lost (and fatal error appeared) due to a firmware bug in the PCIe flush flow.  
The issue was fixed by increasing the flush time and not waiting for PCIe credits to return to default values.
- A rare HW/FW timing race of serdes' power-up sequence.
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- An issue that resulted in temporary packet drops while changing PTP/FCS configuration when the links were up.
- Optimized the virtio data path to reach line speed for Tx bandwidth.
- An issue that resulted in notification indicator mistakenly being reported as FATAL thus, raising false indication.
- Bad cache invalidations of destroyed QPs.
- A configuration issue which flipped the MSB of Partition Key field in CNP packets and led to P\_KEY mismatch between CNP packets and regular packets.
- An issue that resulted in wrong port calibration due to incorrect mapping of the port during initialization stage.

### Enhancements

New features and changes included in version 22.33.1048:

- Enabled 200Gb/s out-of-the-box throughput on crypto capable devices.  
Note: If any crypto offloads is in use, 200Gb/s throughput can be achieved only after the next firmware reset
- Added support for VF migration. The hypervisor can now suspend its VF, meaning from that point the VF cannot perform action such as send/receive traffic or run any command. In this firmware version only the suspend resume mode is supported (on the same VM).

- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added Precision Time Protocol (PTP) support.  
In this version, the support includes:  
16 PTP SQs only  
only 2 ports  
only RT clock mode  
In this version, the following are not supported:  
PTP packet drop  
PTP SQ on VF  
Note: All PTP SQs must be closed before operating LFWP (life fw patch).
- Added support for HW Steering objects dump via resource dump interface.  
This support includes: STC, RTC, STE, modify argument, and modify pattern.
- Added support for VF migration.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Added support for creating a steering definer with a dword selector using create\_match\_definer\_object and the "SELECT" format.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- [Beta] Added HW Steering support for the following:  
set, add and copy inline STC action  
set and copy actions for several fields using modify\_pattern object and inline stc modify action  
FDB mode in HW steering using FDB\_RX and FDB\_TX flow table types  
ASO flow meter action via STC  
flow counter query using ASO WQE  
allocation of large bulks for the objects: STE, ASO flow meter and modify argument  
jumbo match RTC  
count action in STC
- Added support for holdover mode to comply to SyncE specifications (EEC compliance) to limit the maximum phase transient response upon link loss.
- Added support for noise filtering to comply to the SyncE specifications requirements.
- Optimized the performance of vDPA virtio including: throughput, QoS, and accuracy of min/max bandwidth when virtio works with the QoS settings.
- The new vDPA virtio-net Full Emulation capability reduces the switchover time of creating a virtq from scratch during live migration, by creating the virtq beforehand on the target server.  
When switchover happens, the pre-created virtq will be used and modified with necessary parameters.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Added support for advanced ZTR\_RTTCC algorithm based on the Programmable CC platform to achieve better congestion control without dependency on the switch ECN marking.
- DIM is used to tune moderation parameter dynamically for vDPA using an mlxreg command.  
To disable this capability, run:  
mlxreg -d /dev/mst/mt41686\_pciconf0 --reg\_id 0xc00d --reg\_len 0x8 -s "0x4.1:1=0x0"
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Added support for creating a dynamic flex parser on untrusted function, and changed the flex parser cap for untrusted function to the following:  
maximum flex parser node = 2  
maximum dw sample = 4
- Added support for SNAP1 (comm-channel) connection while running on raw ETH link.
- Crypto features can be in either wrapped or unwrapped mode. Meaning, the key can be wrapped or in plaintext when running the CREATE\_DEK PRM command.  
To comply with the requirements specified in FIPS publication, all the created DEKs must be wrapped.  
This feature adds new NV\_CONFIG per device to control this mode, and enables the user to change all the crypto features to wrapped or cleartext.
- [Beta] A new capability that enables the software to directly access ICM and write/modify the DEK objects. Such change improves the DEK object update rate by re-using DEK object instead of creating a new one.  
In addition, added the following:  
New for DEK object: bulk allocation, modify\_dek cmd, and new mode - sw\_wrapped.  
New general object INT\_KEK
- Implemented SLD detection code. Surprise Down Error Reporting Capable value was changed from 1 to 0 in boards where the downstream perst was not controlled thus causing SLD detection not to function properly.

## **Supported Devices and Features**

HPE Part Number	Mellanox Ethernet Only Adapters	PSID
P10180-B21	Mellanox MCX623105AS-VDAT Ethernet 200Gb 1-port QSFP56 Adapter for HPE	MT_0000000435

Mellanox Firmware Package (FWPKG) for Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE  
Version: 22.33.1048 (**Recommended**)  
Filename: 22\_33\_1048-MCX623106AS-CDA\_Ax.pldm.fwpkg

### **Important Note!**

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <http://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <http://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <http://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available at: <https://docs.nvidia.com/networking/display/ConnectX6DxFirmwarev22331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

### **Prerequisites**

FWPKG will work only if the firmware version flashed on the adapter is 22.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

### **Fixes**

**The following issues have been fixed in version 22.33.1048:**

- vDPA traffic unbalance issue in active-backup VF LAG mode.
- An issue that occurred after powering off DC in Multi-Host system which resulted in OOB connection to the BMC getting lost (and fatal error appeared) due to a firmware bug in the PCIe flush flow.  
The issue was fixed by increasing the flush time and not waiting for PCIe credits to return to default values.
- A rare HW/FW timing race of serdes' power-up sequence.
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- An issue that resulted in temporary packet drops while changing PTP/FCS configuration when the links were up.
- Optimized the virtio data path to reach line speed for Tx bandwidth.
- An issue that resulted in notification indicator mistakenly being reported as FATAL thus, raising false indication.
- Bad cache invalidations of destroyed QPs.
- A configuration issue which flipped the MSB of Partition Key field in CNP packets and led to P\_KEY mismatch between CNP packets and regular packets.
- An issue that resulted in wrong port calibration due to incorrect mapping of the port during initialization stage.

### **Enhancements**

**New features and changes included in version 22.33.1048:**

- Enabled 200Gb/s out-of-the-box throughput on crypto capable devices.  
Note: If any crypto offloads is in use, 200Gb/s throughput can be achieved only after the next firmware reset
- Added support for VF migration. The hypervisor can now suspend its VF, meaning from that point the VF cannot perform action such as send/receive traffic or run any command. In this firmware version only the suspend resume mode is supported (on the same VM).
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added Precision Time Protocol (PTP) support.  
In this version, the support includes:  
16 PTP SQs only  
only 2 ports  
only RT clock mode  
In this version, the following are not supported:  
PTP packet drop  
PTP SQ on VF  
Note: All PTP SQs must be closed before operating LFWP (life fw patch).
- Added support for HW Steering objects dump via resource dump interface.  
This support includes: STC, RTC, STE, modify argument, and modify pattern.

- Added support for VF migration.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Added support for creating a steering definer with a dword selector using create\_match\_definer\_object and the "SELECT" format.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- [Beta] Added HW Steering support for the following:
  - set, add and copy inline STC action
  - set and copy actions for several fields using modify\_pattern object and inline stc modify action
  - FDB mode in HW steering using FDB\_RX and FDB\_TX flow table types
  - ASO flow meter action via STC
  - flow counter query using ASO WQE
  - allocation of large bulks for the objects: STE, ASO flow meter and modify argument
  - jumbo match RTC
  - count action in STC
- Added support for holdover mode to comply to SyncE specifications (EEC compliance) to limit the maximum phase transient response upon link loss.
- Added support for noise filtering to comply to the SyncE specifications requirements.
- Optimized the performance of vDPA virtio including: throughput, QoS, and accuracy of min/max bandwidth when virtio works with the QoS settings.
- The new vDPA virtio-net Full Emulation capability reduces the switchover time of creating a virtq from scratch during live migration, by creating the virtq beforehand on the target server. When switchover happens, the pre-created virtq will be used and modified with necessary parameters.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Added support for advanced ZTR\_RTTCC algorithm based on the Programmable CC platform to achieve better congestion control without dependency on the switch ECN marking.
- DIM is used to tune moderation parameter dynamically for vDPA using an mlxreg command. To disable this capability, run:
 

```
mlxreg -d /dev/mst/mt41686_pciconf0 --reg_id 0xc00d --reg_len 0x8 -s "0x4.1:1=0x0"
```
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Added support for creating a dynamic flex parser on untrusted function, and changed the flex parser cap for untrusted function to the following:
  - maximum flex parser node = 2
  - maximum dw sample = 4
- Added support for SNAPI (comm-channel) connection while running on raw ETH link.
- Crypto features can be in either wrapped or unwrapped mode. Meaning, the key can be wrapped or in plaintext when running the CREATE\_DEK PRM command. To comply with the requirements specified in FIPS publication, all the created DEKs must be wrapped. This feature adds new NV\_CONFIG per device to control this mode, and enables the user to change all the crypto features to wrapped or cleartext.
- [Beta] A new capability that enables the software to directly access ICM and write/modify the DEK objects. Such change improves the DEK object update rate by re-using DEK object instead of creating a new one. In addition, added the following:
  - New for DEK object: bulk allocation, modify\_dek cmd, and new mode - sw\_wrapped.
  - New general object INT\_KEK
- Implemented SLD detection code. Surprise Down Error Reporting Capable value was changed from 1 to 0 in boards where the downstream perst was not controlled thus causing SLD detection not to function properly.

## Supported Devices and Features

HPE Part Number	Mellanox Ethernet Only Adapters	PSID
P25960-B21	Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE	MT_0000000437

Mellanox Firmware Package(FWPKG) for HPE Ethernet 10/25Gb 2-port SFP28 MCX512F-ACHT Adapter  
 Version: 16.33.1048 (A) **(Recommended)**  
 Filename: 16\_33\_1048-MCX512F-ACH\_Ax\_Bx.pldm.fwpkg

## Important Note!

**Disclaimer:** Certain software including drivers and documents may be available from NVIDIA. If you select a URL that directs you to <https://www.nvidia.com/>, you are then leaving HPE.com. Please follow the instructions on <https://www.nvidia.com/> to download NVIDIA software or documentation. When downloading the NVIDIA software or documentation, you may be subject to NVIDIA terms

and conditions, including licensing terms, if any, provided on its website or otherwise. HPE is not responsible for your use of any software or documents that you download from <https://www.nvidia.com/>, except that HPE may provide a limited warranty for NVIDIA software in accordance with the terms and conditions of your purchase of the HPE product or solution.

A list of known issues with this release is available at: <https://docs.nvidia.com/networking/display/ConnectX5Firmwarev16331048/Known+Issues>

Note : Flashing PLDM enabled Nvidia network adapters (Ethernet and VPI) may fail to update with FWPKG firmware component when systems have more than two adapters of the same type installed.

### **Prerequisites**

FWPKG will work only if the firmware version flashed on the adapter is 16.27.1016 or later and iLO5 firmware version must be 2.30 or higher.

### **Fixes**

**The following changes have been made in sub-version 16.33.1048(A):**

- Product rebuilt to have the new SHA 384 signature.

**The following issues have been fixed in version 16.33.1048 :**

- A rare issue that caused the QP not to receive a completion.
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- The card occasionally masked some PCIe AER reporting.
- Incorrect flow of credits blockage that prevented booting during DC cycle test.
- An issue with BMC medium migration from SMBUS to PCIe, and increased FIFOs to pass large packets in case of the migration.

### **Enhancements**

**Important :** Security Hardening Enhancements - This release contains important reliability improvements and security hardening enhancements. HPE recommends upgrading your device firmware to this version to improve the firmware security and reliability of your device.

**New features and changes included in version 16.33.1048:**

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

### **Supported Devices and Features**

HPE Part Number	Mellanox Ethernet Only Adapters	PSID
P13188-B21	HPE Ethernet 10/25Gb 2-port SFP28 MCX512F-ACHT Adapter	MT_0000000416

## Prerequisites

Use iLO5 firmware version 2.30 or higher with ConnectX4-Lx firmware version 14.32.1010. Thermal sensor reporting on the adapter will not be functional with older versions of iLO5 firmware.

## Fixes

**The following issues have been fixed in version 14.32.1010:**

- Firmware got into an unresponsive state and caused unexpected behavior when connecting an optical transceiver that support RxLOS and the remote side port was down.
- The system could not create more than 128K QPs.
- On rare occasions, the system got into an unresponsive state when a peer port went down while using an Optical module.
- Packet Pacing rate was used if asymmetric VFs was enabled.
- Incorrect RNR timeout when trying to set it during the rts2rts\_qp transition.
- Issue with RSS on IPSec flows in ConnectX-4 Lx led to performance degradation. In this scenario, the SPI optimization caused packets from a given host to hash to the same CPU core. The fix was to ignore SPI optimization according to I4\_type in ConnectX-4 Lx adapter cards.
- The GetInventory NC-SI command reported leading 0xf in firmware version when it started with 0.

## Enhancements

**Firmware for the following device has been updated to 14.32.1010:**

- P11338-B21 (HPE Ethernet 10Gb 2-port 548SFP+ Adapter)

**New features and changes included in version 14.32.1010:**

- Added 3 new assert filters (Health buffer, NVlog, FW trace). The assert will be exposed now if its severity level is equal to or above the new filter.
- Enabled Rate Limit per VM instead of VM-TC. This capability is implemented by adding support to a new Scheduling element type: rate limit elements that will connect to the rate\_limit and will share its rate limit.
- Added support for asymmetrical VFs per PF. To enable it: PF\_NUM\_OF\_VF\_VALID must be true, and PF\_NUM\_OF\_VF to a non-zero value.
- Limited the external loopback speed to the used module's capabilities.
- Improved linkup time when using the fast linkup capability.
- Added support for the slow\_restart and slow\_restart\_idle parameters to enable Zero Touch RoCE capability.

## Supported Devices and Features

HPE Part Number	Mellanox Ethernet Only Adapters	PSID
P11338-B21	HPE Ethernet 10Gb 2-port 548SFP+ Adapter	HPE0000000038

---

Online Firmware Upgrade Utility (ESXi 8.0) for HPE Mellanox Ethernet only adapters

Version: 1.0.0 (**Recommended**)

Filename: CP052064.compsig; CP052064.zip

## Important Note!

The Firmware Upgrade Utility has been split into 2 packages for Mellanox Ethernet Only NIC adapters, one supporting Synergy platforms and the other supporting ProLiant and Apollo platforms. This package supports Mellanox Ethernet Only NIC adapters on ProLiant and Apollo servers.

## Prerequisites

Use iLO5 firmware version 2.30 or higher with ConnectX4-Lx/ConnectX5 firmware version 14.32.1010/16.32.1010 (or later) respectively. Thermal sensor reporting on the adapter will not be functional with older versions of iLO5 firmware.

## Fixes

**The following issues have been fixed in version 2.42.5044 :**

- An issue that prevented the firmware from detecting a link\_down event thus preventing the IB bond interface from going to a failover mode.

**The following issues have been fixed in version 14.32.1010:**

- Firmware got into an unresponsive state and caused unexpected behavior when connecting an optical transceiver that

support RxLOS and the remote side port was down.

- The system could not create more than 128K QPs.
- On rare occasions, the system got into an unresponsive state when a peer port went down while using an Optical module.
- Packet Pacing rate was used if asymmetric VFs was enabled.
- Incorrect RNR timeout when trying to set it during the rts2rts\_qp transition.
- Issue with RSS on IPsec flows in ConnectX-4 Lx led to performance degradation. In this scenario, the SPI optimization caused packets from a given host to hash to the same CPU core. The fix was to ignore SPI optimization according to l4\_type in ConnectX-4 Lx adapter cards.
- The GetInventory NC-SI command reported leading 0xf in firmware version when it started with 0.

#### **The following issues have been fixed in version 16.33.1048:**

- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- The card occasionally masked some PCIe AER reporting.
- Incorrect flow of credits blockage that prevented booting during DC cycle test.
- An issue with BMC medium migration from SMBUS to PCIe, and increased FIFOs to pass large packets in case of the migration.

### **Enhancements**

#### **Firmware for the following devices has been updated to 2.42.5044 :**

- 779799-B21 (HPE Ethernet 10G 2-port 546FLR-SFP+ Adapter)
- 779793-B21 (HPE Ethernet 10G 2-port 546SFP+ Adapter)

#### **Firmware for the following devices has been updated to 14.32.1010:**

- 817749-B21 (HPE Ethernet 25Gb 2-port 640FLR-SFP28 Adapter)
- 817753-B21 (HPE Ethernet 25Gb 2-port 640SFP28 Adapter)

#### **Firmware for the following device has been updated to 16.33.1048:**

- 874253-B21 (HPE Ethernet 100Gb 1-port 842QSFP28 Adapter)

#### **New features and changes included in version 14.32.1010:**

- Added 3 new assert filters (Health buffer, NVlog, FW trace). The assert will be exposed now if its severity level is equal to or above the new filter.
- Enabled Rate Limit per VM instead of VM-TC. This capability is implemented by adding support to a new Scheduling element type: rate limit elements that will connect to the rate\_limit and will share its rate limit.
- Added support for asymmetrical VFs per PF. To enable it: PF\_NUM\_OF\_VF\_VALID must be true, and PF\_NUM\_OF\_VF to a non-zero value.
- Limited the external loopback speed to the used module's capabilities.
- Improved linkup time when using the fast linkup capability.
- Added support for the slow\_restart and slow\_restart\_idle parameters to enable Zero Touch RoCE capability.

#### **New features and changes included in version 16.33.1048:**

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRQ QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

### **Supported Devices and Features**



HPE Part Number	Mellanox Ethernet Only Adapters	PSID
779793-B21	HPE Ethernet 10Gb 2-port 546SFP+ Adapter	HP_1200111023
779799-B21	HPE Ethernet 10Gb 2-port 546FLR-SFP+ Adapter	HP_2240110004
817749-B21	HPE Ethernet 25Gb 2-port 640FLR-SFP28 Adapter	HP_2690110034
817753-B21	HPE Ethernet 25Gb 2-port 640SFP28 Adapter	HP_2420110034
874253-B21	HPE Ethernet 100Gb 1-port 842QSFP28 Adapter	HPE0000000014

Online Firmware Upgrade Utility (ESXi 8.0) for HPE Mellanox VPI (Ethernet and Infiniband mode) ConnectX4 and ConnectX5 devices on VMware ESXi 8.0

Version: 1.0.0 (**Recommended**)

Filename: CP051774.compsig; CP051774.zip

## Fixes

### **Fixes in version 12.28.2006:**

- Fixes an issue that caused the DCR to be destroyed before the retry option managed to work when the retry timeout is too big. In this case the DCR' time-to-live was increased, and the maximum retry timeout was decreased.
- Increased PHY power consumption limit to 1.5w.
- Fixed an issue that caused PortCounters.PortRcvErr / PPCNT.infiniband\_counters.PortRcvErr not to report port icrc errors.

### **The following issues have been fixed in version 16.33.1048**

- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- The card occasionally masked some PCIe AER reporting.
- Incorrect flow of credits blockage that prevented booting during DC cycle test.
- An issue with BMC medium migration from SMBUS to PCIe, and increased FIFOs to pass large packets in case of the migration.

## Enhancements

### **Firmware for the following devices has been updated to 12.28.2006:**

- 825110-B21 (HPE InfiniBand EDR/Ethernet 100Gb 1-port 840QSFP28 Adapter)
- 825111-B21 (HPE InfiniBand EDR/Ethernet 100Gb 2-port 840QSFP28 Adapter)

### **Firmware for the following devices has been updated to 16.33.1048:**

- 879482-B21 (HPE InfiniBand FDR/Ethernet 40/50Gb 2-port 547FLR-QSFP Adapter)
- 872726-B21 (HPE InfiniBand EDR/Ethernet 100Gb 2-port 841QSFP28 Adapter)

### **New Features and changes included in version 12.28.2006:**

- Increased the maximum XRQ number to 512.

### **New features and changes included in version 16.33.1048:**

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRQ QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

## Supported Devices and Features

HPE Part Number	Device Name	PSID
825110-B21	HPE InfiniBand EDR/Ethernet 100Gb 1-port 840QSFP28 Adapter	HP_2180110032
825111-B21	HPE InfiniBand EDR/Ethernet 100Gb 2-port 840QSFP28 Adapter	HP_2190110032
872726-B21	HPE InfiniBand EDR/Ethernet 100Gb 2-port 841QSFP28 Adapter	HPE0000000009
879482-B21	HPE InfiniBand FDR/Ethernet 40/50Gb 2-port 547FLR-QSFP Adapter	HPE0000000022

Online Firmware Upgrade Utility (ESXi 8.0) for HPE Mellanox VPI (Ethernet and Infiniband mode) ConnectX6 devices on VMware ESXi 8.0

Version: 1.0.0 (**Recommended**)

Filename: CP052055.compsig; CP052055.zip

### Important Note!

**ConnectX-6 VPI supports having one port as InfiniBand and the other port as Ethernet according to the following matrix of combinations.**

Port #2 - InfiniBand				
Port #1 - Ethernet	HDR/HDR100	EDR	FDR	QDR
200GbE/50GbE	supported	not supported	not supported	supported
100GbE/25GbE	supported	not supported	not supported	supported
40GbE/10GbE	supported	not supported	not supported	supported
1GbE	supported	not supported	not supported	supported

Port #2 - Ethernet				
Port #1 - InfiniBand	200GbE/50GbE	100GbE/25GbE	40GbE/10GbE	1GbE
HDR / HDR100	supported	supported	not supported	supported
EDR	supported	supported	not supported	supported
FDR	not supported	not supported	not supported	not supported
QDR/SDR	supported	supported	not supported	supported

### Prerequisites

Use iLO5 firmware version 2.30 or higher with ConnectX6 firmware version 20.32.1010. Thermal sensor reporting on the adapter will not be functional with older versions of iLO5 firmware.

### Fixes

**The following issues have been fixed in version 20.33.1048:**

- An issue where RDMA write may experience performance degradation when working with Adaptive Routing and DCT half-handshake mode.
- An issue that ignored the default value of TX\_SCHEDULER\_BURST when its value in the ini was different than "0".
- Implemented SLD detection code. Surprise Down Error Reporting Capable value was changed from 1 to 0 in boards where the downstream perst was not controlled thus causing SLD detection not to function properly.
- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- RDMA\_WRITE traffic performance degradation that occurred when working with DC on Adaptive Routing network.
- Bad cache invalidations of destroyed QPs.
- A rare case where asserts and ext\_synd appeared in dmesg after performing driver restart.
- An issue related to host isolation on multi-host systems.

### Enhancements

**Firmware for the following devices has been updated to 20.33.1048:**

- HPE InfiniBand HDR/Ethernet 200Gb 1-port QSFP56 PCIe3 x16 MCX653105A-HDAT Adapter - P06154-B21
- HPE InfiniBand HDR100/Ethernet 100Gb 1-port QSFP56 PCIe3 x16 MCX653105A-ECAT Adapter - P06250-B21
- HPE InfiniBand HDR100/Ethernet 100Gb 2-port QSFP56 PCIe3 x16 MCX653106A-ECAT Adapter - P06251-B21

**New features and changes included in version 20.33.1048:**

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.

- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRO QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Crypto features can be in either wrapped or unwrapped mode. Meaning, the key can be wrapped or in plaintext when running the CREATE\_DEK PRM command. To comply with the requirements specified in FIPS publication, all the created DEKs must be wrapped.  
This feature adds new NV\_CONFIG per device to control this mode, and enables the user to change all the crypto features to wrapped or cleartext.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

## Supported Devices and Features

HPE Part Number	Device Name	PSID
P06154-B21	HPE InfiniBand HDR/Ethernet 200Gb 1-port QSFP56 PCIe3 x16 MCX653105A-HDAT Adapter	HPE0000000034
P06250-B21	HPE InfiniBand HDR100/Ethernet 100Gb 1-port QSFP56 PCIe3 x16 MCX653105A-ECAT Adapter	HPE0000000035
P06251-B21	HPE InfiniBand HDR100/Ethernet 100Gb 2-port QSFP56 PCIe3 x16 MCX653106A-ECAT Adapter	HPE0000000036

Online Firmware Upgrade Utility (ESXi 8.0) for HPE Mellanox VPI (Ethernet and Infiniband mode) devices on VMware ESXi 8.0  
Version: 1.0.0 (**Recommended**)  
Filename: CP052070.compsig; CP052070.zip

## Important Note!

### Known Issues in firmware 2.42.5000, 2.42.5056, 2.42.5700:

- When using the Quad Small Form-factor Pluggable (QSFP) module RTX320-581, and performing a driver restart for the firmware upgrade/downgrade to take effect, the link does not come up.  
**Workaround:** Reboot the server.
- Enabling/disabling cq\_timestamp using mlxconfig is not supported.
- In a card with 2 separate LEDs scheme (a Phy LED and a logic LED) only the Phy LED will lit. Meaning, the orange LED will not be active while the ETH link is in an idle mode.
- In SR-IOV setup, using mlxconfig when the Packet Filter (PF) is passed through to a VM requires a reboot of the Hypervisor.
- Downgrading from v2.30.8000 or later to an earlier version than 2.30.8000 requires server reboot.  
**Workaround:** Reboot the server.
- On ConnectX-3 Ethernet adapter cards, there is a mismatch between the GUID value returned by firmware management tools and that returned by fabric/ driver utilities that read the GUID via device firmware (e.g., using ibstat). Mlxburn/flint return 0xffff as GUID while the utilities return a value derived from the MAC address. For all driver/firmware/software purposes, the latter value should be used.  
**Workaround:** Please use the GUID value returned by the fabric/driver utilities (not 0xffff).
- SBR should be asserted for a minimum of 50 milliseconds for the ConnectX-3 adapters.
- On Pilot1 SL230, PCIe link occasionally does not come up at Gen3 speed.
- RHEL6.3 Inbox driver causes kernel panic when SRIOV is enabled on VPI cards due to driver compatibility issue.
- **Workaround:** Set the "do\_ sense=false" parameter in the [IB\_TAB] i.
- In advanced steering mode, side band management connectivity may be lost when having more than 8 QP per mcg.
- When SR-IOV is disabled in the system BIOS, a PCI issue is noticed in Ubuntu v12.04.3 with Linux kernel v3.8 which affects NICs of several manufacturers including Mellanox's, preventing them from operating.  
**Workaround:** Enable SR-IOV in the BIOS.
- Mellanox Firmware Tools (MFT) might leave the flash semaphore locked if the tool operation is forced stopped. The locked semaphore prevents the firmware from accessing the flash and causes firmware hang.  
**Workaround:** Clear the semaphore using MFT command: 'flint -clear\_semaphore'
- Cable Info MAD reports a wrong cable info when using the MC2210411-SR4 module.
- Gen2 failure at temperature sweep up to 10C/min (for MT27518A1-FDIR-BV only)..
- PCIe Gen2 link unstable at temperature sweep of 10C/min for MT27518A1-FDIR-BV.
- Bloom filter is currently not supported.
- When downgrading from firmware v2.11.0000 and using MFT 3.0.0-3, the following message is displayed due to the mlxconfig tool: You are trying to override configurable FW by non-configurable FW. If you continue, old FW configurations will

be cleared, do you want to continue ? (y/n) [n] : y You are trying to restore default configuration, do you want to continue ? (y/n) [n] : y.

- DMFS should not be enabled when working with InfiniBand on MLNX\_OFED-2.0.3
- ConnectX®-3 Pro VF device ID is presented the same as ConnectX®-3 VF device ID due to driver limitations.  
**Workaround:** Use the physical function device ID to identify the device.
- Virtual Product Data (VPD) read-only fields are writable.  
**Workaround:** Do not write to read-only fields if you wish to preserve them.
- When working in Virtual Path Identifier (VPI) mode with port1 FDR and port2 40G, error counters misbehave and increase rapidly.
- Setting the device to 128Byte CQ/EQ stride will cause misbehavior of sideband management resulting in communication loss.
- CQ and EQ cannot be configured to different stride sizes.
- Changing port protocol from ETH to IB on port with NCSI/IPMI enabled while the port is connected to ETH switch is not supported.  
**Workaround:** 1. Unplug the cable from the switch 2. Restart driver 3. Change the protocol via the appropriate tools.
- Adapter card MCX349A-XCCN may experience longer linkup times of a few seconds with specific switches.
- Adapter card MCX349A-XCCN does not respond to ethtool "identify" command (ethtool -p/--identify).
- Remote Desktop Protocol (RDP) over IPv6 is currently not functional.  
**Workaround:** Set the default RoCE mode in the software to RoCE v2 (also when not using RoCE)
- Sniffer QP cannot be removed from the regular rule after adding the QP with insertion scheme equals to "push to that rule".
- Since only a single Boot Entry Vector (BEV) per PCI Physical Function is supported, disabling the first port causes the second port to disappear as well.
- The NIC does not notify the driver of a link-down incident when a cable is unplugged from a NIC port with 56GbE port link.
- 56GbE link is not raised when using 100GbE optic cables.
- When working with MLNX\_OFED v3.3-1.0.0.0, server reboot could get stuck due to a kernel panic in mlx4\_en\_get\_drvinfo() that is called from asynchronous event handler.
- When running ibdump, loopback traffic is mirroring into the kernel driver.
- MAC address that are set from the OS using ifconfig are not reflected in the OCBB buffer.
- The adapter card cannot raise a 10G link vs. a 40GE capable switch port in C7000 enclosure. It can raise a 1G Link and only if the switch port allows it.
- MTUSB communication via I2C header on primary I2C bus is supported only in live-fish mode.

## Fixes

### **Fixes in version 2.42.5000:**

- PortRcvPkts counter was prevented from being cleared after resetting it.
- The system Timed Out on the configuration cycle of the Virtual Functions (VFs) when more than 10 Virtual Functions performed FLR and the completion Time Out value was configured to a range of less than 16 msec.
- The server hangs and results in NMI when running "mlxftop -d mt4103\_pci\_cr0" while restarting the driver in parallel (from a different thread). In this case, the downstream bridge over the device reported completion timeout error.
- In flow\_steering, BMC could not receive a ping over IPV6 after running bmc\_reboot.
- While closing the HCA, the RX packet caused bad access to resources that did not exist, and consequently caused the QPCGW or the irisc to get stuck.
- The master SMLID and the LID was either 0 or 0xFFFF when the port was neither active nor armed.
- ibdump could not capture all MADs packets.
- link did not go up after reboot.
- Fixed a rare issue that cause the PCIe configuration cycle that arrived during the time of sw\_reset to generate 2 completions.
- Network Controller Sideband Interface (NC-SI) did not work when adding the disable\_static\_steering\_ini field in the ini file, due to memory allocation issue for this field in the scratchpad.

### **Fixes in version 2.42.5056:**

- Fixed an issue that resulted in reading from invalid I/O address on handover from UEFI boot to OS boot, when a port was configured as InfiniBand on a VPI adapter device.

## Enhancements

### **Firmware for the following devices are updated to 2.42.5000:**

764282-B21  
764286-B21

### **Firmware for the following devices are updated to 2.42.5056:**

764283-B21  
764284-B21

### **Firmware for the following device is updated to 2.42.5700:**

764285-B21

### **New features in firmware version 2.42.5000:**

- Added support for the following features.
  - new TLV: CX3\_GLOBAL\_CONF to enable/disable timestamp on incoming packets through mlxconfig configuration.
  - User MAC configuration.
  - Automatically collecting mstdump before driver reset.
  - A mechanism to detect DEAD\_IRISC (plastic) from TPT (iron) and raise an assert.

- A new field is added to "set port" command which notifies the firmware what is the user\_mtu size.
- Improved the debug ability for command timeout cases.

#### **New features and changes in firmware version 2.42.5700.**

- Modified the mlx\_cmd\_get\_mlx\_link\_status command return value to return "Link Type = Ethernet" in Ethernet adapter cards.

### **Supported Devices and Features**

#### **Supported Devices:**

HPE Part Number	Device Name	PSID
764282-B21	HPE InfiniBand QDR/Ethernet 10Gb 2-port 544+M Adapter	HPE_1350110023
764283-B21	HPE InfiniBand FDR/Ethernet 10Gb/40Gb 2-port 544+M Adapter	HPE_1360110017
764284-B21	HPE InfiniBand FDR/Ethernet 10Gb/40Gb 2-port 544+QSFP Adapter	HPE_1370110017
764285-B21	HPE InfiniBand FDR/Ethernet 10Gb/40Gb 2-port 544+FLR-QSFP Adapter	HPE_1380110017
764286-B21	HPE InfiniBand QDR/Ethernet 10Gb 2-port 544+FLR-QSFP Adapter	HPE_1390110023

Online Firmware Upgrade Utility (ESXi 8.0) for Mellanox Open Ethernet cards

Version: 1.0.0 (**Recommended**)

Filename: CP052069.compsig; CP052069.zip

#### **Important Note!**

On Adapter Firmware rewrite scenario, SUM will always discover the Mellanox Open adapter firmware smart component as applicable and select it for deployment If the server iLO5 firmware version is older than 2.30.

#### **Prerequisites**

Use iLO5 firmware version 2.30 or higher with ConnectX4-Lx/ConnectX5 firmware version 14.32.1010/16.32.1010 respectively. Thermal sensor reporting on the adapter will not be functional with older versions of iLO5 firmware.

#### **Fixes**

##### **The following issues have been fixed in version 14.32.1010:**

- Firmware got into an unresponsive state and caused unexpected behavior when connecting an optical transceiver that support RxLOS and the remote side port was down.
- The system could not create more than 128K QPs.
- On rare occasions, the system got into an unresponsive state when a peer port went down while using an Optical module.
- Packet Pacing rate was used if asymmetric VFs was enabled.
- Incorrect RNR timeout when trying to set it during the rts2rts\_qp transition.
- Issue with RSS on IPsec flows in ConnectX-4 Lx led to performance degradation. In this scenario, the SPI optimization caused packets from a given host to hash to the same CPU core. The fix was to ignore SPI optimization according to I4\_type in ConnectX-4 Lx adapter cards.
- The GetInventory NC-SI command reported leading 0xf in firmware version when it started with 0.

##### **The following issues have been fixed in version 16.33.1048:**

- Rare lanes skew issue that caused CPU to timeout in Rec.idle.
- The card occasionally masked some PCIe AER reporting.
- Incorrect flow of credits blockage that prevented booting during DC cycle test.
- An issue with BMC medium migration from SMBUS to PCIe, and increased FIFOs to pass large packets in case of the migration.

#### **Enhancements**

##### **Firmware for the following devices has been updated to 14.32.1010:**

- P21930-B21 (HPE Ethernet 10Gb 2-port SFP+ MCX4121A-XCAT Adapter)
- P11341-B21 (HPE Ethernet 10Gb 2-port SFP+ MCX4621A-ACAB OCP3 Adapter)

##### **Firmware for the following device has been updated to 16.33.1048:**

- P21927-B21 (HPE Ethernet 100Gb 2-Port QSFP28 MCX516A-CCHT Adapter)

### New features and changes included in version 14.32.1010:

- Added 3 new assert filters (Health buffer, NVlog, FW trace). The assert will be exposed now if its severity level is equal to or above the new filter.
- Enabled Rate Limit per VM instead of VM-TC. This capability is implemented by adding support to a new Scheduling element type: rate limit elements that will connect to the rate\_limit and will share its rate limit.
- Added support for asymmetrical VFs per PF. To enable it: PF\_NUM\_OF\_VF\_VALID must be true, and PF\_NUM\_OF\_VF to a non-zero value.
- Limited the external loopback speed to the used module's capabilities.
- Improved linkup time when using the fast linkup capability.
- Added support for the slow\_restart and slow\_restart\_idle parameters to enable Zero Touch RoCE capability.

### New features and changes included in version 16.33.1048:

- Added pci\_rescan\_needed field to the MFRL access register to indicate whether a PCI rescan is needed based on the NV configurations issued by the software.  
Note: If the Keep Link Up NV configuration is changed, phyless reset will be blocked.
- Added a new MAD of class SMP that has the attributes hierarchy\_Info as defined in the IB Specification and is used to query the hierarchy information stored on the node and the physical port.
- Added a new register (vhca\_icm\_ctrl access\_reg) to enable querying and limiting the ICM pages in use.
- Enhanced the XRQ QP error information provided to the user in case QP goes into an error state. In such case, QUERY\_QP will provide information on the syndrome type and which side caused the error.
- Replaced the deprecated NetworkPort schema with Port schema in NIC RDE implementation.
- Updated the ibstat status reported when the phy link is down. Now QUERY\_VPORT\_STATE.max\_tx\_speed of UPLINK will not be reported as 0 anymore.
- Disabled the option to send SMPs from unauthorized hosts.
- Enabled the option to modify the ip\_ecn field in the packet header in firmware steering.
- Modified the TX or RX cache invalidation behavior. TX or RX cache invalidation now does not occur automatically but only when the software performs the sync operation using the using sync\_steering command.
- Modified the maximum bulk size per single allocation from "log\_table\_size - log\_num\_unisizes", to allocate any range size, to remove limitations that HWS objects such as counters and modify arguments might encounter.
- Enabled delay drop for hairpin packets. If a hairpin QP is created with delay\_drop\_en enabled, the feature will be enabled across all GVMIs, based on the delay drop status.

### Supported Devices and Features

HPE Part Number	Mellanox Ethernet Only Adapters	PSID
P21930-B21	HPE Ethernet 10Gb 2-port SFP+ MCX4121A-XCHT Adapter	MT_0000000414
P11341-B21	HPE Ethernet 10Gb/25Gb 2-port SFP28 MCX4621A-ACAB OCP3 Adapter	MT_0000000238
P21927-B21	HPE Ethernet 100Gb 2-port QSFP28 MCX516A-CCHT Adapter	MT_0000000417

### Firmware - NVDIMM

Firmware package for HPE Persistent Memory featuring Intel Optane DC Persistent Memory on HPE Gen10 Plus Servers  
Version: 02.02.00.1553 (D) **(Recommended)**  
Filename: dcpmm\_02.02.00.1553.fwpkg

[Top](#)

### Important Note!

This software package contains Intel Optane DC Persistent Memory Firmware version 2.2.0.1553

### Fixes

This product corrects an issue that three different capacities of Intel Optane DC Persistent Memory are identifiable with three individual device GUID.

### Enhancements

- Add ESXi 8.0 support

### Supported Devices and Features

This package supports the following Memory Devices:

- HPE 512GB 3200 Persistent Memory Kit featuring Intel Optane DC Persistent Memory
- HPE 256GB 3200 Persistent Memory Kit featuring Intel Optane DC Persistent Memory

Firmware package for HPE Persistent Memory featuring Intel Optane DC Persistent Memory on HPE Gen10 Servers  
Version: 01.02.00.5446 (C) **(Recommended)**  
Filename: dcpmm\_01.02.00.5446.fwpkg

**Important Note!**

This software package contains Intel Optane DC Persistent Memory Firmware version 1.2.0.5446

**Enhancements**

- Add ESXi 8.0 Support

**Supported Devices and Features**

This package supports the following Memory Devices:

- HPE 512GB 2666 Persistent Memory Kit featuring Intel Optane DC Persistent Memory
- HPE 256GB 2666 Persistent Memory Kit featuring Intel Optane DC Persistent Memory
- HPE 128GB 2666 Persistent Memory Kit featuring Intel Optane DC Persistent Memory

---

**Firmware - Storage Controller**

[Top](#)

Firmware Package - HPE MR216i-a Gen10 Plus Tri Mode Controller  
Version: 52.16.3-4455 (B) **(Recommended)**  
Filename: HPE\_MR216i-a\_Gen10P\_52.16.3-4455.fwpkg

**Important Note!**

**This firmware version to be used on HPE MR216i-a Gen10 Plus Controller.**

**Enhancements**

Support UBM6 and rebuild with SHA384

---

Firmware Package - HPE MR216i-p Gen10 Plus Tri Mode Controller with Gen10 and Gen10 Plus servers  
Version: 52.16.3-4455 (B) **(Recommended)**  
Filename: HPE\_MR216i-p\_Gen10P\_52.16.3-4455.fwpkg

**Important Note!**

**This firmware version to be used on HPE MR216i-p Gen10 Plus Controller.**

**Enhancements**

Support UBM6 and rebuild with SHA384

---

Firmware Package - HPE MR416i-a Gen10 Plus Tri Mode Controller  
Version: 52.16.3-4455 (B) **(Recommended)**  
Filename: HPE\_MR416i-a\_Gen10P\_52.16.3-4455.fwpkg

**Important Note!**

**This firmware version to be used on HPE MR416i-a Gen10 Plus Controller.**

**Enhancements**

Support UBM6

---

Firmware Package - HPE MR416i-p Gen10 Plus Tri Mode Controller with Gen10 and Gen10 Plus servers  
Version: 52.16.3-4455 (B) **(Recommended)**  
Filename: HPE\_MR416i-p\_Gen10P\_52.16.3-4455.fwpkg

### **Important Note!**

**This firmware version to be used on HPE MR416i-p Gen10 Plus Controller.**

### **Enhancements**

Support UBM6 and rebuild with SHA384

---

Firmware Package - HPE SR932i-p Gen10 Plus /SR416i-a Gen10 Plus Controllers  
Version: 03.01.14.062 **(Recommended)**  
Filename: HPE\_SR416\_SR932\_Gen10P\_03.01.14.062.fwpkg

### **Fixes**

- Fixed a performance issue with a RAID 5 volume writing sequential IOs that resulted in IOs taking a long time and LUN resets.
- Fixed an issue where UBM3 backplanes fail to flash.
- Fixed a Predictive Failure drive LED blinking issue.
- Fixed an issue where a re-enabled encrypted single drive RAID 0 volume was reported with state Offline after system reboot.
- Fixed long SATA SSD TRIM causing hang.
- Fixed an issue where events are sent continuously if the host does not respond to PlatformEventMessage.
- Fixed an issue where Redfish Volume Create fails when using 4Kn data drives.
- Fixed an issue with issuing back to back SCSI UNMAP commands to NVMe drives.
- Updated Redfish Drive.Identifiers.DurableName to conform to the standard.
- Updated Redfish Volume.Identifiers.DurableName to conform to the standard.
- Updated Redfish to the 2021.4 schema bundle.

### **Enhancements**

- Added UBM6 backplane support.
- Improved single RAID volume performance for multiple streams of sequential IO.
- Added support to send SCSI timestamp to drives to sync host system time (RTC).
- Added a new HII menu that will attempt to re-enable a previously failed volume whose physical drives are back online.
- Added an Unlock Controller option in the HII menu when controller password is set for Controller Based Encryption (CBE).
- Added a new HII menu option to setup and configure Controller Based Encryption (CBE).
- Added support for the following Redfish ACTION requests:
  - Drive.SecureErase
  - Drive.Reset
  - Storage.ResetToDefaults
- Added support for Redfish PATCH requests for the following properties:
  - Volume.DisplayName
  - Volume.Links.DedicatedSpareDrives
  - Volume.IOPerfModeEnabled
  - Volume.ReadCachePolicy
  - Volume.WriteCachePolicy
  - Drive.LocationIndicatorActive
  - Drive.WriteCacheEnabled
  
  - StorageController.ControllerRates.ConsistencyCheckRatePercent
  - StorageController.ControllerRates.RebuildRatePercent
  - StorageController.ControllerRates.TransformationRatePercent
- Added the following Redfish alerts:
  - DriveOffline
  - DriveMissing
  - DriveOfflineCleared
  - VolumeOffline
  - VolumeOfflineCleared
  - BatteryMissing
  - BatteryFailure
  - BatteryCharging
  
  - BatteryOK
  - ControllerDegraded
  - ControllerFailure
  - ControllerPreviousFailure
  - ControllerPasswordRequired

– ControllerPasswordEntered (changing to ControllerPasswordAccepted in the future)

- Added MaxMembers to Redfish VolumeCollection
- Added 'NativeDriveEncryption' (SED) to Redfish Volume.EncryptionTypes.
- Improved the consistency of I/O latency by caching iLO inventory commands

---

HPE D3600/D3700/D3610/D3710 12Gb SAS Disk Enclosure ROM Flash Component for VMware (ESXi)

Version: 5.04 (E) (**Recommended**)

Filename: CP048111.compsig; CP048111.md5; CP048111.zip

### **Important Note!**

**IMPORTANT:** Firmware updates must be performed during a system maintenance window, with all I/O to the system halted. In single domain configuration, if user hosts an OS in D3000(or any storage box) and flash the SEPs, it will hang/crash everytime as SmartComponent will reset the SEPs after flash/codeload.

**WARNING!** Do not power cycle or restart during the firmware update as this can result in loss of capabilities for this unit. It typically takes several minutes for the firmware to load.

**NOTE:** All firmware flash progress messages are logged to /var/cpq/D3000.log and flash summary is logged to /var/cpq/Component.log.

### **Prerequisites**

**IMPORTANT:** Firmware updates must be performed during a system maintenance window, with all I/O to the system halted.

**WARNING!** Do not power cycle or restart during the firmware update as this can result in loss of capabilities for this unit. It typically takes several minutes for the firmware to load.

**NOTE:** All firmware flash progress messages are logged to /var/cpq/D3000.log and flash summary is logged to /var/cpq/Component.log.

### **Fixes**

**The following fixes were incorporated in this version:**

- The Enabled-ClusterS2D command now completes successfully when executed on a SATA drive within a D3610 disk enclosure for a NonStop solution.
- The smart carrier, which is the drive case for SAS drives, now authenticates in the D3610/D3710 drive enclosure.
- Added new 7-segment error codes E0 and E1 to report issues with Fan modules A and B, respectively. These new codes only apply to the D3610/D3710 and only display when running firmware 5.04.
- If the storage enclosure processor within the I/O module fails, a hard reset (power down and then power up) is executed to ensure the processor comes back online.

Please refer to the [Release Notes](#) for the complete listing of fixes, enhancements, known issues and work-arounds corresponding to this firmware.

### **Supported Devices and Features**

The D3600 / D3700 / D3610 / D3710 Enclosure can be attached to any of the following HPE Storage Controllers and Host Bus Adapters :

- Smart Array P841 Controller
- Smart Array P441 Controller
- Smart HBA H241
- Smart Array P741m Controller
- Smart Array P408e-p Controller
- Smart Array E208e-p Controller
- Smart Array P408e-m Controller

---

HPE D6020 12Gb SAS Disk Enclosure ROM Flash Component for VMware (ESXi)

Version: 2.74 (L) (**Recommended**)

Filename: CP052324.compsig; CP052324.md5; CP052324.zip

### **Important Note!**

**IMPORTANT:** Firmware updates must be performed during a system maintenance window, with all I/O to the system halted. In single domain configuration, if user hosts an OS in D6020(or any storage box) and flash the SEPs, it will hang/crash everytime as

SmartComponent will reset the SEPs after flash/codeload.

**WARNING!** Do not power cycle or restart during the firmware update as this can result in loss of capabilities for this unit. It typically takes several minutes for the firmware to load.

**NOTE:** All firmware flash progress messages are logged to /var/cpq/D6020.log and flash summary is logged to /var/cpq/Component.log.

### **Prerequisites**

**IMPORTANT:** Firmware updates must be performed during a system maintenance window, with all I/O to the system halted.

**WARNING!** Do not power cycle or restart during the firmware update as this can result in loss of capabilities for this unit. It typically takes several minutes for the firmware to load.

**NOTE:** All firmware flash progress messages are logged to /var/cpq/D6020.log and flash summary is logged to /var/cpq/Component.log.

### **Fixes**

**The following fixes were incorporated in this version:**

- Temperature sensors logic inside gSEP model and SES database
- When an IOM is pulled the surviving IOM reports false critical temperatures

Please refer to the [Release Notes](#) for the complete listing of fixes, enhancements, known issues and work-arounds corresponding to this firmware.

### **Supported Devices and Features**

The D6020 Enclosure can be attached to any of the following HPE Storage Controllers and Host Bus Adapters :

- Smart Array P741m Controller
- Smart Array P408e-p Controller
- Smart Array E208e-p Controller
- Smart Array P408e-m Controller

---

HPE D8000 12Gb SAS Disk Enclosure ROM Flash Component for VMware (ESXi)

Version: 0130 (**Recommended**)

Filename: CP051586.compsig; CP051586.md5; CP051586.zip

### **Important Note!**

**IMPORTANT:** Firmware updates must be performed during a system maintenance window, with all I/O to the system halted. In single domain configuration, if user hosts an OS in D8000(or any storage box) and flash the SEPs, it will hang/crash everytime as SmartComponent will reset the SEPs after flash/codeload.

**WARNING!** Do not power cycle or restart during the firmware update as this can result in loss of capabilities for this unit. It typically takes several minutes for the firmware to load.

**NOTE:** All firmware flash progress messages are logged to /var/cpq/D8000.log and flash summary is logged to /var/cpq/Component.log.

### **Prerequisites**

**IMPORTANT:** Firmware updates must be performed during a system maintenance window, with all I/O to the system halted.

**WARNING!** Do not power cycle or restart during the firmware update as this can result in loss of capabilities for this unit. It typically takes several minutes for the firmware to load.

**NOTE:** All firmware flash progress messages are logged to /var/cpq/D8000.log and flash summary is logged to /var/cpq/Component.log.

### **Fixes**

**The following fixes were incorporated in this version:**

- Smart Array reporting "Storage Enclosure FW upgrade Problem Detected".
- PSU version is showing as "No Ver" after inserted the New bel power PSU.
- GEM\_5\_2 Coverity defect fix.

Please refer to the [Release Notes](#) for the complete listing of fixes, enhancements, known issues and work-arounds corresponding to this firmware.

## **Supported Devices and Features**

The D8000 Enclosure can be attached to any of the following HPE Storage Controllers and Host Bus Adapters :

- HPE Smart Array P408e-p Controller
- HPE Smart Array E208e-p Controller

---

Online ROM Flash Component for VMware ESXi - HPE Smart Array P408i-p, P408e-p, P408i-a, E208i-p, E208e-p, E208i-a, P816i-a SR Gen10

Version: 5.32 (B) **(Recommended)**

Filename: CP053855.compsig; CP053855.zip

## **Fixes**

- Fixed an issue where the encrypted data is not accessible for a RAID 50/60 volume when it was failed and healed using "Heal Array".
- Fixed an issue where pointers in PLDM commands were accessed before initialized.
- Fixed an issue where UBM3 backplanes fail to flash.
- Fixed a Predictive Failure drive LED blinking issue.
- Fixed an issue where a re-enabled encrypted single drive RAID 0 logical drive was reported with state Offline after system reboot.
- Fixed long SATA SSD TRIM causing hang.
- Fixed an issue where the product ID of an enclosure was not showing correctly.
- Fixed an issue where the drive bay number for a failed drive is wrong.
- Fixed an issue where the Real Time Clock (RTC) timestamp was not sent to the SES based storage enclosure SEPs attached to internal connectors of the controller.
- Fixed an issue where events are sent continuously if the host does not respond to PlatformEventMessage.
- Fixed an issue where Redfish Volume Create fails when using 4Kn data drives.
- Update Redfish Drive.Identifiers.DurableName to conform to the standard.
- Update Redfish Volume.Identifiers.DurableName to conform to the standard.
- Updated Redfish to the 2021.4 schema bundle.

## **Enhancements**

- Added UBM6 backplane support.
- Added support for SED Local Key Management.
- Added a new HII menu that will attempt to re-enable a previously failed volume whose physical drives are back.
- Added an Unlock Controller option in the HII menu when controller password is set for Controller Based Encryption (CBE).
- Added new HII menu options to configure Controller Based Encryption (CBE).
- Added new HII options to enable and configure SED Local Key Management.
- Added support for the following Redfish ACTION requests:
  - Drive.SecureErase
  - Drive.Reset
  - Storage.ResetToDefaults
- Added support for Redfish PATCH requests for the following properties:
  - Volume.DisplayName
  - Volume.Links.DedicatedSpareDrives
  - Volume.IOPerfModeEnabled
  - Volume.ReadCachePolicy
  - Volume.WriteCachePolicy
  - Drive.LocationIndicatorActive
  - Drive.WriteCacheEnabled
  
  - StorageController.ControllerRates.ConsistencyCheckRatePercent
  - StorageController.ControllerRates.RebuildRatePercent
  - StorageController.ControllerRates.TransformationRatePercent
- Added the following Redfish alerts:
  - DriveOffline
  - DriveMissing
  - DriveOfflineCleared
  - VolumeOffline
  - VolumeOfflineCleared
  - BatteryMissing
  - BatteryFailure
  - BatteryCharging
  
  - BatteryOK
  - ControllerDegraded
  - ControllerFailure
  - ControllerPreviousFailure
  - ControllerPasswordRequired

– ControllerPasswordEntered (changing to ControllerPasswordAccepted in the future)

- Added MaxMembers to Redfish VolumeCollection
- Added 'Reverting' to Redfish Drive.Operations.OperationName used for SED.
- Added 'NativeDriveEncryption' (SED) to Redfish Volume.EncryptionTypes.
- Added support for Redfish Drive.EncryptionStatus for SED.
- Redfish Drive.Status.State will be set to StandbyOffline in the following conditions:
  - SED is Foreign
  - SED is Locked (only for controller owned SEDs)
  - SED is controller owned and controller is waiting on SED adapter password
- Added support for Redfish Volume.Encrypted for SED.

---

Universal Firmware Package for HPE Gen10 Plus Boot Controller NS204i-p, NS204i-d, NS204i-t, NS204i-r  
Version: 1.0.14.1063 (C) **(Recommended)**  
Filename: HPE\_NS204i\_Gen10P\_1.0.14.1063.fwpkg

### **Important Note!**

This firmware version is to be used on NS204i controllers.

Use iLO to flash HPE\_NS204i\_Gen10P\_PLDM\_xxxx.fwpkg above 1.0.14.1055.; continuously HPE offers PLDM Type5 FW flash through .fwpkg file only.

Please find the minimum version required (1.0.14.1055) in below links:

- Windows <https://support.hpe.com/hpesc/public/swd/detail?swItemId=MTX-1b2c98e9d2594b9db679e89bbe#tab-history>
- Linux <https://support.hpe.com/hpesc/public/swd/detail?swItemId=MTX-207ea7e739f048049a66d61008#tab-history>
- VMware [https://support.hpe.com/hpesc/public/swd/detail?swItemId=MTX\\_141038fe565b457ca9fe4d28de#tab-history](https://support.hpe.com/hpesc/public/swd/detail?swItemId=MTX_141038fe565b457ca9fe4d28de#tab-history)

### **Fixes**

- Issue fixed for RDE dictionary reflash error during booting up

### **Enhancements**

Update build environment from SHA512 to SHA384

---

### **Firmware - Storage Fibre Channel**

[Top](#)

HPE Firmware Flash for Emulex Fibre Channel Host Bus Adapters for VMware vSphere 8.0  
Version: 2022.09.01 **(Recommended)**  
Filename: CP050120.compsig; CP050120.zip

### **Important Note!**

This Firmware package contains following firmware versions:

Adapter	Speed	Universal Boot Image	Firmware	UEFI	Boot Bios
HPE SN1200E 16Gb Dual Port Fibre Channel Host Bus Adapter	16Gb	14.0.499.21	14.0.499.21	14.0.499.2	14.0.490.0
HPE SN1200E 16Gb Single Port Fibre Channel Host Bus Adapter	16Gb	14.0.499.21	14.0.499.21	14.0.499.2	14.0.490.0
HPE SN1600E 32Gb Dual Port Fibre Channel Host Bus Adapter	32Gb	14.0.499.21	14.0.499.21	14.0.499.2	14.0.490.0
HPE SN1600E 32Gb Single Port Fibre Channel Host Bus Adapter	32Gb	14.0.499.21	14.0.499.21	14.0.499.2	14.0.490.0
HPE SN1610E 32Gb Single Port Fibre Channel Host Bus Adapter	32Gb	14.0.499.20	14.0.499.20	14.0.499.2	14.0.490.0
HPE SN1610E 32Gb Dual Port Fibre Channel Host Bus Adapter	32Gb	14.0.499.20	14.0.499.20	14.0.499.2	14.0.490.0
HPE SN1700E 64Gb Single Port Fibre Channel Host Bus Adapter	64Gb	14.0.499.20	14.0.499.20	14.0.499.2	14.0.490.0
HPE SN1700E 64Gb Dual Port Fibre Channel Host Bus Adapter	64Gb	14.0.499.20	14.0.499.20	14.0.499.2	14.0.490.0

### **Prerequisites**

Please consult SPOCK for a list of supported configurations available at the following link:

<http://www.hpe.com/storage/spock/>

## Enhancements

This Firmware package contains following firmware versions:

Adapter	Speed	Universal Boot Image	Firmware	UEFI	Boot Bios
HPE SN1200E 16Gb Dual Port Fibre Channel Host Bus Adapter	16Gb	14.0.499.21	14.0.499.21	14.0.499.2	14.0.490.0
HPE SN1200E 16Gb Single Port Fibre Channel Host Bus Adapter	16Gb	14.0.499.21	14.0.499.21	14.0.499.2	14.0.490.0
HPE SN1600E 32Gb Dual Port Fibre Channel Host Bus Adapter	32Gb	14.0.499.21	14.0.499.21	14.0.499.2	14.0.490.0
HPE SN1600E 32Gb Single Port Fibre Channel Host Bus Adapter	32Gb	14.0.499.21	14.0.499.21	14.0.499.2	14.0.490.0
HPE SN1610E 32Gb Single Port Fibre Channel Host Bus Adapter	32Gb	14.0.499.20	14.0.499.20	14.0.499.2	14.0.490.0
HPE SN1610E 32Gb Dual Port Fibre Channel Host Bus Adapter	32Gb	14.0.499.20	14.0.499.20	14.0.499.2	14.0.490.0
HPE SN1700E 64Gb Single Port Fibre Channel Host Bus Adapter	64Gb	14.0.499.20	14.0.499.20	14.0.499.2	14.0.490.0
HPE SN1700E 64Gb Dual Port Fibre Channel Host Bus Adapter	64Gb	14.0.499.20	14.0.499.20	14.0.499.2	14.0.490.0

## Supported Devices and Features

This component is supported on following Emulex Fibre Channel Host Bus adapters:

### 16Gb FC Adapter:

- HPE SN1200E 16Gb Dual Port Fibre Channel Host Bus Adapter
- HPE SN1200E 16Gb Single Port Fibre Channel Host Bus Adapter

### 32Gb FC Adapter:

- HPE SN1600E 32Gb Dual Port Fibre Channel Host Bus Adapter
- HPE SN1600E 32Gb Single Port Fibre Channel Host Bus Adapter
- HPE SN1610E 32Gb Dual port Fibre Channel Host Bus Adapter
- HPE SN1610E 32Gb Single port Fibre Channel Host Bus Adapter

### 64Gb FC Adapter:

- HPE SN1700E 64Gb Dual Port Fibre Channel Host Bus Adapter
- HPE SN1700E 64Gb Single Port Fibre Channel Host Bus Adapter

---

HPE Firmware Flash for QLogic Fibre Channel Host Bus Adapters for VMware vSphere 8.0

Version: 2022.09.01 **(Recommended)**

Filename: CP050119.compsig; CP050119.zip

## Important Note!

This Firmware package contains following firmware versions:

Adapter	Speed	MBI	Firmware	UEFI	Boot Bios
HPE SN1100Q 16GB Dual Port PCIe Fibre Channel Host Bus Adapter	16Gb	1.77.12	9.09.00	7.19	3.64
HPE SN1100Q 16GB Single Port PCIe Fibre Channel Host Bus Adapter	16Gb	1.77.12	9.09.00	7.19	3.64
HPE SN1600Q 32Gb Single Port Fibre Channel Host Bus Adapter	32Gb	1.77.12	9.09.00	7.19	3.64
HPE SN1600Q 32Gb Dual Port Fibre Channel Host Bus Adapter	32Gb	1.77.12	9.09.00	7.19	3.64
HPE SN1610Q 32Gb Dual Port Fibre Channel Host Bus Adapter	32Gb	02.06.27	09.09.00	7.28	0.0
HPE SN1610Q 32Gb Single Port Fibre Channel Host Bus Adapter	32Gb	02.06.27	09.09.00	7.28	0.0

## Prerequisites

Please consult SPOCK for a list of supported configurations available at the following link:

<http://www.hpe.com/storage/spock/>

## Enhancements

This Firmware package contains following firmware versions:

Adapter	Speed	MBI	Firmware	UEFI	Boot Bios
HPE SN1100Q 16GB Dual Port PCIe Fibre Channel Host Bus Adapter	16Gb	1.77.12	9.09.00	7.19	3.64
HPE SN1100Q 16GB Single Port PCIe Fibre Channel Host Bus Adapter	16Gb	1.77.12	9.09.00	7.19	3.64
HPE SN1600Q 32Gb Single Port Fibre Channel Host Bus Adapter	32Gb	1.77.12	9.09.00	7.19	3.64
HPE SN1600Q 32Gb Dual Port Fibre Channel Host Bus Adapter	32Gb	1.77.12	9.09.00	7.19	3.64
HPE SN1610Q 32Gb Dual Port Fibre Channel Host Bus Adapter	32Gb	02.06.27	09.09.00	7.28	0.0
HPE SN1610Q 32Gb Single Port Fibre Channel Host Bus Adapter	32Gb	02.06.27	09.09.00	7.28	0.0

## Supported Devices and Features

This component is supported on following Qlogic Fibre Channel Host Bus adapters:

### 16Gb Fibre Channel Host Bus Adapter:

- HPE SN1100Q 16GB Dual Port PCIe Fibre Channel Host Bus Adapter
- HPE SN1100Q 16GB Single Port PCIe Fibre Channel Host Bus Adapter

### 32Gb Fibre Channel Host Bus Adapter:

- HPE SN1600Q 32Gb Single Port Fibre Channel Host Bus Adapter
- HPE SN1600Q 32Gb Dual Port Fibre Channel Host Bus Adapter
- HPE SN1610Q 32Gb Dual Port Fibre Channel Host Bus Adapter
- HPE SN1610Q 32Gb Single Port Fibre Channel Host Bus Adapter

---

## Software - Management

[Top](#)

HPE Agentless Management Bundle Smart Component on ESXi for Gen10 and Gen10 Plus Servers

Version: 2022.09.01 **(Recommended)**

Filename: cp050760.compsig; cp050760.zip

### Enhancements

#### Agentless Management Service

- Added feature to allow custom trap community name for SMA mode.
- Added support for cpqSasLogDrv MIB OID for the HPE Smart Array P824i-p MR Gen10 Controller.

---

HPE Fiber Channel and Storage Enablement Bundle Smart Component for ESXi 8.0

Version: 2022.09.01 **(Recommended)**

Filename: cp051152.compsig; cp051152.zip

### Enhancements

Supports VMware ESXi 8.0

---

HPE iLO Driver Bundle Smart Component for ESXi 7.0

Version: 2022.09.01 **(Recommended)**

Filename: cp050763.compsig; cp050763.zip

### Fixes

- Fixed issue where ilo driver is failing to acquire contiguous physical memory below 4GB causing userworld apps like hponcfg to be unable to communicate with iLO.

### Enhancements

- Added support for vSphere 8.0

---

## Software - Storage Fibre Channel

[Top](#)

HPE QLogic Fibre Channel driver component for VMware vSphere 8.0

Version: 2022.09.01 **(Recommended)**

Filename: cp050118.compsig; cp050118.zip

## **Important Note!**

This component is intended to be used by HPE applications. It is a zip that contains the same driver deliverable available from the vmware.com and the HPE vibstdepot.hpe.com webpages, plus an HPE specific CPXXXX.xml file.

## **Prerequisites**

Please consult SPOCK for a list of supported configurations available at the following link:

<http://www.hpe.com/storage/spock/>

## **Enhancements**

Driver version 5.3.1.0

## **Supported Devices and Features**

This component is supported on following Qlogic Fibre Channel Host Bus adapters:

### **16Gb Fibre Channel Host Bus Adapter:**

- HPE SN1100Q 16GB Dual Port PCIe Fibre Channel Host Bus Adapter
- HPE SN1100Q 16GB Single Port PCIe Fibre Channel Host Bus Adapter

### **32Gb Fibre Channel Host Bus Adapter:**

- HPE SN1600Q 32Gb Single Port Fibre Channel Host Bus Adapter
- HPE SN1600Q 32Gb Dual Port Fibre Channel Host Bus Adapter
- HPE SN1610Q 32Gb Dual Port Fibre Channel Host Bus Adapter
- HPE SN1610Q 32Gb Single Port Fibre Channel Host Bus Adapter

---

## **Software - System Management**

[Top](#)

HPE Agentless Management Bundle for ESXi for HPE Gen10 and Gen10 Plus Servers

Version: 701.11.8.5 (**Recommended**)

Filename: amsdComponent\_701.11.8.5.22-1\_20314731.zip

## **Enhancements**

### **Agentless Management Service**

- Added feature to allow custom trap community name for SMA mode.
- Added support for cpqSasLogDrv MIB OID for the HPE Smart Array P824i-p MR Gen10 Controller.

---

HPE Fiber Channel and Storage Enablement Component for ESXi 8.0

Version: 3.9.0 (**Recommended**)

Filename: fc-enablement-component\_800.3.9.0.30-1\_20300413.zip

## **Enhancements**

Supports VMware ESXi 8.0

---

Integrated Smart Update Tools for VMware ESXi 8.0

Version: 800.3.0.0 (**Recommended**)

Filename: sutComponent\_800.3.0.0.14-0-signed\_component-20341487.zip

## **Important Note!**

Integrated Smart Update Tools for ESXi 8.0 provides support for firmware and driver updates via iLO Repository

## **Fixes**

See the [iSUT Release Notes](#) for information about the issues resolved in this release

## **Enhancements**

See the [iSUT Release Notes](#) for information about the enhancements in this release.

---

Get connected

[hpe.com/info/getconnected](https://hpe.com/info/getconnected)

Current HPE driver, support, and security alerts delivered directly to your desktop

© Copyright 2022 Hewlett Packard Enterprise Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HPE products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

Trademark acknowledgments, if needed.

Updated October 11 2022

