



NVIDIA ConnectX-8 SuperNIC Firmware Release Notes v40.47.3576 (2025 LTS U3)

Table of Contents

1	Release Notes Update History	5
2	Overview	6
2.1	Firmware Download.....	6
3	Firmware Compatible Products.....	7
3.1	Supported Devices	7
3.2	Driver Software, Tools and Switch Firmware	7
4	Changes and New Features.....	9
4.1	Customer Affecting Changes	9
4.1.1	Changes in This Release	9
4.1.2	Changes Planned for Future Releases.....	9
4.1.3	Changes in Earlier Releases	9
4.1.4	Discontinued Features	10
4.2	Declared Unsupported Features.....	10
5	Bug Fixes in this Firmware Version	12
6	Known Issues	14
7	PreBoot Drivers (FlexBoot/UEFI)	18
7.1	FlexBoot Changes and New Features.....	18
7.2	UEFI Changes and Major New Features	18
8	Validated and Supported Cables and Switches	19
8.1	Validated and Supported Cables and Modules	19
8.1.1	Cables Lifecycle Legend	19
8.1.2	XDR / 800GbE / 1600GbE Cables.....	20
8.1.3	NDR / 400GbE / 800GbE Cables	20
8.1.4	HDR / 200GbE Cables	25
8.1.5	EDR / 100GbE Cables	26
8.1.6	Supported 3rd Party Cables and Modules	27
8.1.7	Tested Switches	28
8.1.8	XDR / 800GbE Switches	28
8.1.9	NDR / 400GbE Switches	28
8.1.10	HDR / 200GbE Switches	29
9	Release Notes History	30
9.1	Changes and New Feature History	30

9.2	Bug Fixes History.....	35
10	Legal Notices and 3rd Party Licenses.....	45



This LTS version has been tested and is fully supported with the following GB versions: GB200 v1.3.10 and GB300 v1.0.10.



This is a long-term support (LTS) release. LTS is the practice of maintaining a software product for an extended period of time (up to three years) to help increase product stability. LTS releases include bug fixes and security patches.

1 Release Notes Update History

Version	Date	Description
40.47.3576	June 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](#).

3 Firmware Compatible Products

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

- InfiniBand - EDR, HDR100², HDR², NDR200², NDR², XDR²
- Ethernet - 25GbE, 50GbE¹, 100GbE¹, 200GbE², 400GbE²
- PCI Express Gen6., supporting backwards compatibility for v5.0, v4.0, v3.0, v2.0 and v1.1

¹. Speed that supports both NRZ and PAM4 modes in Force mode and Auto-Negotiation mode.

². Speed that supports PAM4 mode only.



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 [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	40.47.3576 / 40.47.2682 / 40.47.1088
DOCA-HOST	3.2.3 / 3.2.2 Note: For the list of the supported Operating Systems, please refer to the driver's Release Notes.
WinOF-2	25.10.51000 / 25.10.50020 / 25.10.50020 Note: For the list of the supported Operating Systems, please refer to the driver's Release Notes.
MFT	4.34.1-25 / 4.34.1-18 / 4.34.1-10 Note: For the list of the supported Operating Systems, please refer to the driver's Release Notes.
FlexBoot	3.8.201
UEFI	14.40.10
MLNX-OS	3.12.6000 onwards
Quantum-2 FW	31.2016.2054 onwards
NVOS	25.02.6000 onwards

	Supported Version
Quantum-3 FW (part of NVOS)	35.2016.2080 onwards

4 Changes and New Features



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
40.47.3576	
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.
Bug Fixes	See <i>Bug Fixes</i> 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
40.47.1026	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
4895260 / 4683339	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: 40.47.1026
	Fixed in Release: 40.47.3576
5043712 / 5073671	Description: Fixed an issue where a firmware assert was raised because of a write timeout on the GPIO expander.
	Keywords: GPIO expander
	Detected in version: 40.47.2682
	Fixed in Release: 40.47.3576
5043581 / 4666196	Description: Fixed a PCIe switch initialization issue where upstream traffic was dropped due to ambiguous hardware routing on uninitialized downstream ports, causing downstream devices to disappear from the host.
	Keywords: PCIe switch initialization
	Detected in version: 40.47.2682
	Fixed in Release: 40.47.3576
4970550 / 4843342	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: 40.46.3004
	Fixed in Release: 40.47.3576
4934268 / 4867111	Description: Fixed an issue where RoCEv2 TX traffic could be dropped when running the DOCA spcx_cc application due to stale PCC steering extension entries.
	Keywords: spcx_cc cc_steering_ext
	Detected in version: 40.47.1026
	Fixed in Release: 40.47.3576
4982511 / 4947006	Description: Fixed two rare race conditions in the PCC steering extension when the feature is enabled via path migration, empty RTT, or a custom header.
	Keywords: PCC steering extension

Internal Ref.	Issue
	Detected in version: 40.47.1026
	Fixed in Release: 40.47.3576
5071700 / 4928046	Description: Fixed a corner-case race condition triggered by DSP resets, causing the link to hang.
	Keywords: DSP reset
	Detected in version: 40.47.1088
	Fixed in Release: 40.47.3576
4850603 / 4907678	Description: Fixed a rare race condition in the SUSPEND_VHCA command that could cause the command to hang.
	Keywords: Live Migration
	Detected in version: 40.47.1088
	Fixed in Release: 40.47.3576
4986774 / 4774410	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: 40.47.1026
	Fixed in Release: 40.47.3576

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)

ConnectX-8 has the same feature set and limitations as ConnectX-7 adapter card. For the list of ConnectX-7 Known Issues, please go to <https://docs.nvidia.com/networking/software/adapter-firmware/index.html#connectx-7>.

The below are limitations related to ConnectX-8 only.

Internal Ref.	Issue
4919262 / 4901395	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. If the BIOS initiates a new TFTP session using different source and destination UDP ports, the UEFI driver fails to update the firmware with these new ports, resulting in dropped packets.
	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: 40.47.3576
5059333	Description: The DC transport service does not function on MGX systems with ConnectX-8 when using an Ethernet link type.
	Workaround: N/A
	Keywords: DC, MGX, ETH
	Detected in version: 40.47.3576
4892822 / 4638811	Description: In some cases where ConnectX-8 is connected to an MP3 switch, the link may not come up on the switch side after toggling the ports. ConnectX-8 reports opcode 14, indicating it is detecting remote faults from the partner and that the partner is not raising the link.
	Workaround: N/A
	Keywords: Link down, opcode 14, AC/DC cycles
	Detected in version: 40.47.2682
4884739	Description: Link failures may occasionally be observed at PAM4 speeds over optical interfaces in rare cases.
	Workaround: Toggle the port (admin down/up).
	Keywords: PAM4 speeds, optical interfaces

Internal Ref.	Issue
	Detected in version: 40.47.1026
4604969	Description: Probe packets might be dropped at the transmission stage when multiple congestion control flows are active.
	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

Internal Ref.	Issue
	Detected in version: 40.47.1026
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
4270913	Description: When using DC InfiniBand transport, SRQ usage requires distributing QPs across four SRQs. This means assigning each QP to its own SRQ when there are up to four QPs.
	Workaround: If there are more than four QPs, they must be distributed equally among the four available SRQs.
	Keywords: DC InfiniBand transport, SRQ, QPs
	Detected in version: 40.45.1200
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 <code>mlxconfig ROCE_CC_PRIO_MASK_P\$port</code> or <code>sysfs ecn/roce_rp/enable/\$port</code> .
	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 <code>mlxfwreset</code>).

Internal Ref.	Issue
	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](#)
 - [8.1.2 XDR / 800GbE / 1600GbE Cables](#)
 - [8.1.3 NDR / 400GbE / 800GbE Cables](#)
 - [8.1.4 HDR / 200GbE Cables](#)
 - [8.1.5 EDR / 100GbE Cables](#)
 - [8.1.6 Supported 3rd Party Cables and Modules](#)
 - [8.1.7 Tested Switches](#)
 - [8.1.8 XDR / 800GbE Switches](#)
 - [8.1.9 NDR / 400GbE Switches](#)
 - [8.1.10 HDR / 200GbE 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

IB Data Rate	Eth Data Rate	NVIDIA SKU	Marketing Description	LifeCycle Phase
XDR	1600GE	980-9IAM1-00X001	NVIDIA Active copper cable, 1600Gbps to 1600Gbps, OSFP, 1.1m, RHS to RHS, standard package	P-Rel
XDR	800GE	980-9IAT0-00XM00	NVIDIA single port transceiver for ConnectX-8 Mezz Card, 800Gbps, OSFP, MPO, 1310nm SMF, EML, up to 500m, RHS	EVT

8.1.3 NDR / 400GbE / 800GbE Cables

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	LifeCycle Phase
NA	400GE	980-9I31M-00NM00	MMS4X00-NM400-T	NVIDIA single port transceiver, 400Gbps, 400GbE, OSFP, MPO12 APC, 1310nm SMF, up to 500m, flat top	Prototype
NDR	NA	980-9I924-00N002	MCP7Y00-N002	NVIDIA passive copper splitter cable, IB twin port NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP, 2m	P-Rel
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	P-Rel
NDR	NA	980-9I920-00N02A	MCP7Y00-N02A	NVIDIA passive copper splitter cable, IB twin port NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP, 2.5m	P-Rel
NDR	NA	980-9I73U-000003	MFP7E10-N003	NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 3m	MP
NDR	NA	980-9I73V-000005	MFP7E10-N005	NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 5m	MP
NDR	NA	980-9I57W-000007	MFP7E10-N007	NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 7m	MP
NDR	NA	980-9I57X-00N010	MFP7E10-N010	NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 10m	MP
NDR	NA	980-9I57Y-000015	MFP7E10-N015	NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 15m	MP
NDR	NA	980-9I57Z-000020	MFP7E10-N020	NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 20m	MP
NDR	NA	980-9I570-00N030	MFP7E10-N030	NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 30m	MP
NDR	NA	980-9I570-00N035	MFP7E10-N035	NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 35m	MP
NDR	NA	980-9I570-00N040	MFP7E10-N040	NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 40m	MP

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	LifeC ycle Phas e
NDR	NA	980-9I57 Y-00N050	MFP7E10 -N050	NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 50m	MP
NDR	NA	980-9I57 1-00N003	MFP7E20 -N003	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 3m	MP
NDR	NA	980-9I57 2-00N005	MFP7E20 -N005	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 5m	MP
NDR	NA	980-9I57 3-00N007	MFP7E20 -N007	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 7m	MP
NDR	NA	980-9I55 4-00N010	MFP7E20 -N010	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 10m	MP
NDR	NA	980-9I55 5-00N015	MFP7E20 -N015	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 15m	MP
NDR	NA	980-9I55 6-00N020	MFP7E20 -N020	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 20m	MP
NDR	NA	980-9I55 7-00N030	MFP7E20 -N030	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 30m	MP
NDR	NA	980-9I55 Z-00N050	MFP7E20 -N050	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 50m	MP
NDR	NA	980-9I55 A-00N003	MFP7E30 -N003	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 3m	MP
NDR	NA	980-9I55 B-00N005	MFP7E30 -N005	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 5m	MP
NDR	NA	980-9I58 C-00N007	MFP7E30 -N007	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 7m	MP
NDR	NA	980-9I58 D-00N010	MFP7E30 -N010	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 10m	MP
NDR	NA	980-9I58 E-00N015	MFP7E30 -N015	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 15m	MP
NDR	NA	980-9I58 F-00N020	MFP7E30 -N020	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 20m	MP
NDR	NA	980-9I58 G-00N030	MFP7E30 -N030	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 30m	MP
NDR	NA	980-9I58 O-00N030	MFP7E30 -N040	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 40m	MP
NDR	NA	980-9I58 H-00N050	MFP7E30 -N050	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 50m	MP
NDR	NA	980-9I58 1-00N050	MFP7E30 -N060	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 60m	MP
NDR	NA	980-9I58 2-00N050	MFP7E30 -N070	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 70m	MP

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	Life Cycle Phase
NDR	NA	980-9I58I-00N100	MFP7E30-N100	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 100m	MP
NDR	NA	980-9I58J-00N150	MFP7E30-N150	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 150m	MP
NDR	NA	980-9I58K-00N003	MFP7E40-N003	NVIDIA passive fiber cable, SMF, MPO12 APC to 2xMPO12 APC, 3m	MP
NDR	NA	980-9I58L-00N005	MFP7E40-N005	NVIDIA passive fiber cable, SMF, MPO12 APC to 2xMPO12 APC, 5m	MP
NDR	NA	980-9I58M-00N007	MFP7E40-N007	NVIDIA passive fiber cable, SMF, MPO12 APC to 2xMPO12 APC, 7m	MP
NDR	NA	980-9I58N-00N010	MFP7E40-N010	NVIDIA passive fiber cable, SMF, MPO12 APC to 2xMPO12 APC, 10m	MP
NDR	NA	980-9I56O-00N015	MFP7E40-N015	NVIDIA passive fiber cable, SMF, MPO12 APC to 2xMPO12 APC, 15m	MP
NDR	NA	980-9I56P-00N020	MFP7E40-N020	NVIDIA passive fiber cable, SMF, MPO12 APC to 2xMPO12 APC, 20m	MP
NDR	NA	980-9I56Q-00N030	MFP7E40-N030	NVIDIA passive fiber cable, SMF, MPO12 APC to 2xMPO12 APC, 30m	MP
NDR	NA	980-9I56R-000050	MFP7E40-N050	NVIDIA passive fiber cable, SMF, MPO12 APC to 2xMPO12 APC, 50m	MP
NDR	NA	980-9I601-00N003	MCA4J80-N003-FTF	NVIDIA Active copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 3m, flat to finned	MP
NDR	NA	980-9IA0J-00N002	MCP4Y10-N002-FLT	NVIDIA passive Copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 2m, flat top	MP
NDR	NA	980-9IA0L-00N00A	MCP4Y10-N00A-FLT	NVIDIA Passive Copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 0.5m, flat top	MP
NDR	NA	980-9IA0R-00N01A	MCP4Y10-N01A-FLT	NVIDIA passive Copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 1.5m, flat top	MP
NDR	NA	980-9I433-00N001	MCP7Y00-N001-FLT	NVIDIA passive copper splitter cable, IB twin port NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP, 1m, flat top	P-Rel
NDR	NA	980-9I925-00N002	MCP7Y00-N002-FLT	NVIDIA passive copper splitter cable, IB twin port NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP, 2m, flat top	P-Rel
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	P-Rel

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	Life Cycle Phase
NDR	800GE	980-9I929-00N002	MCP7Y10-N002	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112, 2m, fin to flat	P-Rel
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	P-Rel
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	P-Rel
NDR	800GE	980-9I80C-00N002	MCP7Y40-N002	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xQSFP112, 2m, fin to flat	P-Rel
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	P-Rel
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	P-Rel
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
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	P-Rel
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	P-Rel
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	P-Rel
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	P-Rel
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	P-Rel
NDR	NA	980-9I60O-00N003	MCA4J80-N003-FLT	Active copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 3m, flat top	MP
NDR	800GE	980-9I948-00N004	MCA7J60-N004	NVIDIA active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 4m, fin to flat	P-Rel
NDR	800GE	980-9I949-00N005	MCA7J60-N005	NVIDIA active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 5m, fin to flat	P-Rel
NDR	800GE	980-9I81B-00N004	MCA7J65-N004	NVIDIA Active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112, 4m, fin to flat	P-Rel
NDR	800GE	980-9I81C-00N005	MCA7J65-N005	NVIDIA Active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112, 5m, fin to flat	P-Rel

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	Life Cycle Phase
NDR	NA	980-9IA0G-00N001	MCP4Y10-N001-FLT	NVIDIA Passive Copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 1m, flat top	MP
NDR	NA	980-9IA0H-00N001	MCP4Y10-N001-FTF	NVIDIA Passive Copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 1m, flat to finned	MP
NDR	NA	980-9I432-00N001	MCP7Y00-N001	NVIDIA passive copper splitter cable, IB twin port NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP, 1m	P-Rel
NDR	NA	980-9I92N-00N003	MCP7Y00-N003	NVIDIA passive copper splitter cable, IB twin port NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP, 3m	P-Rel
NDR	800GE	980-9I928-00N001	MCP7Y10-N001	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112, 1m, fin to flat	P-Rel
NDR	800GE	980-9I80P-00N003	MCP7Y10-N003	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112, 3m, fin to flat	P-Rel
NDR	800GE	980-9I80B-00N001	MCP7Y40-N001	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xQSFP112, 1m, fin to flat	P-Rel
NDR	800GE	980-9I75R-00N003	MCP7Y40-N003	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xQSFP112, 3m, fin to flat	P-Rel
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
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
NDR	400GE	980-9I693-00NS00	MMA1Z00-NS400	NVIDIA single port transceiver, 400Gbps, QSFP112, MPO12 APC, 850nm MMF, up to 50m, flat top	P-Rel
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
NDR	400GE	980-9I51S-00NS00	MMA4Z00-NS400	NVIDIA single port transceiver, 400Gbps, OSFP, MPO12 APC, 850nm MMF, up to 50m, flat top	MP
NDR	NA	980-9I068-00NM00	MMS1X00-NS400	NVIDIA single port transceiver, 400Gbps, NDR, QSFP112, MPO, 1310nm SMF, up to 500m, flat top	Prototype
NDR	NA	980-9I301-00NM00	MMS4X00-NM-FLT	NVIDIA twin port transceiver, 800Gbps, 2xNDR, OSFP, 2xMPO12 APC, 1310nm SMF, up to 500m, flat top	P-Rel
NDR	NA	980-9I30I-00NM00	MMS4X00-NS-FLT	NVIDIA twin port transceiver, 800Gbps, 2xNDR, OSFP, 2xMPO12 APC, 1310nm SMF, up to 100m, flat top	MP
NDR	400GE	980-9I31N-00NM00	MMS4X00-NS400	NVIDIA single port transceiver, 400Gbps, OSFP, MPO12 APC, 1310nm SMF, up to 100m, flat top	MP
NA	400GE	980-9I693-F4NS00	MMA1Z00-NS400-T	SINGLE PORT TRANSCEIVER, 400GBPS, 400GbE, QSFP112, MPO12 APC, 850NM MMF, UP TO 50M, FLAT TOP	P-Rel

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	Life Cycle Phase
NA	400GE	980-9I51S-F4NS00	MMA4Z00-NS400-T	SINGLE PORT TRANSCEIVER, 400GBPS, 400GbE, OSFP, MPO12 APC, 850NM MMF, UP TO 50M, FLAT TOP	P-Rel



Starting with firmware version xx.48.10xx, the twin-port transceiver MMA4Z00-NS-FLT may not function properly and can cause link-up failures.

For further information, see Known Issues 4835832.

8.1.4 HDR / 200GbE Cables

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	Life Cycle Phase
NA	200GE	980-9I54C-00V001	MCP1650-V001E30	Mellanox Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 1m, black pulltab, 30AWG	EOL [HVM]
NA	200GE	980-9I54D-00V002	MCP1650-V002E26	Mellanox Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 2m, black pulltab, 26AWG	EOL [HVM]
NA	200GE	980-9I54H-00V00A	MCP1650-V00AE30	Mellanox Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 0.5m, black pulltab, 30AWG	EOL [HVM]
NA	200GE	980-9I54I-00V01A	MCP1650-V01AE30	Mellanox Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 1.5m, black pulltab, 30AWG	EOL [HVM]
NA	200GE	980-9I54L-00V02A	MCP1650-V02AE26	Mellanox Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 2.5m, black pulltab, 26AWG	EOL [HVM]
HDR	200GE	980-9I39E-00H001	MCP7H50-H001R30	Nvidia Passive copper splitter cable, 200Gbps to 2x100Gbps, QSFP56 to 2xQSFP56, 1m	HVM
HDR	200GE	980-9I549-00H002	MCP1650-H002E26	Nvidia Passive Copper cable, up to 200Gbps, QSFP56 to QSFP56, 2m	HVM
HDR	200GE	980-9I54A-00H00A	MCP1650-H00AE30	Nvidia Passive Copper cable, up to 200Gbps, QSFP56 to QSFP56, 0.5m	HVM
HDR	200GE	980-9I46K-00H001	MCP7Y60-H001	NVIDIA passive copper splitter cable, 400(2x200)Gbps to 2x200Gbps, OSFP to 2xQSFP56, 1m, fin to flat	MP
HDR	200GE	980-9I46L-00H002	MCP7Y60-H002	NVIDIA passive copper splitter cable, 400(2x200)Gbps to 2x200Gbps, OSFP to 2xQSFP56, 2m, fin to flat	MP

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	LifeCycle Phase
HDR	400GE	980-9I93N-00H001	MCP7Y70-H001	NVIDIA passive copper splitter cable, 400(2x200)Gbps to 4x100Gbps, OSFP to 4xQSFP56, 1m, fin to flat	MP
HDR	400GE	980-9I93O-00H002	MCP7Y70-H002	NVIDIA passive copper splitter cable, 400(2x200)Gbps to 4x100Gbps, OSFP to 4xQSFP56, 2m, fin to flat	MP
HDR	NA	980-9I055-00H000	MMS1W50-HM	Mellanox transceiver, IB HDR, up to 200Gb/s, QSFP56, LC-LC, 1310nm, FR4	MP
HDR	200GE	980-9I548-00H001	MCP1650-H001E30	Nvidia Passive Copper cable, up to 200Gbps, QSFP56 to QSFP56, 1m	HVM
HDR	200GE	980-9I54B-00H01A	MCP1650-H01AE30	Nvidia Passive Copper cable, up to 200Gbps, QSFP56 to QSFP56, 1.5 m	HVM

8.1.5 EDR / 100GbE Cables

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	LifeCycle Phase
NA	100GE	980-9I62V-00C002	MCP1600-C002E30N	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 2m, Black, 30AWG, CA-N	HVM
NA	100GE	980-9I62O-00C003	MCP1600-C003E30L	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 3m, Black, 30AWG, CA-L	HVM
NA	100GE	980-9I62C-00C01A	MCP1600-C01AE30N	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 1.5m, Black, 30AWG, CA-N	HVM
NA	100GE	980-9I62I-00C02A	MCP1600-C02AE30L	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 2.5m, Black, 30AWG, CA-L	HVM
EDR	NA	980-9I62Q-00E001	MCP1600-E001E30	Mellanox Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 1m, Black, 30AWG	HVM
EDR	NA	980-9I62Z-00E005	MCP1600-E005E26	Mellanox Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 5m, Black, 26AWG	HVM
NA	100GE	980-9I62O-00C001	MCP1600-C001E30N	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 1m, Black, 30AWG, CA-N	HVM
NA	100GE	980-9I625-00C005	MCP1600-C005E26L	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 5m, Black, 26AWG, CA-L	HVM
NA	100GE	980-9I62Z-00C003	MCP1600-C003E26N	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 3m, Black, 26AWG, CA-N	EOL [HVM]

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	LifeCycle Phase
NA	100GE	980-9I627-00C00A	MCP1600-C00AE30N	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 0.5m, Black, 30AWG, CA-N	EOL [HVM]
EDR	100GE	980-9I62V-00C003	MCP1600-E003	Mellanox Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 3m 26AWG	EOL [HVM]
EDR	NA	980-9I62U-00E002	MCP1600-E002E30	Mellanox Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 2m, Black, 30AWG	HVM

8.1.6 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.

Data Rate	Cable OPN	Description
800GE	RTXM600-7xx	800G OSFP112-2x400G QSFP112 Active Optical Cable (rev: R1)
800GE	DME8821X-ECxx	800G OSFP112-2x400G QSFP112 Active Optical Cable (rev: 05)
800GE	C-OSG8CNSxxx-N00	INNOLIGHT 800G DR8 OSFP TO 2X400G QSFP112 DR4 BREAKOUT AOC
800GE	EOLO-138HG-5H-DR2	EOPTOLINK 800G OSFP MODULE
800GE	RTXM600-710	ACCELINK 800G OSFP TO 2X400G QSFP112 BREAKOUT AOC
800GE	T-RS8CNT-NMT	INNOLIGHT 800G DR8 OSFP RHS, DUAL MPO-12 APC
800GE	EOLO-138HG-5H-DR2	EOPTOLINK 800G 2XDR4 MODULE
800GE	AGP800C11311S50	ACCELIGHT 2X400G DR4 OSFP MODULE
800GE	FTCE4517E1PCA-2N	CREDO 800G OSFP TO 2X400G QSFP112 ACC
400GE	AQQLBCQ4EDLA1729	AOI 400G FR4 QSFP112 MODULE
400GE	ATRF-C020	HGTECH 200G QSFP56 AOC 20M
400GE	C-GD4CNS010-N00	INNOLIGHT 400G QSFP112 TO 400G QSFP-DD AOC
400GE	CTF4XFR4CS1-01	INNOLIGHT 400G-FR4 MODULE
400GE	EOLO-134HG-5H-B	EOPTOLINK 400G OSFP DR4 MODULE
400GE	RTXM600-610	ACCELINK 400G QSFP-DD TO QSFP112 AOC
400GE	T-GQ4CNT-N00	INNOLIGHT 400G QSFP112 FR4 MODULE, LC
400GE	T-RS4CNH-NFL	INNOLIGHT 400G (BRCM LASER)
400GE	T-RS4CNH-NFM	INNOLIGHT 400G (SUMI LASER)

Data Rate	Cable OPN	Description
400GE	T-OH4CNT-N00	INNOLIGHT 400G DR4+ QSFP112 MODULE
400GE	EOLQ-134HG-5H-DR2	EOPTOLINK 400G DR4 MODULE
200GE	EOLQ-132HG-5H-M3	EOPTOLINK 200G QSFP112 DR2 MODULE
200GE	QSFP-200-CU3M	CISCO 200G QSFP56 DAC 3M
200GE	RTXM500-301-F1	ACCELINK 200G QSFP56 SR4
200GE	RTXM600-338-R0	ACCELINK 200G QSFP112 VR2 MODULE
200GE	T-FX4FNS-N00	INNOLIGHT 200G QSFP56 SR4 MODULE
200GE	T-GP2CNH-NR0	INNOLIGHT 200G QSFP112 DR2 MODULE, MPO-12
200GE	TR-HM4M085V-CF21	CREALIGHT 200G QSFP112 VR2 MODULE
200GE	QSFP-200-CU3M	200G QSFP56 TO QSFP56 PASSIVE COPPER CABLE, 3M
200GE	EOLQ-132HG-5H-M3	EOPTOLINK 200G DR2 QSFP112
200GE	T-GP2CNH-NR0	INNOLIGHT 200G DR2 QSFP112

8.1.7 Tested Switches

8.1.8 XDR / 800GbE Switches

Speed	NVIDIA SKU	Legacy OPN	Description	LifeCycle Phase
XDR	920-9B3 4F-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	920-9B3 6F-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
800 GbE	920-9N4 2F-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

8.1.9 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

Speed	NVIDIA SKU	Legacy OPN	Description	LifeCycle Phase
400GbE	920-9N301-00xB-xxx	SN4700	NVIDIA Spectrum-3 based 400GbE, 1U Open Ethernet switch, 32xQSFP-DD ports, x86 CPU, standard depth	MP

8.1.10 HDR / 200GbE Switches

Speed	NVIDIA SKU	Legacy OPN	Description	LifeCycle Phase
200GbE	920-9N201-00F7-ON1	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

- [Changes and New Feature History](#)
- [Bug Fixes 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
40.47.1088	
Bug Fixes	See <i>Bug Fixes in this Firmware Version</i> 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.

Feature/Change	Description
40.47.1026	
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.
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 GVM 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 <i>Bug Fixes in this Firmware Version</i> 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 .
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.

Feature/Change	Description
40.46.1006	
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 <i>Bug Fixes in this Firmware Version</i> section.

Feature/Change	Description
40.45.1200	
Bug Fixes	See <i>Bug Fixes in this Firmware Version</i> section.

Feature/Change	Description
40.44.1036	
Static Split 8x100G ConnectX-8 to Spectrum-4 with SM Modules	A static split of 8x100G channels from a ConnectX-8 SuperNIC to a Spectrum-4 switch allows the system to use Single Mode (SM) optical modules for high-speed data transmission across a long-distance fiber link. This setup is typically used in high-performance networks where there is a need for high throughput (e.g., 800G in total bandwidth) with low latency, such as in data centers or high-performance computing environments.
DOCA Telemetry	DOCA Telemetry enables users to monitor and collect data related to the performance, health, and behavior of systems or applications running on DOCA. To optimize for a faster sampling period, it is recommended to configure all PCIe-related Diagnostic Data IDs sequentially, one after another to prevent a prolonged sampling period.
PCIe Switch fwreset	Added support for a new synchronized flow, including a tool and driver, to perform a fwreset on setups with a PCIe switch configuration.
PTP	Unified PTP is now supported across different VFs on the same PF.
Dual-Mode Temperature Compensated Crystal Oscillator (DC-TCXO) and Synchronous Ethernet (SyncE) Source	DC-TCXO is used now as the source of timing for SyncE, providing an accurate and stable clock for the synchronized operation of network devices that rely on Ethernet for timing.
DPA Application Signing	Allows DOCA applications signed with OEM/NVIDIA certificate private keys to be loaded onto the DPA engine, after the OEM/NVIDIA root certificates are installed on the NIC.
Data-Path Accelerator (DPA)	The DPA hardware version is now exposed as a new capability, labeled "dpa_platform_version."

Feature/Change	Description
40.44.1036	
Block SMP Traffic	Added a new NV config (SM_DISABLE, default 0) which, when enabled, blocks SMP traffic that does not originate from the SM.
Dynamic Long Cables	Added the ability to set cable length as a parameter in the PFCC access register. The cable length is used in the calculation of RX lossless buffer parameters, including size, Xoff, and Xon thresholds.
Bug Fixes	See <i>Bug Fixes in this Firmware Version</i> section.

Feature/Change	Description
40.44.0212	
Segment on PCIe Switch	Added support for Segment on PCIe switch.
AER on PCIe Switch Bridge	Added support for AER on PCIe switch bridge.
Bug Fixes	See <i>Bug Fixes in this Firmware Version</i> section.

Feature/Change	Description
40.44.0208	
General	This is the initial firmware release of NVIDIA® ConnectX®-8 SuperNIC. ConnectX-8 has the same feature set as ConnectX-7 adapter card. For the list of the ConnectX-7 firmware features, please see ConnectX-7 Firmware Release Notes. The features described here are new features in addition to the ConnectX-7 set.
Link Speed	NVIDIA® ConnectX®-8 SuperNIC supports 800Gb/s or XDR IB or 2 x 400GbE link speeds. Note: 800GbE link speed is not supported on a single port.
Planarized Topology Network	ConnectX®-8 SuperNIC uses planarized topology network to reach Extended Data Rate (XDR) performance.
Direct NIC-GPU Datapath	To read/write data directly from the GPU and to overcome grace CPU PCIe bandwidth issue a direct NIC-GPU datapath is required. To do so, the HCA exposes a side DMA engine as an additional PCIe function which is called "Data Direct". This additional DMA engine allows vHCA access data buffers using MKEY through it, providing multiple PCIe data path interfaces. Such behavior is needed in a scenario where different memory region requires different PCIe data path, i.e NUMA (Non Uniform Memory Access) systems. A vHCA is allowed to use a Data Direct function if <code>HCA_CAP.data_direct</code> is set. To use the Data Direct interface, the vHCA should create an MKEY with the <code>data_direct</code> bit set. The MKEY returned enables access through the side DMA engine. The MKEY access mode must be PA. It supports only the following fields: <code>a, rw, rr, lw, lr, relaxed_ordering_write, relaxed_ordering_read, mkey[7:0], length64, pd, start_addr, len</code>. All other fields are reserved.

Feature/Change	Description
40.44.0208	
Congestion Control	Congestion Control provides performance isolation when multiple applications running on the same cluster. Additionally, it prevents congestion spreading when there is a slow receiver, reduce latency in the cluster, improves fairness, prevents parking-lot effects and packet's drop in lossy networks.
Multiple Encapsulation/Decapsulation Operation on a Packet	This capability enables the encapsulation table to be opened on both the FDB and the NIC tables together.
Crypto Algorithms	Extended the role-based authentication to cover all crypto algorithms. Now the TLS , IPsec , MACsec , GCM , mem2mem , and NISP work when nv_crypto_conf.crypto_policy = CRYPTO_POLICY_FIPS_LEVEL_2 , meaning all cryptographic engines can also work in wrapped mode and not only in plaintext mode.
RoCE: Adaptive Timer	Enabled ADP timer to allow the user to configure RC or DC qp_timeout values lower than 16.
Multiple-Window in DPA Mode	Multi-window capability is now supported in DPA mode.
Doorbell Less QP	The new capability enables the user to send a queue without a doorbell record. To create a doorbell less QP/SP, set send_dbr_mode = 1 in qp/sq ctx as defined in the PRM.
Packet's Flow Label Fields	The flow_label fields can be set, added or copied from the packet.
ODP Event	The following prefetch fields are available ODP event: pre_demand_fault_pages, post_demand_fault_pages
Jump from NIC_TX to FDB_TX	The user can jump from NIC_TX to FDB_TX table and bypass the ACL table using the 'table_type_valid' and 'table_type' fields available in the steering action (STC) "Jump To Flow" table.

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
4824209 / 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: 40.47.1088
	Fixed in Release: 40.47.2682

Internal Ref.	Issue
4882128 / 4881141 / 4882127	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: 40.47.1088
	Fixed in Release: 40.47.2682
4892848 / 4864767	Description: Fixed an issue where PCIe messages could be dropped when a function undergoes FLR.
	Keywords: PCIe messages, FLR
	Detected in version: 40.47.1088
	Fixed in Release: 40.47.2682
4532684 / 4532684	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: 40.47.1088
	Fixed in Release: 40.47.2682
4794290 / 4621747	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: 40.47.1088
	Fixed in Release: 40.47.2682
4806969 / 4804664	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: 40.47.1088
	Fixed in Release: 40.47.2682
4824535 / 4809134	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: 40.47.1088
	Fixed in Release: 40.47.2682
4849832 / 4890248	Description: Fixed an issue where internal hardware error 0x07 occurred randomly during device unbind.
	Keywords: HW error 0x07

Internal Ref.	Issue
4861039 / 4705918	Detected in version: 40.47.1088
	Fixed in Release: 40.47.2682
	Description: Fixed an issue where PTP could converge to an incorrect time/offset and report an inaccurate path delay.
	Keywords: PTP
	Detected in version: 40.46.1006
	Fixed in Release: 40.47.2682

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
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

Internal Ref.	Issue
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
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

Internal Ref.	Issue
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
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 Sswitch/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.

Internal Ref.	Issue
	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. 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.
	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

Internal Ref.	Issue
	Fixed in Release: 40.46.1006
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

Internal Ref.	Issue
4436922	Description: Fixed DC InfiniBand functionality. Keywords: DC Detected in version: 40.45.1020 Fixed in Release: 40.45.1200

Internal Ref.	Issue
4321324	Description: Fixed an issue in SD where sending an ATS translation address request could lead to a completion timeout. Keywords: PCI

Internal Ref.	Issue
	Detected in version: 40.44.0208 Fixed in Release: 40.45.1020
4318537	Description: Fixed an issue where the AI and HAI parameters of the ZTR_RTTCC algorithm, when configured by users, were automatically overwritten upon link speed changes. With this fix, if AI/HAI values were tuned for link speeds other than 100Gb/s, users should now divide those values by (link_speed / 100) to maintain consistent congestion control algorithm behavior. Keywords: Congestion control, ZTR_RTTCC Detected in version: 40.44.0208 Fixed in Release: 40.45.1020
-	Description: Although ConnectX-8 SuperNIC is defined to work at Gen6 x16 (default) or Gen5 x32, in firmware v40.44.0204, the default configuration is Gen5 x32. Changing between the modes is done by an NVConfig command. Keywords: Gen6, Gen5, PCIe Detected in version: 40.44.0208 Fixed in Release: 40.45.1020
4369943	Description: Fixed an issue where the BMC would fail to read the RDE Port resource when operating in Ethernet mode. Keywords: BMC, RDE Detected in version: 40.44.1036 Fixed in Release: 40.45.1020
4389140	Description: Updated PLDM sensor ID values to align with the latest DSP2054 specifications. Keywords: PLDM Detected in version: 40.44.1036 Fixed in Release: 40.45.1020
4161849	Description: The Green LED remains solid and does not blink when running traffic. Keywords: Green LED Detected in version: 40.44.0212 Fixed in Release: 40.45.1020
4158184	Description: The Lane Error Status may occasionally appear in Configuration space. It can be safely ignored as it does not have any impact on device performance. Users are encouraged to monitor their systems, but this condition does not warrant any immediate action unless other issues arise. Keywords: Lane Error Status Detected in version: 40.44.0208 Fixed in Release: 40.45.1020
4274669	Description: Fixed a race condition that could prevent the application from transmitting when VQoS is enabled. Keywords: VQoS Detected in version: 40.44.1036

Internal Ref.	Issue
	Fixed in Release: 40.45.1020
4199274	Description: Fixed an issue where RTT packets with any destination MAC address were incorrectly treated as having a valid destination MAC. The new firmware now discards RTT packets if their destination MAC does not match the port's MAC.
	Keywords: RTT, destination MAC
	Detected in version: 40.44.1036
	Fixed in Release: 40.45.1020
4260394	Description: Aligned the command channel entry offset in the GVMI context to a cache line boundary to avoid issues with GET/SET_GVMI_FW_CONTEXT_SUB_STRUCT operations, which require read and write actions to be contained within a single cache line.
	Keywords: Firmware assert
	Detected in version: 40.44.1036
	Fixed in Release: 40.45.1020
4268027	Description: Increased RX lossless buffer xoff and size to extend the delay before triggering Pause frames.
	Keywords: RX lossless buffer
	Detected in version: 40.44.1036
	Fixed in Release: 40.45.1020

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.47.3576	• License • 3rd Party Notice • 3rd Party Unify Notice • HCA Firmware EULA
DOCA-Host	3.2.3	• License • 3rd Party Notice • 3rd Party Unify Notice
MFT FreeBSD	4.34.1-25	• 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.



Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation and/or its affiliates in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2026 NVIDIA Corporation & affiliates. All Rights Reserved.

