NVIDIA ConnectX-8 SuperNIC Firmware Release Notes v40.50.1002

Table of Contents

1 Release Notes Update History

Version	Date	Description
40.50.1002	August 2026	Initial release of this Release Notes version.

2 Overview

The NVIDIA ConnectX-8 SuperNIC leverages NVIDIA's next-generation adapter architecture to deliver unparalleled end-to-end 800Gb/s networking with performance isolation, essential for efficiently managing generative AI clouds. It provides 800Gb/s data throughput with PCI Express (PCIe) Gen6, offering up to 48 lanes for various use cases such as PCIe switching inside NVIDIA GPU systems. It also supports advanced NVIDIA In-Network Computing, MPI_Alltoall, as well as fabric enhancement features like quality of service and congestion control. The ConnectX-8 SuperNIC, featuring single-port OSFP224 and dual-port 112 connectors for the adapters, is compatible with various form factors, including OCP 3.0 and Card Electromechanical (CEM) PCIe x16. ConnectX-8 SuperNIC also supports NVIDIA Socket Direct™ 16-lane auxiliary card expansion.

2.1 Firmware Download

Please visit Firmware Downloads.

2.2 Document Revision History

A list of the changes made to this document are provided in Document Revision History.

3 Firmware Compatible Products

These are the release notes for the NVIDIA® ConnectX®-8 SuperNIC firmware. ConnectX®-8 supports the following protocols:

- InfiniBand - XDR, NDR, HDR, EDR, SDR
- Ethernet - 400GAUI4 C2M, 200GAUI2 C2M, 100GAUI1 C2M, 100GAUI1_SFE C2M, 400GAUI4 CR4, 200GAUI2 CR2, 200GAUI4 CR4, 100GAUI1 CR1, 100GAUI1_SFE CR1, 100GAUI2 CR2, CAUI4 CR4, 50GAUI1 CR1, 50GAUI2_RS_NS CR2, 50GAUI2_FC_NS CR2, 25GAUI1_RS C2M, 25GAUI1_FC C2M, 25GAUI1 C2M
- PCI Express Gen6, supporting backwards compatibility for v5.0, v4.0, v3.0, v2.0 and v1.1

⚠ When connecting an NVIDIA-to-NVIDIA adapter card in ETH PAM4 speeds, Auto-Neg should always be enabled.

⚠ Some of the Ethernet protocol speeds listed above are not supported by the firmware.. For more information, see **(40.49.1000) Known Issues** 4031430.

3.1 Supported Devices

Refer to the hardware documentation for the list of supported devices.

3.2 Driver Software, Tools and Switch Firmware

The following are the drivers' software, tools, switch/HCA firmware versions tested that you can upgrade from or downgrade to when using this firmware version:

	Supported Version
ConnectX-8 Firmware	40.50.1002 / 40.49.1014 / 40.48.1000
DOCA-HOST	3.5 0 / 3.4.0 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.
WinOF-2	26.7.50000 / 26.1.50000 / 25.10.51000 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.

	Supported Version
MFT	4.37.0-154 / 4.36.0-147 / 4.35.0-159 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.
FlexBoot	3.9.101
UEFI	14.42.11
MLNX-OS	3.12.6000 onwards
Quantum-2 FW	31.2016.2054 onwards
NVOS	25.02.6000 onwards
Quantum-3 FW (part of NVOS)	35.2016.2080 onwards

4 Changes and New Features

Security Hardening Enhancements: This release contains important reliability improvements and security hardening enhancements. NVIDIA recommends upgrading your devices' firmware to this release to improve the devices' firmware security and reliability.

To generate PLDM packages for firmware updates, users must install and use the MFT version that corresponds with the respective firmware release.

Feature/Change	Description
40.50.1002	

Dual Steering Multiplane Emulation Groups	Added a new capability to open two steering multiplane emulation groups. To enable this feature, the following parameters must be applied in addition to the standard configuration required for multiplane emulation:
	• **Port Mapping** (`Array[phy_port] = logical port number`):
	  • `MODULE_SPLIT_M0[0..7] = 1, 2, 3, 4, 5, 6, 7, 8` (Respectively)
	  • `MODULE_SPLIT_M0[8..15] = FF`
	• **Physical Functions:**
	  • `NUM_OF_PF = 8` (**Note:** `NUM_PFS` must equal the number of logical ports)
	• **Plane Configuration** (**Note:** `NUM_PLANES` must equal the number of logical ports on the plane group):
	  • `NUM_OF_PLANES_P1 = 4` (Logical port 1 has 4 planes, spanning logical ports 1-4)
	  • `NUM_OF_PLANES_P5 = 4` (Logical port 5 has 4 planes, spanning logical ports 5-8)
	  • `NUM_OF_PLANES_P# = 0` (For all other ports)
	• **Load Balancing (Packet Steering):**
	  • `LOAD_BALANCE_MODE_P1 = 3`
	  • `LOAD_BALANCE_MODE_P5 = 3`

Module Precoding Control	Module precoding control allows precoding to be enabled or disabled through PHY TLVs on both the host and the module, aligning NIC and module behavior with host PRM precoding control.
Modify-header Action Dependency	Added the missing source dependency check (SRC) for modify-header parallel actions, preventing odd actions from reading a field modified by the preceding even action.
FEC Alignment	FEC alignment is supported for lane speeds up to 100G per lane.
Bug Fixes	See *Bug Fixes in this Firmware Version* section.

4.1 Customer Affecting Changes

4.1.1 Changes in This Release

This section provides a list of changes that took place in the current version and break compatibility/interface, discontinue support for features and/or OS versions, etc.

Introduced in Version	Description
N/A	N/A

4.1.2 Changes Planned for Future Releases

This section provides a list of changes that will take place in a future version of the product and will break compatibility/interface, discontinue support for features and/or OS versions, etc.

Planned for Version	Description
N/A	N/A

4.1.3 Changes in Earlier Releases

This section provides a list of changes that took place throughout the past two major releases that broke compatibility/interface, discontinued support for features and/or OS versions, etc.

For an archive of all changes, please refer to the Release Notes History section.

Introduced in Version	Description

40.48.1000	When using an optics module, if the required speed is lower than the module's maximum supported speed, the user must configure the desired port link width.
	Note: This configuration must be repeated after every system reboot.
	To avoid reconfiguring it manually after each reboot, configure the requested speed using the following commands:
	`mlxconfig -d <device> set PHY_RATE_MASK_OVERRIDE_P<port>=1`
	`mlxfwreset -d <device> r -y --no_mst`
	`mlxconfig -d <device> set PHY_RATE_MASK_P<port>=<PTYS.speed_mask>`
	`mlxfwreset -d <device> r -y --no_mst`
	Where **<PTYS.speed_mask>** is the speed mask defined in the PRM.

	Starting with this release, ConnectX CoRIMs/cbor files are available through the NVIDIA Attestation RIM service. RIM service documentation is available here. Bundling ConnectX CoRIMs with the firmware zip files will be retired in a future release, and customers should plan to transition to the NVIDIA RIM service.

Introduced in Version	Description
40.47.1026	To align with updated Microsoft UEFI Secure Boot requirements and the upcoming end-of-life of the 2011 Certificate Authority (CA), NVIDIA is transitioning to the 2023 CA. To ensure successful loading of the Expansion ROM (ExpROM) during the UEFI Secure Boot process, system BIOS and operating system trust stores must be updated to include the 2023 CA. **Note:** When performing a firmware update of ConnectX and BlueField devices the new certificate is required for Secure Boot. To continue supporting Secure Boot, systems must be updated to recognize the "Microsoft Option ROM UEFI CA 2023."
40.46.1006	Renamed firmware-generated PLDM images to include the firmware name and PSID.
40.45.1020	Downgrading the ConnectX-8 SuperNIC firmware to a version earlier than the April release (40.45.1020) is not supported
	The default CNP moderation behavior has been changed from per-flow to per-port to align with previous device generations. A dedicated mlxconfig parameter, ROCE_CC_CNP_MODERATION, is available to modify this configuration if needed.
40.44.1036	In ConnectX-8, the DIAG_COUNTER interface has been changed from the following set of commands: • SET_DIAGNOSTIC_PARAMS • QUERY_DIAGNOSTIC_PARAMS • QUERY_DIAGNOSTIC_COUNTERS • ICMD_SET_DIAGNOSTIC_PARAMS • ICMD_QUERY_DIAGNOSTIC_PARAMS • ICMD_QUERY_DIAGNOSTIC_COUNTERS to: • DIAG_DATA_OWNERSHIP • DIAG_DATA_PARAMS_CONTEXT • DIAG_DATA_ID_LIST • DIAG_DATA_QUERY The old interface will now return zero values when queried.

4.1.4 Discontinued Features

List of features which are supported in previous generations of hardware devices.

- Ethernet:
  - T10 Data Integrity Field (DIF)
  - CRC
  - Transport Layer Security (TLS) handshake
  - NVMe over TCP acceleration
- InfiniBand:
  - FDR and lower speeds

For inquiries regarding mitigation, please contact NVIDIA Support.

4.2 Declared Unsupported Features

The following are the unsupported features for ConnectX-8 SuperNIC in this firmware version:

- Zero Touch Tuning (ZTT)
- Hot reset
- Segment on PCIe switch
- LAG Bonding with Q-Affinity
- ISSU
- Read-on-LAN is not supported on SKU 900-9X85E-00EX-MJA

5 Bug Fixes in this Firmware Version

Internal Ref.	Issue
5232646	**Description:** Resolved an issue where PLDM update from the auxiliary card side could fail on devices in multi-host configurations.
	Keywords: PLDM
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5073671	**Description:** Resolved an issue where a firmware assert could be raised due to a GPIO expander write timeout.
	Keywords: GPIO
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5130705 / 5131942 / NVbug 6376684	**Description:** Resolved an issue that could cause CoRIM FLM test failures during SBIOS and SMA validation on MGX Vera platform/system configurations.
	Keywords: CoRIM
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5139055 / 5072582 / NVbug 6427860	**Description:** Resolved an issue where firmware did not validate all port health threshold slots. Firmware now rejects unsupported threshold values instead of silently accepting them.
	Keywords: Firmware threshold
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5173627	**Description:** Resolved an issue where critical messages could be dropped due to queue overflow.
	Keywords: Queue overflow
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5041711	**Description:** Resolved an issue where, after toggling a port, the peer could incorrectly remain in the PHY-up state.
	Keywords: Port toggle
	Detected in version: 40.49.1014

Internal Ref.	Issue
	Fixed in Release: 40.50.1002
4971613 / 4852107	**Description:** Resolved an issue where, after setting FORCE on a specific speed and rebooting, mlxlink could incorrectly show all default speeds as enabled instead of showing only the forced speed.
	Keywords: mlxlink
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
4936278	**Description:** Added support for 10G and 25G using limited LLU request messages via the disable_b2b flag, with clause49_en enabled and the numerator/denominator configured to a 1:1 ratio.
	Keywords: LLU, 10G/25G speed
	Detected in version: 40.48.1000
	Fixed in Release: 40.50.1002
5042147 / 5041441	**Description:** Updated the Best Known Values (BKV) to version 1.1, improving the overall stability and robustness of the link-training reset flow during boot-up.
	Keywords: BKV
	Detected in version: 40.48.1000
	Fixed in Release: 40.50.1002
5012336 / NVbug 6140150	**Description:** Resolved a potential issue that could result in lower-than-expected bidirectional RDMA bandwidth in PCIe switch topologies.
	Keywords: RDMA bandwidth
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
4851684	**Description:** Resolved an issue where RDMA workloads using null memory regions with MKEY_BY_NAME enabled could experience up to a 30% throughput drop. The internal null/ dump-fill key was incorrectly routed through the slower signature datapath. Workloads not using MKEY_BY_NAME or null memory regions were unaffected.
	Keywords: RDMA
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
4851684	**Description:** Resolved an issue where, under high-stress LRO traffic, an internal FIFO could enter a hardware deadlock state, preventing the NIC from receiving additional traffic. The fix detects and releases the deadlock if it occurs, restoring the NIC to a functional state.

Internal Ref.	Issue
	Keywords: LRO traffic
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5174504	**Description:** Resolved an issue where intensive parallel QP INIT2RTR / DESTROY operations could incur high operation latency and reduce connection establishment rates when PCC was enabled.
	Keywords: PCC, INIT2RTR, DOCA PCC
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5044357	**Description:** Resolved an issue where a RISC could become stuck in a fw_memcpy loop called by the multi-ASIC I2C mailbox function. This could occur if firmware read invalid parameters from I2C HW fields. Parameter validation was added before calling fw_memcpy, along with an error indication when the parameters are invalid.
	Keywords: RISC
	Detected in version: 40.47.1026
	Fixed in Release: 40.50.1002
5166648	**Description:** Resolved an issue involving an uninitialized stack read and an inverted EtherType check in the RX parser.
	Keywords: RX parser
	Detected in version: 40.47.1026
	Fixed in Release: 40.50.1002
4937640	**Description:** Resolved an issue where a PPCC mlxreg read operation could succeed when executed on a PF that is not the port owner, which could have security implications.
	Keywords: PCC, PPCC
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5161601	**Description:** Resolved an issue where multiplane emulation mode conditioning used the raw num_pfs value before invalid values were reduced during FMT initialization. The fix reuses the PF reduction calculation after logical-port discovery, without updating PCI settings before FMT.
	Keywords: PF, FMT initialization
	Detected in version: 40.49.1014

Internal Ref.	Issue
	Fixed in Release: 40.50.1002
4670019	**Description:** Resolved an issue where GPIO expander write operations could take more than 200 ms.
	Keywords: GPIO expander
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5154717 / 5152076 / 5154806 / NVbug 6422398	**Description:** Resolved an issue where, when a transceiver module held its I2C management bus line stuck (SDA or SCL), the firmware detected the stuck bus and reset the module. After the reset, the module remained in low-power mode and its datapath was not re-established, preventing the link from coming up. On a shared module, all associated ports could drop simultaneously.
	Keywords: Cables
	Detected in version: 40.45.3034
	Fixed in Release: 40.50.1002
4923806 / 4867111 / 4927875	**Description:** Resolved an issue that could cause RoCE communication failures between specific NIC pairs after repeated queue pair (QP) create/destroy cycles. In the affected cases, traffic could enter a bad state and fail with transport retry counter exceeded errors. This fix prevents the NIC from entering that condition.
	Keywords: RDMA; connection failure
	Detected in version: 40.46.3036
	Fixed in Release: 40.50.1002
5155445	**Description:** Resolved an issue with SyncE frequency adjustment overflow when using a third-party secondary vendor.
	Keywords: SyncE frequency
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5125089	**Description:** Resolved an issue where an incorrect Flex Parser configuration could trigger an assert. A syndrome was added to notify HWS when an invalid configuration is provided.
	Keywords: Flex Parser configuration
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5093695	**Description:** Resolved an issue where using mlxconfig to configure all modules with the 0xff module split value could reset the modules to the default configuration.

Internal Ref.	Issue
	Keywords: mlxconfig
	Detected in version: 40.48.1000
	Fixed in Release: 40.50.1002
5090639	**Description:** Resolved an issue in single-domain mode where an FLR on a GVMI did not clear the per-domain usage bits in `SP_core_diag_multi_domain->pcie_link_latency_max_read` and `SP_core_diag_multi_domain->pcie_link_latency_min_read`. As a result, when the next GVMI took ownership and attempted to configure these counters, it could receive `DIAG_DATA_LIST_WRITE_STATUS_NO_RESOURCES`.
	Keywords: PCIe Link Latency Arch counters
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002
5089348	**Description:** Resolved an issue where the MFT database was not updated to configure profile 12 for the Flex Parser.
	Keywords: MFT, Flex Parser
	Detected in version: 40.48.1000
	Fixed in Release: 40.50.1002
4950259	**Description:** Unblocked 4-port, single-LAG configurations by updating the INI setting nv_config.global.power_cap.max_num_port=4, and fixed the Queue Affinity error seen in LAG mode.
	Keywords: LAG, Queue Affinity
	Detected in version: 40.48.1000
	Fixed in Release: 40.50.1002
5077326 / 4902171	**Description:** Resolved an issue where a slow module response to the NCSI GetModuleSerialData command could incorrectly trigger a dead IRISC watchdog assert. The firmware now handles this waiting state correctly, so slow module responses no longer cause a false watchdog hang indication.
	Keywords: NCSI
	Detected in version: 40.44.1036
	Fixed in Release: 40.50.1002
5025688	**Description:** Resolved an incorrect INI settings on one downstream link of CX8_Rithmax_CX8aas_BP_4P. All DSLs now operate in Gen6.
	Keywords: INI settings
	Detected in version: 40.48.1000

Internal Ref.	Issue
	Fixed in Release: 40.50.1002
4947006 / 4980956 / 4982511	**Description:** Resolved two rare race conditions in the PCC steering extension when the feature is enabled through path migration, empty RTT, or a custom header. Under rare timing conditions during QP creation or destruction, these issues could cause traffic disruption or resource leaks.
	Keywords: PCC steering extension
	Detected in version: 40.49.1014
	Fixed in Release: 40.50.1002

6 Known Issues

VF Network Function Limitations in SR-IOV Legacy Mode & in Switchdev Mode

Dual Port Device	Single Port Device
127 VF per PF (254 functions)	127

VF+SF Network Function Limitations in Switchdev Mode

Dual Port Device	Single Port Device
• 127 VF per PF (254 functions) • 512 PF+VF+SF per PF (1024 functions)	• 127 VF (127 functions) • 512 PF+VF+SF per PF (512 functions)

Internal Ref.	Issue
5234743 / 5145460	**Description:** Query Port Telemetry v2 (CMD 0x14) reports tag 0x20 (rx_los_mask), which should not be populated.
	Workaround: N/A
	Keywords: Query Port Telemetry v2
	Detected in version: 40.50.1002
5234750	**Description:** The Plug Voltage sensor does not read the value from the second module.
	Workaround: N/A
	Keywords: Plug Voltage sensor
	Detected in version: 40.50.1002
5230758 / 4553130 / 5231785	**Description:** For LPO cables, link-up may fail after the port on the peer device is toggled.
	Workaround: Toggle the port on the SuperNIC side.
	Keywords: LPO cables
	Detected in version: 40.50.1002
5227817 / NVbug 6618275	**Description:** Running applications concurrently on both GPU and CPU with different numbers of QPs, for example 1K QPs on the GPU and 1 QP on the CPU, is not supported when using hardware QoS bandwidth sharing.
	Workaround: N/A
	Keywords: GB300, QoS, GPU, CPU
	Detected in version: 40.50.1002
5085252 / 5090080	**Description:** Following a hot reset, a false Surprise Link Down (SLD) alarm may be reported intermittently on the Downstream Port (DSP) through the Advanced Error Reporting (AER) register. This is a false alarm and has no functional impact on link stability or data integrity.
	Workaround: N/A
	Keywords: Surprise Link Down (SLD)

Internal Ref.	Issue
	Detected in version: 40.49.1014
5133249	**Description:** A recent update to the default unmapped local port value causes the LED flow to query an invalid port, resulting in a failure of the LED logic.
	Workaround: N/A
	Keywords: LED indication
	Detected in version: 40.49.1014
4901395 / 4919262	**Description:** When PXE is enabled for both Legacy (FlexBoot) and UEFI drivers, the UEFI driver supports only a single TFTP session after the DHCP handshake.
	Workaround: For scenarios that require multiple concurrent TFTP sessions, disable PXE filters via the HCA HII menu.
	Keywords: PXE, Filter enabled, UEFI
	Detected in version: 40.49.1014
5072222	**Description:** In DirectNIC and SocketDirect topologies, PCIe TLP Analyzer visibility is limited to the host's management domain. Traffic passing through an unmanaged host link, such as a DirectNIC connection to a GPU, is not visible to the local tracer.
	Workaround: N/A
	Keywords: PCIe TLP Analyzer
	Detected in version: 40.49.1014
4893639	**Description:** Direct Memory Access (DMA) protection is not supported in the ipxe.efi driver.
	Workaround: N/A
	Keywords: DMA
	Detected in version: 40.48.1000
4835832	**Description:** In rare cases, certain module types may experience link-up failures.

	Workaround: Configure the requested speed using the following commands:
	`mlxconfig -d <device> set PHY_RATE_MASK_OVERRIDE_P<port>=1`
	`mlxfwreset -d <device> r -y --no_mst`
	`mlxconfig -d <device> set PHY_RATE_MASK_P<port>=<PTYS.speed_mask> (speed mask is defined in PRM)`
	`mlxfwreset -d <device> r -y --no_mst`

	Keywords: Cables
	Detected in version: 40.48.1000
4594515 / 4622127 / NVbug 5460310	**Description:** Running perftest with ib_send_bw using a message size of 4KB and 1000 QPs does not reach the expected full line rate.
	Workaround: N/A
	Keywords: perftest
	Detected in version: 40.48.1000
4604969	**Description:** Probe packets might be dropped at the transmission stage when multiple congestion control flows are active.

Internal Ref.	Issue
	Workaround: N/A
	Keywords: PCC, RTT, probe
	Detected in version: 40.47.1026
4496642	**Description:** The timestamps (t2, t4) of the received RTT probes are taken from the free-running clock, even when ROCE_CC_RTT_TIMESTAMP_FORMAT is set to 0x02. The format of all RTT probe timestamps can be found in HCA_CAP.rtt_timestamp_format.
	Workaround: N/A
	Keywords: RTT RTC timestamp
	Detected in version: 40.47.1026
4705241	**Description:** When Quantum-2 is part of an XDR topology, serving as a leaf switch connected to NDR-based hosts, a bandwidth degradation of approximately 3–7 Gb/s is expected.
	Workaround: N/A
	Keywords: XDR, NDR, Quantum-2
	Detected in version: 40.47.1026
4705948	**Description:** When using DC as the InfiniBand transport type to perform an ib_read RDMA operation between ConnectX-7 (NDR) and ConnectX-8, a bandwidth degradation of approximately 25% may be observed when using a low number of QPs (1–16). The performance degradation diminishes as the number of QPs increases.
	Workaround: N/A
	Keywords: DC, ib_read RDMA, NDR, performance
	Detected in version: 40.47.1026
4685736	**Description:** Creating a DPA process that allocates a 128 MB data segment and loads a dynamic library may fail with syndrome 0xdc30ac.
	Workaround: Limit the DPA application's data segment size to 64 MB.
	Keywords: DPA
	Detected in version: 40.47.1026
4618452	**Description:** When only a partial number of fiber lanes are connected to a module (i.e., not all supported lanes are populated), the module re-insertion process may experience an extended link-up time of approximately one minute. This occurs because the module requires additional time to detect that only a subset of the lanes is connected.
	Workaround: N/A
	Keywords: Cables
	Detected in version: 40.47.1026
4270913	**Description:** For DC InfiniBand transport, using SRQs requires distributing QPs across four SRQs. Each QP must be assigned to its own SRQ when there are up to four QPs.
	Workaround: If more than four QPs are used, they must be evenly distributed across the four available SRQs.
	Keywords: DC InfiniBand transport, SRQ, QPs
	Detected in version: 40.47.1026

Internal Ref.	Issue
4665802	**Description:** Due to a re-burst scheduler limitation, the expected “speed of light” value for single QP RDMA Write on ConnectX-8 devices cannot be calculated. Currently, verification is limited to identifying performance degradations relative to previous firmware versions.
	Workaround: N/A
	Keywords: Single QP RDMA Write
	Detected in version: 40.47.1026
4412310	**Description:** Split operation (port splitting) on the second port is not supported on the ConnectX-8 SuperNIC model 900-9X81Q-00CN-ST0.
	Workaround: N/A
	Keywords: Port splitting
	Detected in version: 40.45.1020
4394475	**Description:** The existing congestion control configuration applies globally, rather than on a per-priority basis.
	Workaround: Ensure that the configuration values for all priorities are aligned in either `mlxconfig ROCE_CC_PRIO_MASK_P$port` or `sysfs ecn/roce_rp/enable/$port`.
	Keywords: Congestion control, ROCE_CC_PRIO
	Detected in version: 40.45.1020
4240828	**Description:** Simultaneous access to the same PDDR module info page is not supported.
	Workaround: N/A
	Keywords: PDDR module
	Detected in version: 40.45.1020
4230775	**Description:** Due to a known issue, telemetry rate must be set to lower than 3 minutes.
	Workaround: N/A
	Keywords: Telemetry rate
	Detected in version: 40.44.0212
4038325 / 4031430 / 4038341 / 4046105	**Description:** Connecting to systems with NRZ speeds of 1,10, 25, 40, 50, or 100Gb/s is not supported.
	Workaround: N/A
	Keywords: NRZ, Connectivity
	Detected in version: 40.44.0208
4201405	**Description:** Upgrading to firmware 40.44.0xxx from any previous Engineering Sample version requires power cycling the driver and not just resetting it (using mlxfwreset).
	Workaround: N/A
	Keywords: Upgrade, power cycle, reset
	Detected in version: 40.44.0208

7 PreBoot Drivers (FlexBoot/UEFI)

7.1 FlexBoot Changes and New Features

For further information, please refer to the FlexBoot Release Notes.

7.2 UEFI Changes and Major New Features

For further information, please refer to the UEFI Release Notes.

8 Validated and Supported Cables and Switches

8.1 Validated and Supported Cables and Modules

8.1.1 Cables Lifecycle Legend

Lifecycle Phase	Definition
EOL	End of Life
LTB	Last Time Buy
HVM	GA level
MP	GA level
P-Rel	GA level
Preliminary	Engineering Sample
Prototype	Engineering Sample

⚠ For information on any cable limitations, please refer to Known Issues section.

⚠ The cables below are supported by all of the Supported Devices when used with a compatible connector. Use of copper cables on mezzanine board platforms requires NVIDIA approval for the specific product architecture.

8.1.2 XDR / 800GbE / 1600GbE Cables

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Description	Lifecycle Phase
OSFP	XDR	1600GE	980-9IAM1-00X001	N/A	NVIDIA Linear Active Copper Cable, 1600Gbps to 1600Gbps, OSFP, 1.1m, RHS to RHS, NIC to NIC	P-Rel

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Description	Lifecycle Phase
OSFP	XDR	1600GE	980-9IAM2-00X001	N/A	NVIDIA Active copper cable, 1600Gbps to 1600Gbps, OSFP, 1.1m, RHS to RHS, NVlink	P-Rel
OSFP to 2xOSFP	XDR	1600GE	980-9IAO5-00X001	N/A	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 2x800Gb/s, OSFP to 2xOSFP, 1m, IHS to RHS	P-Rel
OSFP to 2xOSFP	XDR	1600GE	980-9IAO5-00X002	N/A	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 2x800Gb/s, OSFP to 2xOSFP, 2m, IHS to RHS	P-Rel
OSFP to 2xOSFP	XDR	1600GE	980-9IAO5-00X01A	N/A	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 2x800Gb/s, OSFP to 2xOSFP, 1.5m, IHS to RHS	P-Rel
OSFP to 2xOSFP	XDR	1600GE	980-9IAQ5-00X001	N/A	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 4x400Gb/s, OSFP to 4xOSFP, 1m, IHS to RHS	Prototype
OSFP to 2xOSFP	XDR	1600GE	980-9IAQ5-00X002	N/A	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 4x400Gb/s, OSFP to 4xOSFP, 2m, IHS to RHS	Prototype
OSFP to 2xOSFP	XDR	1600GE	980-9IAQ5-00X01A	N/A	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 4x400Gb/s, OSFP to 4xOSFP, 1.5m, IHS to RHS	Prototype
OSFP	XDR	NA	980-9IAT0-00XM00	N/A	NVIDIA single port transceiver, DR4, 800Gbps, OSFP, MPO(APC), 1310nm SMF, up to 500m, RHS, FRO, XDR Only	P-Rel

8.1.3 NDR / 400GbE / 800GbE Cables

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Description	Lifecycle Phase
OSFP	NDR	NA	980-9I600-00N003	MCA4J80-N003-FLT	Active copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 3m, flat to flat	MP
OSFP	NDR	NA	980-9I601-00N003	MCA4J80-N003-FTF	NVIDIA Active copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 3m, flat to fin	MP
O22O	NDR	800GE	980-9I948-00N004	MCA7J60-N004	NVIDIA active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 4m, fin to flat	EOL [P-Rel]

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Description	Lifecycle Phase
O22O	NDR	800GE	980-9I949-00N005	MCA7J60-N005	NVIDIA active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 5m, fin to flat	EOL [P-Rel]
OSFP to 2xQSFP112	NDR	800GE	980-9I81B-00N004	MCA7J65-N004	NVIDIA Active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112, 4m, fin to flat	EOL [P-Rel]
OSFP to 2xQSFP112	NDR	800GE	980-9I81C-00N005	MCA7J65-N005	NVIDIA Active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112, 5m, fin to flat	EOL [P-Rel]
OSFP	NDR	800GE	980-9IA0G-00N001	MCP4Y10-N001-FLT	NVIDIA passive Copper cable, 800(2x400)Gbps, OSFP to OSFP, 1m, flat to flat	LTB [MP]
OSFP	NDR	800GE	980-9IA0H-00N001	MCP4Y10-N001-FTF	NVIDIA passive Copper cable, 800(2x400)Gbps, OSFP to OSFP, 1m, fin to flat	LTB [MP]
OSFP	NDR	800GE	980-9IA0J-00N002	MCP4Y10-N002-FLT	NVIDIA passive Copper cable, 800(2x400)Gbps, OSFP to OSFP, 2m, flat to flat	LTB [MP]
OSFP	NDR	800GE	980-9IA0L-00N00A	MCP4Y10-N00A-FLT	NVIDIA passive Copper cable, 800(2x400)Gbps, OSFP to OSFP, 0.5m, flat to flat	LTB [MP]
OSFP	NDR	800GE	980-9IA0R-00N01A	MCP4Y10-N01A-FLT	NVIDIA passive Copper cable, 800(2x400)Gbps, OSFP to OSFP, 1.5m, flat to flat	LTB [MP]
OSFP to 2xOSFP	NDR	800GE	980-9I432-00N001	MCP7Y00-N001	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP,1m, fin to flat	P-Rel
OSFP to 2xOSFP	NDR	800GE	980-9I433-00N001	MCP7Y00-N001-FLT	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP,1m, flat to flat	LTB [P-Rel]
OSFP to 2xOSFP	NDR	800GE	980-9I924-00N002	MCP7Y00-N002	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 2m, fin to flat	P-Rel
OSFP to 2xOSFP	NDR	800GE	980-9I925-00N002	MCP7Y00-N002-FLT	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 2m, flat to flat	LTB [P-Rel]
OSFP to 2xOSFP	NDR	800GE	980-9I92N-00N003	MCP7Y00-N003	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 3m, fin to flat	P-Rel
OSFP to 2xOSFP	NDR	NA	980-9I926-00N01A	MCP7Y00-N01A	NVIDIA passive copper splitter cable, IB twin port NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP,1.5m	EOL [P-Rel]
OSFP to 2xOSFP	NDR	NA	980-9I927-00N01A	MCP7Y00-N01A-FLT	NVIDIA passive copper splitter cable, IB twin port NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP,1.5m, flat top	EOL [P-Rel]

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Description	Lifecycle Phase
OSFP to 2xOSFP	NDR	NA	980-9I92O-00N02A	MCP7Y00-N02A	NVIDIA passive copper splitter cable, IB twin port NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP, 2.5m	EOL [P-Rel]
OSFP to 2xQSFP112	NDR	800GE	980-9I928-00N001	MCP7Y10-N001	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112,1m, fin to flat	LTB [P-Rel]
OSFP to 2xQSFP112	NDR	800GE	980-9I929-00N002	MCP7Y10-N002	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112,2m, fin to flat	LTB [P-Rel]
OSFP to 2xQSFP112	NDR	800GE	980-9I80P-00N003	MCP7Y10-N003	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112,3m, fin to flat	LTB [P-Rel]
OSFP to 2xQSFP112	NDR	800GE	980-9I80A-00N01A	MCP7Y10-N01A	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112,1.5m, fin to flat	EOL [P-Rel]
OSFP to 2xQSFP112	NDR	800GE	980-9I80Q-00N02A	MCP7Y10-N02A	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112,2.5m, fin to flat	EOL [P-Rel]
OSFP to 4xQSFP112	NDR	800GE	980-9I80B-00N001	MCP7Y40-N001	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xQSFP112, 1m, fin to flat	LTB [P-Rel]
OSFP to 4xQSFP112	NDR	800GE	980-9I80C-00N002	MCP7Y40-N002	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xQSFP112, 2m, fin to flat	LTB [P-Rel]
OSFP to 4xQSFP112	NDR	800GE	980-9I75R-00N003	MCP7Y40-N003	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xQSFP112, 3m, fin to flat	LTB [P-Rel]
OSFP to 4xQSFP112	NDR	800GE	980-9I75D-00N01A	MCP7Y40-N01A	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xQSFP112, 1.5m, fin to flat	EOL [P-Rel]
OSFP to 4xQSFP112	NDR	800GE	980-9I75S-00N02A	MCP7Y40-N02A	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xQSFP112, 2.5m, fin to flat	EOL [P-Rel]
OSFP to 4xOSFP	NDR	800GE	980-9I75E-00N001	MCP7Y50-N001	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 1m, fin to flat	P-Rel
OSFP to 4xOSFP	NDR	800GE	980-9I75F-00N001	MCP7Y50-N001-FLT	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 1m, flat to flat	LTB [P-Rel]
OSFP to 4xOSFP	NDR	800GE	980-9I46G-00N002	MCP7Y50-N002	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 2m, fin to flat	P-Rel
OSFP to 4xOSFP	NDR	800GE	980-9I46H-00N002	MCP7Y50-N002-FLT	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 2m, flat to flat	LTB [P-Rel]

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Description	Lifecycle Phase
OSFP to 4xOSFP	NDR	800GE	980-9I46T-00N003	MCP7Y50-N003	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 3m, fin to flat	P-Rel
OSFP to 4xOSFP	NDR	800GE	980-9I46I-00N01A	MCP7Y50-N01A	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 1.5m, fin to flat	EOL [P-Rel]
OSFP to 4xOSFP	NDR	800GE	980-9I46J-00N01A	MCP7Y50-N01A-FLT	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 1.5m, flat to flat	EOL [P-Rel]
OSFP to 4xOSFP	NDR	800GE	980-9I46U-00N02A	MCP7Y50-N02A	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 2.5m, fin to flat	EOL [P-Rel]
QSFP112	NDR	400GE	980-9I693-00NS00	MMA1Z00-NS400	NVIDIA single port transceiver, 400Gbps, QSFP112, MPO12 APC, 850nm MMF, up to 50m, flat top	P-Rel
OSFP	NDR	800GE	980-9I51A-00NS00	MMA4Z00-NS-FLT	NVIDIA twin port transceiver, 800(2x400)Gbps, OSFP, 2xMPO12 APC, 850nm MMF, up to 50m, flat top	MP
OSFP	NDR	400GE	980-9I51S-00NS00	MMA4Z00-NS400	NVIDIA single port transceiver, 400Gbps, OSFP, MPO12 APC, 850nm MMF, up to 50m, flat top	MP
QSFP112	NDR	400GE	980-9I068-00NM00	MMS1X00-NS400	NVIDIA single port transceiver, 400Gbps, QSFP112, MPO, 1310nm SMF, up to 500m, flat top	P-Rel
OSFP	NDR	800GE	980-9I301-00NM00	MMS4X00-NM-FLT	NVIDIA twin port transceiver, 800(2x400)Gbps, OSFP, 2xMPO12 APC, 1310nm SMF, up to 500m, flat top	P-Rel
OSFP	NDR	NA	980-9I30I-00NM00	MMS4X00-NS-FLT	NVIDIA twin port transceiver, 800Gbps,2xNDR, OSFP, 2xMPO12 APC, 1310nm SMF, up to 100m, flat top	EOL [MP]
OSFP	NDR	400GE	980-9I31N-00NM00	MMS4X00-NS400	NVIDIA single port transceiver, 400Gbps, OSFP, MPO12 APC, 1310nm SMF, up to 100m, flat top	MP

8.1.4 HDR / 200GbE Cables

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Description	Lifecycle Phase
QSFP56	HDR	200GE	980-9I548-00H001	MCP1650-H001E30	Nvidia Passive Copper cable, up to 200Gbps, QSFP56 to QSFP56, 1m	HVM
QSFP56	HDR	200GE	980-9I549-00H002	MCP1650-H002E26	Nvidia Passive Copper cable,up to 200Gbps, QSFP56 to QSFP56, 2m	HVM

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Description	Lifecycle Phase
QSFP56	HDR	200GE	980-9I54A-00H00A	MCP1650-H00AE30	Nvidia Passive Copper cable, up to 200Gbps, QSFP56 to QSFP56, 0.5m	EOL [HVM]
QSFP56	HDR	200GE	980-9I54B-00H01A	MCP1650-H01AE30	Nvidia Passive Copper cable, up to 200Gbps, QSFP56 to QSFP56, 1.5 m	HVM
QSFP56 to 2xQSFP56	HDR	200GE	980-9I39E-00H001	MCP7H50-H001R30	Nvidia Passive copper splitter cable, 200Gbps to 2x100Gbps, QSFP56 to 2xQSFP56, 1m	EOL [HVM]
QSFP56 to 2xQSFP56	HDR	400GE	980-9I46K-00H001	MCP7Y60-H001	NVIDIA passive copper splitter cable, 400(2x200)Gbps to 2x200Gbps, OSFP to 2xQSFP56, 1m, fin to flat	EOL [MP]
QSFP56 to 2xQSFP56	HDR	400GE	980-9I46L-00H002	MCP7Y60-H002	NVIDIA passive copper splitter cable, 400(2x200)Gbps to 2x200Gbps, OSFP to 2xQSFP56, 2m, fin to flat	EOL [MP]
QSFP56	HDR	NA	980-9I055-00H000	MMS1W50-HM	Mellanox transceiver, IB HDR, up to 200Gb/s, QSFP56, LC-LC, 1310nm, FR4	MP

8.1.5 Supported 3rd Party Cables and Modules

Third-party devices that have not been qualified by NVIDIA may be used; however, please be aware that no performance guarantees are provided.
Any issues that arise will require initiating a new feature request process for third-party support.

Form Factor	Eth Data Rate	Cable P/N	Description
OSFP to 2xQSFP112	800GE	C-OSG8CNSxxx-N00	INNOLIGHT 800G DR8 OSFP TO 2X400G QSFP112 DR4 BREAKOUT AOC
OSFP to 2xQSFP112	800GE	RTXM600-710	ACCELINK 800G OSFP TO 2X400G QSFP112 BREAKOUT AOC
OSFP RHS	800GE	T-RS8CNT-NMT	INNOLIGHT 800G DR8 OSFP RHS, DUAL MPO-12 APC
OSFP RHS	800GE	EOLO-138HG-5H-DR2	EOPTOLINK 800G 2XDR4 MODULE
OSFP RHS	800GE	AGP800C11311S50	ACCELIGHT 2X400G DR4 OSFP MODULE
OSFP RHS	800GE	FTCE4517E1PCA-2N	CREDO 800G OSFP TO 2X400G QSFP112 ACC
OSFP RHS	800GE	EOLO-138HG-G2-RMT	EOPTOLINK 800G DR8 OSFP RHS

Form Factor	Eth Data Rate	Cable P/N	Description
OSFP IHS to RHS	800GE	DMS8821X-Ecxx rev1	HISENSE 800G OSFP TO RHS AOC
OSFP IHS to RHS	800GE	RTXM600-96xx rev1	WTD 800G OSFP TO RHS AOC
OSFP IHS to RHS	800GE	TF-O330xx-C003 rev1	CREALIGHTS 800G OSFP TO RHS AOC
OSFP RHS	800GE	MTRO-H9F6VR-01	HGTECH 800G 2XVR4 OSFP RHS
OSFP RHS	800GE	RTXM600-565-R0	ACCELINK 800G VR8 OSFP RHS DUAL MPO-12
OSFP RHS	800GE	EOLO-138HG-5H-2B	EOPTOLINK 800G 2XDR4 OSFP RHS
OSFP RHS	800GE	T-RH8CNH-NFM	INNOLIGHT 800G 2XDR4 OSFP RHS
OSFP to 2xQSFP112	800GE	DME8821X-ECxx (rev: 05)	HISENSE 800G OSFP IHS TO 2X400G QSFP112 AOC
OSFP to 2xQSFP112	800GE	RTXM600-7xx (rev: R1)	WTD 800G OSFP IHS TO 2X400G QSFP112 AOC
OSFP to 2xQSFP112	800GE	TF-O330xx-C003(rev:R1)	CREALIGHT 800G OSFP IHS TO 2X400G QSFP112 AOC
OSFP RHS	800GE	SPQ1-8E8-8SO-COD	800G OSFP 2VR4 RHS MODULE
QSFP112	400GE	AQQLBCQ4EDLA1729	AOI 400G FR4 QSFP112 MODULE
QSFP56	400GE	ATRF-C020	HGTECH 200G QSFP56 AOC 20M
QSFP112 to QSFP-DD	400GE	C-GD4CNS010-N00	INNOLIGHT 400G QSFP112 TO 400G QSFP-DD AOC
QSFP112	400GE	CTF4XFR4CS1-01	INNOLIGHT 400G-FR4 MODULE
OSFP RHS	400GE	EOLO-134HG-5H-B	EOPTOLINK 400G OSFP DR4 MODULE
QSFP112	400GE	RTXM600-610	ACCELINK 400G QSFP-DD TO QSFP112 AOC
QSFP112	400GE	T-GQ4CNT-N00	INNOLIGHT 400G QSFP112 FR4 MODULE, LC
OSFP RHS	400GE	T-RS4CNH-NFL	INNOLIGHT 400G (BRCM LASER)
OSFP RHS	400GE	T-RS4CNH-NFM	INNOLIGHT 400G (SUMI LASER)
OSFP RHS	400GE	T-OH4CNT-N00	INNOLIGHT 400G DR4+ QSFP112 MODULE
OSFP RHS	400GE	EOLO-134HG-5H-DR2	EOPTOLINK 400G DR4 MODULE
OSFP IHS to RHS	400GE	NND1DM-F330	AMPHENOL 400G_8X DAC OSFP TO RHS 2.2M
QSFP112	400GE	GX-QS400GOCDxxB1	ACCELINK 400G QSFP112 AOC
QSFP112	400GE	XC-QS400GOCDxxB1	HISENSE 400G QSFP112 AOC
QSFP112	400GE	NDYRYF-0001	AMPHENOL 400G QSFP112 DAC

Form Factor	Eth Data Rate	Cable P/N	Description
QSFP112	400GE	NDYRYH-0005	AMPHENOL 400G QSFP112 DAC
QSFP112	200GE	EOLQ-132HG-5H-M3	EOPTOLINK 200G QSFP112 DR2 MODULE
QSFP56	200GE	RTXM500-301-F1	ACCELINK 200G QSFP56 SR4
QSFP112	200GE	RTXM600-338-R0	ACCELINK 200G QSFP112 VR2 MODULE
QSFP56	200GE	T-FX4FNS-N00	INNOLIGHT 200G QSFP56 SR4 MODULE
QSFP112	200GE	T-GP2CNH-NR0	INNOLIGHT 200G QSFP112 DR2 MODULE, MPO-12
QSFP112	200GE	TR-HM4M085V-CF21	CREALIGHT 200G QSFP112 VR2 MODULE
QSFP56	200GE	QSFP-200-CU3M	200G QSFP56 TO QSFP56 PASSIVE COPPER CABLE, 3M
QSFP112	200GE	MTRQ-4S104-02	HGTECH--200G SR4
QSFP56	200GE	LMTQF022-SD-R	LUXSHARE 200G DAC
QSFP56 to DSFP	100GE	DMM8211X-DCxx rev:11	ACCELINK 100G AOC QSFP56 TO DSFP

8.1.6 Tested Switches

8.1.7 XDR / 800GbE Switches

Speed	NVIDIA SKU	Legacy OPN	Description	LifeCycle Phase
XDR	920-9B34F-00RX-FS0	Q3200-RA	Quantum-3 based Two-Adjoining XDR InfiniBand Switches, Q3200-RA, 2U, with 36 XDR Ports over 18 OSFP cages per Switch, 4 Power Supplies (Power Cords Not Included), Standard Depth, Managed, C2P Airflow, Rail Kit	Prototype
XDR	920-9B36F-00RX-8S0	Q3400-RA	NVIDIA Quantum-3 based XDR InfiniBand Switch, Q3400-RA, 4U, 144 XDR Ports over 72 OSFP Cages, 8 Power Supplies (Power Cords Not Included), Standard Depth, Managed, C2P Airflow, Rail Kit	Prototype
800GbE	920-9N42F-00RI-xxx	SN5600	NVIDIA Spectrum-4 based 800GbE 2U Open Ethernet switch with ONIE and NOS Authentication, 64 OSFP ports and 1 SFP28 port, 2 power supplies (AC), x86 CPU, Secure-boot, standard depth, C2P airflow, Tool-less Rail Kit	P-Rel

8.1.8 NDR / 400GbE Switches

Speed	NVIDIA SKU	Legacy OPN	Description	LifeCycle Phase
NDR	920-9B210-00FN-xxx	QM9700	NVIDIA Quantum 2 based NDR InfiniBand Switch, 64 NDR ports, 32 OSFP ports, 2 Power Supplies (AC), Standard depth, Managed, P2C airflow, Rail Kit	MP
400GbE	920-9N301-00xB-xxx	SN4700	NVIDIA Spectrum-3 based 400GbE, 1U Open Ethernet switch, 32xQSFP-DD ports, x86 CPU, standard depth	MP

8.1.9 HDR / 200GbE Switches

Speed	NVIDIA SKU	Legacy OPN	Description	LifeCycle Phase
200GbE	920-9N201-00F7-0N1	MSN3700	NVIDIA Spectrum-2 based 100GbE 1U Open Ethernet Switch with ONIE, 32 QSFP28 ports, 2 Power Supplies (AC), x86 CPU, standard depth, P2C airflow, Rail Kit	EOL

9 Release Notes History

9.1 Changes and New Feature History

⚠ This section includes history of changes and new feature of 3 major releases back. For older releases history, please refer to the relevant firmware versions.

Feature/Change	Description
40.49.1014	
PCC/ZTR-RTT Congestion-Control Histogram Collection	Added support for Congestion-Control histogram collection in the PCC/ZTR-RTT algorithm. After enabling this capability, customers can read RATE and RTT histogram counters for PCC-managed flows.
ZTR_RTTCC Tunable Probe Timeout	Added a new parameter to the ZTR_RTTCC algorithm to define the probe-packet timeout threshold.
Port LEDs	Firmware now supports two LEDs per transceiver module on the GB300 OSFP board. When configured for multiple LEDs per module, LED control (including the Module LED Control Register and host-visible indications) drives the correct LED for the relevant module and lane. Boards with a single LED per module are unchanged when **num_leds_per_module** is set to 1.
Link up/Down dmesg Messages	When powering down the OSFP port, the link up/down dmesg messages were previously sent only to the primary ASIC, they are now sent to the secondary ASIC as well.
PSP Transport Packet Reformat	Added support for a new packet reformat type for PSP transport packets. This reformat removes the PSP trailer and the UDP + PSP transport headers, and updates the IP Protocol field based on the PSP next_header value.
Persistent CRDT Token	CRDT token behavior has changed and is now persistent. Previously, debug firmware was allowed only if a CRDT token was applied temporarily, meaning it was erased after reset, or if another debug firmware was already running. Once a user moved to production firmware, a new token was required to switch back to debug firmware. With the new behavior, the token is saved to flash, meaning it remains readable via MDSR even after reset. Once installed, the token allows the user to move between debug and production firmware until it is explicitly revoked using the MDSR set command.
Extended Recovery Support for XDR Optics	Added support for extended recovery on XDR optics when the peer device also supports the capability. This enhancement improves link recovery behavior by allowing both sides of the connection to use the extended recovery mechanism, helping increase resiliency in supported XDR optic configurations.
Bug Fixes	See *Bug Fixes in this Firmware Version* section.

Feature/Change	Description
40.48.1000	
DOCA PCC	The DOCA PCC NP application now enables the NIC to insert the RTT response transmit timestamp in hardware, reducing software-induced jitter and improving the accuracy and consistency of RTT measurements.
DOCA PCC API: DSCP Query for PCC Flows in QP Mode	Introduces a DOCA PCC device API that enables retrieval of the DSCP value associated with a PCC flow when PCC operates in QP mode (for example, when `ROCE_CC_SHAPER_COALESCE_P1=2` and `ROCE_CC_SHAPER_COALESCE_P2=2` are configured via mlxconfig).
DPA Process Limit Update	The system-wide limit for DPA processes has been reduced to **30**. This total includes both user processes across all GVMIs and internal ProgCC processes. The `max_dpa_processes` value reported to the user is calculated as: `max_dpa_processes=30-number_of_progcc_processes`
Host Rate Limiting Support Above 255 Gbps	Host rate limiting has been extended to support bandwidth values above 255 Gbps. To remove the previous cap, a new `max_bw_value_msb` field was added to `est_global`, providing additional MSB bits to represent higher bandwidth values. With this enhancement, firmware and host tooling can correctly configure and report rate limits beyond 255 Gbps on high-speed links.
PLDM PDR Repository Change Event Support	PLDM now supports the PDR Repository Change event type, enabling notification to the BMC when PDRs change. With this flow, the BMC can detect cable insertion/removal events. Refer to DSP0248 for details.
BMC Write-Protection Check on Firmware Update Failure	Added a validation step during firmware updates to detect whether the BMC is asserting write protection, helping diagnose and prevent update failures.
Parallel Save/Load Support for VF Migration	Added support for running save and load operations in parallel, enabling multiple contexts (e.g., multiple VFs) to be checkpointed and restored concurrently instead of serially. This reduces overall migration time and improves scalability in environments that need to migrate or recover many VFs at once.
NVGRE VSID Modify-Header Support	Extended packet modify-header operations to support `set` and `copy` actions on the NVGRE VSID (Virtual Subnet Identifier). A new field, `TUNNEL_HDR_DW_2` (0x84), enables dynamic VSID modification, adding header rewrite support for NVGRE tunnel traffic in addition to existing filtering capabilities.
mlxlink	mlxlink `show_links` now reports the full PCIe identifier (domain/segment + BDF), improving device-to-link mapping and avoiding ambiguous/duplicate BDF entries on multi-domain systems.
Bug Fixes	See *Bug Fixes in this Firmware Version* section.

Feature/Change	Description
40.47.1088	
Bug Fixes	See *Bug Fixes in this Firmware Version* section.

Feature/Change	Description
40.47.1026	
Lane Margin	Lane Margin is a signal integrity diagnostic feature that measures the electrical “eye margin” of high-speed serial lanes, the physical data paths that carry bits over interfaces like PCIe, SerDes, or Ethernet links.
PCIe Trace Function	Added a new NVLOG TLV type to support PCIe logger functionality. This enhancement enables logging and debugging of PCIe-related events through the NVLOG infrastructure, improving traceability and issue analysis.
Passing Metadata Registers between the NIC Layer and the E-Switch (esw) Layer	This enhancement enables seamless metadata propagation across layers, allowing flow steering rules and packet processing logic to share contextual information such as flow identifiers, source context, or policy tags. It improves coordination between NIC and E-Switch pipelines, enabling more flexible traffic handling and advanced offload capabilities.
DPA (Data Path Accelerator) Partition Creation	Access control was added to ensure that only the VHCA instance that created a DPA partition is permitted to modify or delete it.
DPA Manifest	A new DPA Manifest mechanism was introduced to define and manage application permissions.
DPA TIMER	DPA TIMER functionality has been exposed through the MTCTR access register, allowing direct access by applications.
Parallel Suspends of VFs	Added support for parallel suspend operations across multiple VFs.
RTT RTC Timestamp	Added support for using the real-time clock to fill the request and response timestamps in hardware-generated RTT packets. To enable this feature, set `REAL_TIME_CLOCK_ENABLE` in **mlxconfig** and configure `ROCE_CC_RTT_TIMESTAMP_FORMAT` to `0x02 (REAL_TIME)`. For additional details, see Known Issue 4496642 in the Known Issues section.
PCC NP IFA2 GNS	Enables customers to specify the corresponding GNS values that will be forwarded to the DOCA PCC NP feature. When multiple slots are configured with IFA2, the GNS settings in `pcc_config` and `pcc_np_config` must be identical across all slots using IFA2.
Enhanced Error Recovery for GGA QPs	When a GGA QP encounters a memory access (address translation) issue in one VM or Function, it no longer enters an error state. Instead, the QP now recovers from the error, sends an error CQE to the software, and continues serving other VMs and Functions. Unlike RDMA QPs, the error CQE may redundantly reference a valid mkey, therefore, the software should re-construct all mkeys that received error CQE notifications.
Enable/Disable ECN in Upstream	Added the ability to enable or disable ECN in the upstream by allowing the MODIFY_CONG_STATUS and QUERY_CONG_STATUS commands in mlx5_fwctl.
Link Speed per-lane	Enabled 50G per-lane link speed and improved LED behavior for clearer network status indication. Traffic LED now blinks when traffic is active and reflects accurate link status. Scan chain secondary LED behavior fixed, now correctly reflects link activity instead of remaining constantly on 900-9X85E-00EX-MJA.

Feature/Change	Description
40.47.1026	
API to Write PSP Master Key	Added a new API to write PSP Master Key. This API allows writing a new PSP Master Key, which will be used to generate new SPI/key pairs. The previous key remains valid for decryption until the key rotation process is completed.
ADP-RETX Timeout Profile	Firmware now allows the ADP-RETX timeout profile to be configured even when there are open QPs.
PCI Logs	PCI logs are now reported via the existing NC-SI OEM command `Get Log Info` (Command = 0x0, Parameter = 0x2F).
Adaptive Hot-plug System (AHS)	Added support for Adaptive Hotplug System (AHS) alongside the existing NHP solution, enhancing hotplug flexibility and system adaptability.
Additional STE Action	The ASO object pointer size has been increased from 24 bits to 32 bits, eliminating the previous limitation of ~16 million ASO objects per GVMI and enabling significantly greater scalability for future expansions.
NV Configuration	Added an NV configuration option to allow disabling XDR. **Note:** Disabling SDR or enabling configurations not supported by the INI file remains unsupported.
MVCAP (Multi-Version Capability)	Added support for MVCAP (Multi-Version Capability) functionality enabling improved compatibility and version management across multiple components.
Bug Fixes	See *Bug Fixes in this Firmware Version* section.

Feature/Change	Description
40.46.3048	
Security Hardening Enhancements	This release contains important reliability improvements and security hardening enhancements. NVIDIA recommends upgrading your devices firmware to this release to improve the devices' firmware security and reliability.

Feature/Change	Description
40.46.1006	
PCIe TLP Processing Hints (TPH) and Steering Tag (ST)	Enabled PCIe TLP Processing Hints (TPH) and Steering Tag (ST) during MKey creation. **Note:** The steering tag index in the MKey creation must reference an MSIX entry containing the actual steering tag value.
PCIe Congestion Events	Added support for the general PCIe congestion object to monitor and receive events related to inbound and outbound PCIe congestion. A threshold can be configured to specify when the firmware should send an event to the software. This capability is activated by setting the mlxconfig parameter `PCIE_CONGESTION_MONITOR` .

Feature/Change	Description
40.46.1006	
RDMA QP	When an RDMA QP encounters a memory access an issue caused by address translation, it can recover without transitioning to an error state. The QP will send an error CQE to notify the software while continuing to serve other VMs and functions.
PPCNT Counters	Firmware now supports new counters in the PPCNT register to track multicast and unicast packets transmitted and received. The counters include: • `port_multicast_xmit_pkts_high` • `port_multicast_xmit_pkts_low` • `port_multicast_rcv_pkts_high` • `port_multicast_rcv_pkts_low` • `port_unicast_xmit_pkts_high` • `port_unicast_xmit_pkts_low` • `port_unicast_rcv_pkts_high` • `port_unicast_rcv_pkts_low`
Safely Identify DPUs/SmartNICs is a Machine and PCIe Slot	A new access register is introduced that accepts a type, length, and R/W command. • **Write operation:** Allocates a new ICMC buffer of the specified size (aligned to 64B) and stores the provided data. If a buffer for the given type already exists, the data in the ICMC is overwritten, and the locked area is adjusted accordingly • **Read operation:** If a buffer exists, its data is copied out. If not, the access register returns a size of 0 or an explicit error The length can be stored within the data in the ICMC, and the type is mapped to 256B chunks (due to access register limitations), so the VA of the buffer is calculated as `(base + (type << 8))`. The first 4 bytes store a validity flag and the length. If length storage is unnecessary (e.g., null-terminated data), a hardware read can use a cache-line hit as a validity bit. This feature is designed for limited use cases and does not address multi-host scenarios or broader ICMC utilization implications.
Latency Histogram Counter	Introduced a new latency histogram counter that measures the distribution of read operation latencies from our device to the PCI link, providing better visibility into PCI read performance and potential bottlenecks.
Incoming NC-SI Messages Validation for the payload_len Field	Added an extra validation for the payload_len field in incoming NC-SI messages. Previously, invalid packets might have been accepted; now, such packets are silently dropped.
RSS with Crypto Offload	Added support for RSS with crypto offload enabling the NIC to parallelize packet processing across CPU cores while performing encryption/decryption in hardware. Additionally, introduced a new l4_type_ext parameter with values: 0 (None), 1 (TCP), 2 (UDP), 3 (ICMP).
SPDM	Updated SPDM measurements report to version 1.1.
Bug Fixes	See *Bug Fixes in this Firmware Version* section.

9.2 Bug Fixes History

⚠ This section includes history of bug fixes of 3 major releases back. For older releases history, please refer to the relevant firmware versions Release Notes.

Internal Ref.	Issue
4869576	**Description:** Fixed the rate limiter for XDR clusters.
	Keywords: Rate limiter
	Detected in version: 40.44.1036
	Fixed in Release: 40.49.1014
4976460	**Description:** Fixed an issue that caused low performance when the link was not saturated in many-to-one traffic scenarios by optimizing the plane load balancer.
	Keywords: Performance
	Detected in version: 40.46.3048
	Fixed in Release: 40.49.1014
4812446 / 4917369	**Description:** Fixed an issue where CREATE_PARSE_GRAPH_NODE incorrectly applied validation intended for WA mode, in which `header_length_field_offset` is adjusted by firmware. Since ConnectX-8 and newer adapters use normal mode only, these checks were redundant and could reject valid input. Validation is now mode-specific: normal-mode rules are applied on ConnectX-8 and newer adapters, while WA-mode rules are applied on older adapters.
	Keywords: Parse graph node
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4812446	**Description:** Fixed an issue where firmware did not enforce the adapter limit for the number of flow match samples per parse graph node across the device. As a result, creation could succeed even when the total exceeded the hardware-supported limit. Firmware now enforces this limit before allocating samples. Creating a parse graph node that would exceed the maximum number of hardware flow match samples now fails with a “no resources” error.
	Keywords: Parse graph node
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014

Internal Ref.	Issue
4864823 / 4858750	**Description:** Fixed an issue where, when Flex Parser overwrite native arc was enabled, destroying or closing a parse graph node on a native protocol arc did not restore the native protocol parser. As a result, the native parser remained unavailable on the device until a reset or another recovery action was performed. Native protocol parsing is now restored as expected when the parse graph node is closed or destroyed. Native parsers are disabled only while the Flex Parser graph owns the arc.
	Keywords: Flex Parser
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4980585 / NVBug 6084453	**Description:** Fixed an issue where NSM Get-Port-Network-Addresses returned an invalid format on B200/B300 HMC. It now follows the documented non-compact format.
	Keywords: NSM Get-Port-Network-Addresses
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4928056 / 4999011 / 4999400	**Description:** Fixed a corner-case race condition triggered by DSP resets, causing the link to hang.
	Keywords: DSP reset
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4964566 / 4957757	**Description:** Fixed a DEAD IRISC assert that could occur during TLV NV_DATA flash access by suspending the watchdog while waiting for flash IPC (until timeout), preventing the assert on TLV access.
	Keywords: DEAD IRISC assert
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4835832	**Description:** In rare cases, certain module types may experience link-up failures.
	Keywords: Cables
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4859294	**Description:** Fixed an issue where, after toggling all ports, one port could become stuck in the Receiver Detect state.
	Keywords: Port toggling
	Detected in version: 40.48.1000

Internal Ref.	Issue
	Fixed in Release: 40.49.1014
4657767 / 4658776 / 4874764 / 4874765	**Description:** Fixed an issue where repeatedly writing NVCONFIG TLVs could cause excessive NV_DATA partition swaps during garbage collection. This rapid cycling could accelerate flash wear (end-of-life at 100,000 erases) and potentially render the device inoperable. Firmware now avoids unnecessary physical writes by returning OK when the requested configuration already exists in flash, and increases the maximum supported NV_DATA partition swaps from 100,000 to 200,000.
	Keywords: NVCONFIG TLVs
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4983638 / NVBug 6090735	**Description:** Fixed an issue where NSM did not function reliably in multihost configurations when a DMA physical function (PF) was present.
	Keywords: NSM, multihost configurations, DMA physical function (PF)
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4871267 / 4871254	**Description:** Fixed an issue where ICM_RES_HW_DMFS_ENCAP_H_FW was allocated per GVMI, preventing some RTTs from using it.
	Keywords: GVMI, RTT
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4860860	**Description:** Fixed an issue where queue pairs (QPs) created during a PCC process transition could miss congestion-control (CC) information, preventing them from being fully managed.
	Keywords: DOCA, PCC, QP, Congestion Control
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4774410 / 4773595	**Description:** Fixed an issue where the LAG layer steering table had only one strict-SQ entry, causing strict no-port-select traffic to be routed to a single default port. With LOAD_BALANCE_MODE_P1=3, this could lead to link flapping because LACP packets were transmitted only from that default port. The fix expands the LAG layer steering table and adds additional strict-port entries to distribute strict traffic correctly.
	Keywords: Link Flapping
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014

Internal Ref.	Issue
4789601 / 4850200 / NVBug 5736447	**Description:** Fixed an issue where RDMA traffic could stall in large-scale deployments for certain source IP and UDP source-port combinations when DOCA PCC was active and no congestion-control algorithm was configured in algorithm slot 0.
	Keywords: RDMA, DOCA PCC, Congestion Control Algorithm
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4796182	**Description:** Fixed an issue where the live migration target did not receive a port state change event on the resume VHCA command. The target now generates this event so software that depends on port state is notified of any changes.
	Keywords: Live migration
	Detected in version: 40.47.1026
	Fixed in Release: 40.49.1014
4859742 / 4847702	**Description:** Fixed an issue during hitless upgrade where, after the SACK generation/handler fence, firmware could mark the old port configuration ID as invalid. If SACK causes were still active, a race could cause SACK ISRs to stop unexpectedly. The fix is to always return BUSY while handover is active.
	Keywords: Hitless upgrade
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4873533 / 4947034	**Description:** Fixed an issue where, when using multiplane ZTRCC congestion control with multiple flows, the RTT timeout counter in the PPCC register could increase even when the network was not congested.
	Keywords: Congestion control, multiplane, SPCX CC, RTT, PPCC
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4876790 / 4822829	**Description:** Fixed an issue where a firmware race between packet receive and QP cleanup could move a QP to error when reopening the same QP and sending the same MSN. This could occur when interrupting traffic (e.g., Ctrl+C) and running many iterations until the same QP/MSN combination is reused.
	Keywords: Firmware race
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4905774	**Description:** Fixed an issue where the FMT Static PF verifier did not account for the Tools PF when verifying the DMA PF.

Internal Ref.	Issue
	Keywords: FMT Static PF verifier
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014
4683339 / 4780301 / 4895260	**Description:** Fixed an issue where QPs established before loading DOCA PCC could exhibit inconsistent algorithm-selected behavior between ports in LAG mode after DOCA PCC is loaded.
	Keywords: Congestion Control, DOCA PCC
	Detected in version: 40.45.1020
	Fixed in Release: 40.49.1014
4843342 / 4970550	**Description:** Fixed an issue where one-to-one RoCE traffic using a single QP might not achieve line rate on some platforms when using the default ROCE_CC_COMPATIBILITY_MODE setting in mlxconfig.
	Keywords: Congestion control, PCC, ROCE_CC_COMPATIBILITY_MODE
	Detected in version: 40.47.1026
	Fixed in Release: 40.49.1014
4892822 / 4638811	**Description:** Fixed an issue where, in some configurations with ConnectX-8 connected to an MP3 switch, th e link might fail to come up on the switch side after toggling ports. In these cases, ConnectX-8 reported opcode 14, indicating detected remote faults and that the partner was not bringing t he link up.
	Keywords: Link down, opcode 14, AC/DC cycles
	Detected in version: 40.47.1026
	Fixed in Release: 40.49.1014
4783261 / 4658799	**Description:** Fixed an issue where ACS errors could occur due to a race condition when performing LDE/ SBR while MCTP traffic to the device was active.
	Keywords: ACS, MCTP
	Detected in version: 40.47.1026
	Fixed in Release: 40.49.1014

Internal Ref.	Issue
4881141 / 4882127 / 4882128	**Description:** Fixed a potential routing error when IOVAs assigned to the NIC overlap with the PCIe address space. In certain configurations, an IOVA could fall into an “unclaimed address” range, within a PCIe switch’s Upstream Port (USP) window but outside any Downstream Port (DSP) aperture. **Note:** Since this is a kernel issue, to ensure packets are routed correctly in these scenarios, users must enable the ACS Unclaimed Request Redirect bit in the PCIe bridges’ Access Control Services (ACS) capability via the kernel.
	Keywords: ACS Unclaimed Request Redirect, IOVA
	Detected in version: 40.47.1026
	Fixed in Release: 40.49.1014
4952219 / NVBug 6026050	**Description:** Fixed an issue where partially enabling ACS enhanced capability could cause unexpected driver and system behavior. ACS enhanced capability remains disabled and will not be enabled.
	Keywords: ACS
	Detected in version: 40.48.1000
	Fixed in Release: 40.49.1014

Internal Ref.	Issue

4823907 / NVbug 5742181	**Description:** Fixed an issue where, in certain configurations with the ConnectX-8 PCIe switch enabled, downstream devices (including GPUs) might not be detected and could drop from the PCI bus, with GPU sensors/properties reporting nan. This was caused by the device not receiving the required PERST# assertion during initialization, and was seen only when PCIe settings were manually modified via mlxconfig (e.g., restricting link speed/width or ASPM on specific PCI buses).
	Note: On legacy firmware, additional configuration steps may still be required, as detailed below.
	If you cannot update the firmware immediately, you can restore device detection using one of the following options:
	• **Option 1:** Reset configuration
	  • Reset the device configuration to defaults:
	`mlxconfig -d <device> -y reset`
	• Option 2: Manual PERST configuration (B300)
	  • Manually configure the PERST parameters:
	`mlxconfig -d <device> set PCI_BUS10_CONTROL_EN=1`
	`mlxconfig -d <device> set PCI_BUS10_PERST_SOURCE=2`
	`mlxconfig -d <device> set PCI_BUS10_PERST_GPIO=8`

	Keywords: ConnectX-8 PCIe, GPU, PERST# assertion
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000

Internal Ref.	Issue
4786813	**Description:** Fixed an issue where the DPA kernel used unsafe ICM access during process creation/modification, which could cause the DPA kernel to hang during FLR.
	Keywords: DPA kernel, FLR
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4804664 / 4806969	**Description:** Fixed an issue in the User Debugger “query caps” where it returned only the number of capabilities, not the capability bitmap.
	Keywords: User Debugger “query caps”
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4813862 / 4146077	**Description:** Fixed an issue where CR dumps could time out when accessing xpl_top addresses across all three pcores.
	Keywords: CR dump
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4833440	**Description:** Fixed an issue where the Virtio and NVMe EMU_MNG settings were exposed incorrectly, which could cause confusion when using mlxconfig.
	Keywords: Virtio and NVMe emulation, mlxconfig
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4756451	**Description:** Fixed an issue where the PHY LED could show green during the initializing state when active speed was set to full speed. In IB mode, the initializing-state LED should be amber only.
	Keywords: PHY LED
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4768546	**Description:** Fixed an issue where, on multi-PF-per-port systems, a PF FLR could impact the traffic bandwidth of other PFs on the same port.
	Keywords: PF FLR
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000

Internal Ref.	Issue
4705918	**Description:** Fixed an issue where PTP could converge to an incorrect time/offset and report an inaccurate path delay.
	Keywords: PTP
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4657792 / NVbug 5567725	**Description:** Fixed an issue where, in Flit Mode, the device could become unresponsive when receiving malformed or invalid traffic from a link partner.
	Keywords: Flit Mode
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4484662	**Description:** Fixed an issue where mlxlink reported 0 values for SNR (media and host) due to incorrect local port mapping in firmware and an incorrect page number used by MFT.
	Keywords: mlxlink
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4621747 / 4792742 / 4794290 / NVbug 5502241	**Description:** Fixed an issue where parallel accesses to the MCIA register could return incorrect data. In some hosts running `ethtool -m <interface>` repeatedly (e.g., once per second), this could intermittently report `Identifier: 0x00` (unknown/no module), causing health checks to fail.
	Keywords: MCIA register
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4686284 / NVbug 5607036	**Description:** Implemented IB extended port telemetry counters via the NSM Type 1 Get Port Telemetry Counters command, adding counters 19 and 20: `NSM_LINK_ERROR_RECOVERY_COUNTER_CNTR_ID` and `NSM_LINK_DOWNED_COUNTER_CNTR_ID`.
	Keywords: IB extended port telemetry counters, NSM
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4758174 / NVbug 5698200	**Description:** Fixed a rare attestation certificate signature formatting issue by removing an unnecessary leading zero byte in the “r” or “s” value.
	Keywords: Attestation certificate signature format
	Detected in version: 40.47.1026

Internal Ref.	Issue
	Fixed in Release: 40.48.1000
4532684 / 4635872 / 4794865 / 4794866 / 4794867 / NVbug 5385446	**Description:** Fixed an issue by improving the ADP-RETX algorithm to avoid re-arming without performing a retransmission.
	Keywords: ADP-RETX algorithm
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4542516 / 4554220	**Description:** Fixed an issue where, in certain Gen6 setups, RDMA READ bidirectional traffic required at least 5 QPs to reach full wire speed.
	Keywords: RDMA READ bidirectional traffic
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4554763 / 4808657	**Description:** Fixed an issue affecting single-process, unidirectional RDMA READ to GPU memory (4 QPs, 128KB messages) by enabling **`ZERO_TOUCH_TUNING_ENABLE`** via MLXCONFIG.
	Keywords: Zero Touch Tuning, mlxconfig
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4608214	**Description:** Fixed an issue where probe packets might not be sent under heavy traffic.
	Keywords: PCC, ZTR_RTTCC, probe
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4450570 / 4780432 / 4780433	**Description:** Fixed an issue where the root complex sent MCTP-over-PCI messages before a BDF was assigned, causing responses to be sent with BDF 0. The fix ensures that MCTP messages routed by ID are ignored until a valid BDF is assigned.
	Keywords: MCTP-over-PCI, BDF, MCTP messages
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000
4809134 / 4824635	**Description:** Fixed an issue where the steering tables were not updated after enabling partial Spectrum-X capabilities (BTH.AR) via LLPD.
	Keywords: Steering tables, LLDP
	Detected in version: 40.47.1026

Internal Ref.	Issue
	Fixed in Release: 40.48.1000
4797308 / NVbug 5706024	**Description:** Fixed an issue where an intX message was sent with a Requester ID of 0, causing an ACS violation at the root port. The fix uses the correct BDF as the Requester ID instead of 0.
	Keywords: intX message
	Detected in version: 40.47.1026
	Fixed in Release: 40.48.1000

Internal Ref.	Issue
4608544	**Description:** Fixed an issue where, in rare live migration scenarios, a delayed doorbell triggered a false timeout alarm.
	Keywords: Live migration, doorbell, timeout alarm
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1088
4648642	**Description:** Fixed a rare issue in which destroying PCC NP configuration objects could result in assert 0x8175 being logged in dmesg.
	Keywords: Assert 0x8175, PCC NP
	Detected in version: 40.47.1026
	Fixed in Release: 40.47.1088
4655971	**Description:** Fixed the PCIe counters to correctly report event values in nanoseconds.
	Keywords: DOCA Telemetry Diagnostics
	Detected in version: 40.47.1026
	Fixed in Release: 40.47.1088
4690503	**Description:** Fixed an issue where creating a DPA process that uses 128 MB of data caused the dynamic library to fail with syndrome 0xdc30ac. The BSS section of the DPA application is now limited to 64 MB.
	Keywords: DPA process, BSS
	Detected in version: 40.47.1026
	Fixed in Release: 40.47.1088

Internal Ref.	Issue
4884739	**Description:** Resolved an issue where link failures could occasionally occur at PAM4 speeds over optical interfaces in rare cases.
	Keywords: PAM4 speeds, optical interfaces
	Detected in version: 40.46.1006

Internal Ref.	Issue
	Fixed in Release: 40.47.1026
4570205	**Description:** Fixed a firmware issue where the ZTR_RTTCC algorithm parameters AI and HAI did not support a sufficient range.
	Keywords: PCC, ZTR_RTTCC
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026
4629077	**Description:** Fixed an issue where coalescing regular SX events with SX RTT events under ZTR_RTTCC could keep improper event fields, which could impact congestion control behavior.
	Keywords: PCC, ZTR_RTTCC
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026
4683328	**Description:** Fixed an issue in the ZTR_RTTCC algorithm where probe-abortion handling could behave improperly under high-stress network conditions, ensuring proper congestion control and stable traffic performance.
	Keywords: PCC, ZTR_RTTCC
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026
4501554	**Description:** Fixed an assertion failure that could occur with the E-Switch uplink in specific configurations where the e-switch was disabled and Path Migration was active or GVMIs were using SRQ loopback in SQs. The issue occurred because the firmware attempted to perform cleanup operations when the uplink configuration lacked sufficient capacity. Now, when the E-Switch is disabled and no actions are available in the uplink STE, the firmware connects to the uplink STE instead of copying it.
	Keywords: Path migration, steering
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026
4506854	**Description:** Added Scaling Factor "read" field. To obtain correct values in mlxlink, MFT version 4.33.0 or later is required.
	Keywords: Scaling Factor, mlxlink, MFT
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026

Internal Ref.	Issue
4468319	**Description:** Fixed an issue where the ConnectX-8 downstream port failed to send a NACK when rejecting an L1 entry request from the upstream port.
	Keywords: NACK, downstream port
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026
4550782	**Description:** Fixed an issue on GB200 systems with two symmetrical ConnectX-8 SuperNICs, which caused DPN numbering differences on the HCA upstream port. Legacy drivers accessed with dpn=0,0,0, which could result in attempts to access the wrong DPN node in socket-direct systems. The firmware now automatically determines the correct pcie_index based on the accessed link in direct-NIC systems.
	Keywords: DPN numbering
	Detected in version: 40.45.1020
	Fixed in Release: 40.47.1026
4571079	**Description:** Fixed an issue where invoking the resourcedump tool with segment type DPA_PROCESS_LST returned invalid data when the parameter n1 == 1 and no processes existed on the current vhca_id. The fix adds a proper check, and the resourcedump tool now reports the correct error in this scenario.
	Keywords: DPA PROCESS, RESOURCE DUMP
	Detected in version: 40.45.1020
	Fixed in Release: 40.47.1026
4529293	**Description:** Fixed an issue where, during failover or restart, the SM sending a PortInfo MAD to the HCA firmware triggered reinitialization of port buffers, momentarily halting ingress traffic and causing packet drops. The firmware now avoids reconfiguring port buffers when the new configuration matches the current one.
	Keywords: OpenSM
	Detected in version: 40.45.1020
	Fixed in Release: 40.47.1026
4641215	**Description:** Fixed a rare issue where MFRL operations could fail due to a timeout.
	Keywords: MFRL
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026

Internal Ref.	Issue
4683346	**Description:** Fixed an issue where, under the ZTR_RTTCC algorithm, a flow that reached its minimum rate due to heavy congestion would not recover its rate once the congestion cleared.
	Keywords: PCC, ZTR_RTTCC
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026
4575692	**Description:** Fixed an issue where a missing interrupt from the module IO (Expander) could prevent the module from being raised.
	Keywords: Module IO (Expander)
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026
4620765	**Description:** Fixed an issue where reading debug registers could cause link BER (Bit Error Rate) degradation over time.
	Keywords: BER
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026
4658434	**Description:** Fixed an issue where ports connected via 4 or 8 lanes and configured for 200G_2x (using only 2 lanes) would fail to link when using a mix of new firmware (with “Non Tx-Squelch” support) and older firmware versions. **Note:** Please make sure on both sides, switch (local device) and Ssitch/NIC (peer device) you: • Deploy the new firmware release versions as a matched bundle on both Switch and NIC devices. • Configure the port to use 2 lanes (instead of 4 or 8 lanes) while keeping the 200G_2x speed setting.
	Keywords: Port speed
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026
4401684	**Description:** Fixed an issue in Arch diagnostic data counters where the pcie_link_outbound_data_bytes counter was incorrectly returning only zero values.
	Keywords: Arch diagnostic data counters
	Detected in version: 40.45.1020
	Fixed in Release: 40.47.1026
4575696	**Description:** Fixed an issue where multiple long-running process registers could cause aborted access and timeouts, the internal state is now properly handled.

Internal Ref.	Issue
	Keywords: ibdiagnet2
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026
4583940	**Description:** Fixed an issue where enabling the CCMAD custom header on one PCC probe slot caused other slots to malfunction when multiple slots were configured. **Note:** If using firmware versions older than the 40.47.10xx GA release, disable the CCMAD custom header when multiple probe slots are enabled.
	Keywords: PCC CCMAD custom header
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.1026
4208960	**Description:** A packet may be parsed incorrectly, if a driver uses the `header_length_field_mask` when creating a `PARSE_GRAPH_NODE` object, and the mask value is not composed of continuous bits or does not commence at the least significant bit.
	Keywords: PARSE GRAPH NODE, Flex Parser
	Detected in version: 40.44.0208
	Fixed in Release: 40.47.1026
4610740	**Description:** Fixed a firmware issue where a CQE error with vendor_syndrome RDE_MAL_WQE (0xd6) could cause traffic disruption on the affected QP.
	Keywords: RDMA, transport
	Detected in version: 40.45.1020
	Fixed in Release: 40.47.1026

Internal Ref.	Issue
4603774	**Description:** Fixed an issue where the adapter card could drop NC-SI over MCTP commands when padding bytes were present after the NC-SI checksum.
	Keywords: NC-SI
	Discovered in Version: 40.46.1006
	Fixed in Release: 40.46.3048

Internal Ref.	Issue
4286902	**Description:** Fixed a race condition in DPA process termination during the exception flow, where a failed process could be missed and not reported to the user.

Internal Ref.	Issue
	Keywords: DPA
	Detected in version: 40.45.1020
	Fixed in Release: 40.46.1006
4401109	**Description:** Fixed an issue where RTTs on IFA1 were not sent when IFA1 and IFA2 were configured in cumulative slots.
	Keywords: PCC, multi probe, IFA
	Detected in version: 40.45.1020
	Fixed in Release: 40.46.1006
4486431	**Description:** Fixed an issue where issuing multiple parallel queries of DPA_THREAD objects with the same object ID could fail.
	Keywords: DPA
	Detected in version: 40.45.1020
	Fixed in Release: 40.46.1006
4443601	**Description:** Fixed a firmware issue where PXE failed to boot when both LAG ports were up.
	Keywords: PXE, LAG
	Detected in version: 40.45.1020
	Fixed in Release: 40.46.1006
4475307	**Description:** Fixed an issue where PCC DCQCN used incorrect parameter values when link speed was 400Gbps or higher.
	Keywords: PCC DCQCN, congestion control.
	Detected in version: 40.45.1020
	Fixed in Release: 40.46.1006
4480427	**Description:** Fixed incorrect calculation of start address and mode for the CQE buffer in DPA CQ, which could cause CQEs to be written to the wrong address when the buffer is not 4K-aligned and spans a second page boundary.
	Keywords: CQ, CQE Buffer, DPA
	Detected in version: 40.45.1020
	Fixed in Release: 40.46.1006
4402022	**Description:** Fixed an issue where Wake-on-LAN (WoL) may not function correctly on certain multihost configurations.
	Keywords: Wake-on-LAN (WoL)
	Detected in version: 40.45.1020
	Fixed in Release: 40.46.1006
4445479	**Description:** Added a fixed estimated power value for Hvdd in the INI configuration.
	Keywords: Hvdd
	Detected in version: 40.45.1020
	Fixed in Release: 40.46.1006

Internal Ref.	Issue
4426779	**Description:** Updated the handling of the PLDM Type-5 'Activate Firmware' command to ensure the update flow does not fail when 'self-contained activation' is requested. Although the 'Self Contained Activation Is Not Supported' completion code will still be returned, the component will now be successfully marked as pending.
	Keywords: PLDM
	Detected in version: 40.44.1036
	Fixed in Release: 40.46.1006
4318063	**Description:** When running the PTP4L application, the path delay displays inconsistent values after each fwreset and rerun, resulting in a non-constant PPS offset that fails to meet Class B/C requirements.
	Keywords: PTP
	Detected in version: 40.45.1020
	Fixed in Release: 40.46.1006
4163634	**Description:** When connecting a Quantum-3 switch system (with a port split into 8 ports) to a ConnectX-8 single port SuperNIC, the link will not be established.
	Workaround: Configure the Quantum-3 switch system port to be split into 2 or 4 ports, or set the ConnectX-8 to operate in multiplane mode.
	Keywords: Port split, Quantum-3 **Detected in version:** 40.44.0212
	Fixed in Release: 40.46.1006
4366117	**Description:** Configuring a small MTU leads to fragmentation of packets critical for the PXE boot process. As a result, the PXE boot filters mistakenly discard these packets, causing the PXE boot to fail.
	Keywords: PXE boot filters
	Detected in version: 40.45.1020
	Fixed in Release: 40.46.1006

10 Legal Notices and 3rd Party Licenses

The following are the drivers' software, tools and HCA firmware legal notices and 3rd party licenses.

Product	Version	Legal Notices and 3rd Party Licenses
Firmware	xx.50.1002	• License • 3rd Party Notice • 3rd Party Unify Notice • HCA Firmware EULA
DOCA-Host	3.5.0	• License • 3rd Party Notice • 3rd Party Unify Notice
MFT FreeBSD	4.37.0-154	• License • 3rd Party Notice • 3rd Party Unify Notice
MFT Linux		• License • 3rd Party Notice • 3rd Party Unify Notice
MFT VMware		• License • 3rd Party Notice • 3rd Party Unify Notice
MFT Windows		• License • 3rd Party Notice • 3rd Party Unify Notice

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. Neither NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: “NVIDIA”) make any representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice. Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

NVIDIA