



# **NVIDIA ConnectX-8 SuperNIC Firmware Release Notes v40.47.1088 (2025 LTS U1)**

# Table of contents

Firmware Compatible Products	3
Changes and New Features	5
Customer Affecting Changes	5
Declared Unsupported Features	8
Bug Fixes in this Firmware Version	9
Known Issues	10
PreBoot Drivers (FlexBoot/UEFI)	11
Validated and Supported Cables and Switches	12
Release Notes History	41
Changes and New Feature History	41
Bug Fixes History	49
Legal Notices and 3rd Party Licenses	51

## Info

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.

## Release Notes Update History

Version	Date	Description
40.47.1088	December 2025	Initial release of this Release Notes version.

## 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.

## Firmware Download

Please visit [Firmware Downloads](#).

---

# Firmware Compatible Products

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

- InfiniBand - EDR, HDR100<sup>2</sup>, HDR<sup>2</sup>, NDR200<sup>2</sup>, NDR<sup>2</sup>, XDR<sup>2</sup>
- Ethernet - 25GbE, 50GbE<sup>1</sup>, 100GbE<sup>1</sup>, 200GbE<sup>2</sup>, 400GbE<sup>2</sup>
- PCI Express Gen6., supporting backwards compatibility for v5.0, v4.0, v3.0, v2.0 and v1.1

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

<sup>2</sup>. Speed that supports PAM4 mode only.

## Note

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

## Note

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

## Supported Devices

Refer to the hardware [documentation](#) for the list of supported devices.

## 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:

	<b>Supported Version</b>
ConnectX-8 Firmware	40.47.1088 / 40.47.1026 / 40.46.3048
DOCA-HOST	3.2.1 / 3.2.0 <b>Note:</b> 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 <b>Note:</b> For the list of the supported Operating Systems, please refer to the driver's Release Notes.
MFT	4.34.1-10 / 4.34.0-145 / 4.33.0-169 <b>Note:</b> For the list of the supported Operating Systems, please refer to the driver's Release Notes.
FlexBoot	3.8.201
UEFI	14.40.10
NVOS	25.02.6000 onwards
Quantum-3 FW (part of NVOS)	35.2016.2080 onwards

---

# Changes and New Features

## Note

**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.

## Info

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
	40.47.1088
<b>Bug Fixes</b>	See <i>Bug Fixes in this Firmware Version</i> section.

## Customer Affecting Changes

### 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.

<b>Introduced in Version</b>	<b>Description</b>
N/A	N/A

## 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.

<b>Planned for Version</b>	<b>Description</b>
N/A	N/A

## 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.

<b>Introduced in Version</b>	<b>Description</b>
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.

Introduced in Version	Description
40.44.1036	<p>In ConnectX-8, the DIAG_COUNTER interface has been changed from the following set of commands:</p> <ul style="list-style-type: none"> <li>• SET_DIAGNOSTIC_PARAMS</li> <li>• QUERY_DIAGNOSTIC_PARAMS</li> <li>• QUERY_DIAGNOSTIC_COUNTERS</li> <li>• ICMD_SET_DIAGNOSTIC_PARAMS</li> <li>• ICMD_QUERY_DIAGNOSTIC_PARAMS</li> <li>• ICMD_QUERY_DIAGNOSTIC_COUNTERS</li> </ul> <p>to:</p> <ul style="list-style-type: none"> <li>• DIAG_DATA_OWNERSHIP</li> <li>• DIAG_DATA_PARAMS_CONTEXT</li> <li>• DIAG_DATA_ID_LIST</li> <li>• DIAG_DATA_QUERY</li> </ul> <p>The old interface will now return zero values when queried.</p>

## 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

**Note**

For inquiries regarding mitigation, please contact [NVIDIA Support](#).

## 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

---

# Bug Fixes in this Firmware Version

Internal Ref.	Issue
4608544	<b>Description:</b> Fixed an issue where, in rare live migration scenarios, a delayed doorbell triggered a false timeout alarm.
	<b>Keywords:</b> Live migration, doorbell, timeout alarm
	<b>Detected in version:</b> 40.46.1006
	<b>Fixed in Release:</b> 40.47.1088
4648642	<b>Description:</b> Fixed a rare issue in which destroying PCC NP configuration objects could result in assert 0x8175 being logged in dmesg.
	<b>Keywords:</b> Assert 0x8175, PCC NP
	<b>Detected in version:</b> 40.47.1026
	<b>Fixed in Release:</b> 40.47.1088
4655971	<b>Description:</b> Fixed the PCIe counters to correctly report event values in nanoseconds.
	<b>Keywords:</b> DOCA Telemetry Diagnostics
	<b>Detected in version:</b> 40.47.1026
	<b>Fixed in Release:</b> 40.47.1088
4690503	<b>Description:</b> 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.
	<b>Keywords:</b> DPA process, BSS
	<b>Detected in version:</b> 40.47.1026
	<b>Fixed in Release:</b> 40.47.1088

---

# Known Issues

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

### **VF+SF Network Function Limitations in Switchdev Mode**

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.

---

# PreBoot Drivers (FlexBoot/UEFI)

## FlexBoot Changes and New Features

For further information, please refer to the [FlexBoot Release Notes](#).

## UEFI Changes and Major New Features

For further information, please refer to the [UEFI Release Notes](#).

---

# Validated and Supported Cables and Switches

## Validated and Supported Cables and Modules

### 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

 **Note**

For information on any cable limitations, please refer to [Known Issues](#) section.

**Note**

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.

## 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	980-9IATO-00XM00	NVIDIA single port transceiver for ConnectX-8 Mezz Card, 800Gbps, OSFP, MPO, 1310nm SMF, EML, up to 500m, RHS	EVT

## NDR / 400GbE / 800GbE Cables

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
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

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
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
NDR	NA	980-9I57Y-00N050	MFP7E10-N050	NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 50m	MP
NDR	NA	980-9I571-00N003	MFP7E20-N003	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 3m	MP

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
NDR	NA	980-9I572-00N005	MFP7E20-N005	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 5m	MP
NDR	NA	980-9I573-00N007	MFP7E20-N007	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 7m	MP
NDR	NA	980-9I554-00N010	MFP7E20-N010	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 10m	MP
NDR	NA	980-9I555-00N015	MFP7E20-N015	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 15m	MP
NDR	NA	980-9I556-00N020	MFP7E20-N020	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 20m	MP
NDR	NA	980-9I557-00N030	MFP7E20-N030	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 30m	MP
NDR	NA	980-9I55Z-00N050	MFP7E20-N050	NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 50m	MP

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
NDR	NA	980-9I55A-00N003	MFP7E30-N003	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 3m	MP
NDR	NA	980-9I55B-00N005	MFP7E30-N005	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 5m	MP
NDR	NA	980-9I58C-00N007	MFP7E30-N007	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 7m	MP
NDR	NA	980-9I58D-00N010	MFP7E30-N010	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 10m	MP
NDR	NA	980-9I58E-00N015	MFP7E30-N015	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 15m	MP
NDR	NA	980-9I58F-00N020	MFP7E30-N020	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 20m	MP
NDR	NA	980-9I58G-00N030	MFP7E30-N030	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 30m	MP
NDR	NA	980-9I58O-00N030	MFP7E30-N040	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 40m	MP
NDR	NA	980-9I58H-00N050	MFP7E30-N050	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 50m	MP

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
NDR	NA	980-9I581-00N050	MFP7E30-N060	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 60m	MP
NDR	NA	980-9I582-00N050	MFP7E30-N070	NVIDIA passive fiber cable, SMF, MPO12 APC to MPO12 APC, 70m	MP
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

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
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

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
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

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
NDR	800GE	980-91929-00N002	MCP7Y10-N002	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112,2m, fin to flat	P-Rel
NDR	800GE	980-9180A-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	980-9180Q-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	980-9180C-00N002	MCP7Y40-N002	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xQSFP112, 2m, fin to flat	P-Rel

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
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

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
NDR	800GE	980-9175F-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	980-9146H-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	980-9146J-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	980-91600-00N003	MCA4J80-N003-FLT	Active copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 3m, flat top	MP
NDR	800GE	980-91948-00N004	MCA7J60-N004	NVIDIA active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 4m, fin to flat	P-Rel

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
NDR	800GE	980-91949-00N005	MCA7J60-N005	NVIDIA active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 5m, fin to flat	P-Rel
NDR	800GE	980-9181B-00N004	MCA7J65-N004	NVIDIA Active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112, 4m, fin to flat	P-Rel
NDR	800GE	980-9181C-00N005	MCA7J65-N005	NVIDIA Active copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xQSFP112, 5m, fin to flat	P-Rel
NDR	NA	980-91A0G-00N001	MCP4Y10-N001-FLT	NVIDIA Passive Copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 1m, flat top	MP
NDR	NA	980-91A0H-00N001	MCP4Y10-N001-FTF	NVIDIA Passive Copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 1m, flat to finned	MP

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
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

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
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

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
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

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	LifeCycle Phase
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
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

## HDR / 200GbE Cables

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	LifeCycle 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]

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
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

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
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
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

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
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

## **EDR / 100GbE Cables**

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
NA	100GE	980-9I62V-00C002	MCP1600-C002E30N	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 2m, Black, 30AWG, CA-N	HVM
NA	100GE	980-9I620-00C003	MCP1600-C003E30L	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 3m, Black, 30AWG, CA-L	HVM

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
NA	100GE	980-9162C-00C01A	MCP1600-C01AE30N	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 1.5m, Black, 30AWG, CA-N	HVM
NA	100GE	980-9162I-00C02A	MCP1600-C02AE30L	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 2.5m, Black, 30AWG, CA-L	HVM
EDR	NA	980-9162Q-00E001	MCP1600-E001E30	Mellanox Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 1m, Black, 30AWG	HVM
EDR	NA	980-9162Z-00E005	MCP1600-E005E26	Mellanox Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 5m, Black, 26AWG	HVM
NA	100GE	980-9162O-00C001	MCP1600-C001E30N	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 1m, Black, 30AWG, CA-N	HVM

<b>IB Data Rate</b>	<b>Eth Data Rate</b>	<b>NVIDIA SKU</b>	<b>Legacy P/N</b>	<b>Marketing Description</b>	<b>LifeCycle Phase</b>
NA	100GE	980-91625-00C005	MCP1600-C005E26L	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 5m, Black, 26AWG, CA-L	HVM
NA	100GE	980-9162Z-00C003	MCP1600-C003E26N	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 3m, Black, 26AWG, CA-N	EOL [HVM]
NA	100GE	980-91627-00C00A	MCP1600-C00AE30N	Mellanox Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 0.5m, Black, 30AWG, CA-N	EOL [HVM]
EDR	100GE	980-9162V-00C003	MCP1600-E003	Mellanox Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 3m 26AWG	EOL [HVM]
EDR	NA	980-9162U-00E002	MCP1600-E002E30	Mellanox Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 2m, Black, 30AWG	HVM

## Supported 3rd Party Cables and Modules

### Note

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	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

<b>Data Rate</b>	<b>Cable OPN</b>	<b>Description</b>
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

<b>Data Rate</b>	<b>Cable OPN</b>	<b>Description</b>
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)
400GE	T-OH4CNT-N00	INNOLIGHT 400G DR4+ QSFP112 MODULE
400GE	EOLO-134HG-5H-DR2	EOPTOLINK 400G DR4 MODULE
200GE	EOLQ-132HG-5H-M3	EOPTOLINK 200G QSFP112 DR2 MODULE

<b>Data Rate</b>	<b>Cable OPN</b>	<b>Description</b>
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

<b>Data Rate</b>	<b>Cable OPN</b>	<b>Description</b>
200GE	T-GP2CNH-NR0	INNOLIGHT 200G DR2 QSFP112

## Tested Switches

### XDR / 800GbE Switches

<b>Speed</b>	<b>NVIDIA SKU</b>	<b>Legacy OPN</b>	<b>Description</b>	<b>LifeCycle Phase</b>
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

Speed	NVIDIA SKU	Legacy OPN	Description	LifeCycle Phase
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

## 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

## 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

---

# Release Notes History

- [Changes and New Feature History](#)
- [Bug Fixes History](#)

## Changes and New Feature History

### Note

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
40.47.1026	
<b>Lane Margin</b>	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.
<b>PCIe Trace Function</b>	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.
<b>Passing Metadata Registers between the NIC Layer and the E-Switch (esw) Layer</b>	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.

Feature/Change	Description
<b>DPA (Data Path Accelerator) Partition Creation</b>	Access control was added to ensure that only the VHCA instance that created a DPA partition is permitted to modify or delete it.
<b>DPA Manifest</b>	A new DPA Manifest mechanism was introduced to define and manage application permissions.
<b>DPA TIMER</b>	DPA TIMER functionality has been exposed through the MTCTR access register, allowing direct access by applications.
<b>Parallel Suspend of VFs</b>	Added support for parallel suspend operations across multiple VFs.
<b>RTT RTC Timestamp</b>	<p>Added support for using the real-time clock to fill the request and response timestamps in hardware-generated RTT packets.</p> <p>To enable this feature, set <code>REAL_TIME_CLOCK_ENABLE</code> in <b>mlxconfig</b> and configure <code>ROCE_CC_RTT_TIMESTAMP_FORMAT</code> to <code>0x02 (REAL_TIME)</code>.</p> <p>For additional details, see Known Issue 4496642 in the Known Issues section.</p>
<b>PCC NP IFA2 GNS</b>	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 <code>pcc_config</code> and <code>pcc_np_config</code> must be identical across all slots using IFA2.
<b>Enhanced Error Recovery for GGA QPs</b>	<p>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.</p> <p>Unlike RDMA QPs, the error CQE may redundantly reference a valid mkey, therefore, the software should reconstruct all mkeys that received error CQE notifications.</p>
<b>Enable/Disable ECN in Upstream</b>	Added the ability to enable or disable ECN in the upstream by allowing the <code>MODIFY_CONG_STATUS</code> and <code>QUERY_CONG_STATUS</code> commands in <code>mlx5_fwctl</code> .

Feature/Change	Description
<b>Link Speed per-lane</b>	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.
<b>API to Write PSP Master Key</b>	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.
<b>ADP-RETX Timeout Profile</b>	Firmware now allows the ADP-RETX timeout profile to be configured even when there are open QPs.
<b>PCI Logs</b>	PCI logs are now reported via the existing NC-SI OEM command <code>Get Log Info</code> (Command = 0x0, Parameter = 0x2F).
<b>Adaptive Hot-plug System (AHS)</b>	Added support for Adaptive Hotplug System (AHS) alongside the existing NHP solution, enhancing hotplug flexibility and system adaptability.
<b>Additional STE Action</b>	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.
<b>NV Configuration</b>	Added an NV configuration option to allow disabling XDR. <b>Note:</b> Disabling SDR or enabling configurations not supported by the INI file remains unsupported.
<b>MVCAP (Multi-Version Capability)</b>	Added support for MVCAP (Multi-Version Capability) functionality enabling improved compatibility and version management across multiple components.
<b>Bug Fixes</b>	See <i>Bug Fixes in this Firmware Version</i> section.

Feature/Change	Description
	40.46.3048

Feature/Change	Description
<b>Security Hardening Enhancements</b>	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	
<b>PCIe TLP Processing Hints (TPH) and Steering Tag (ST)</b>	Enabled PCIe TLP Processing Hints (TPH) and Steering Tag (ST) during MKey creation. <b>Note:</b> The steering tag index in the MKey creation must reference an MSIX entry containing the actual steering tag value.
<b>PCIe Congestion Events</b>	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 <code>PCIE_CONGESTION_MONITOR</code> .
<b>RDMA QP</b>	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.
<b>PPCNT Counters</b>	Firmware now supports new counters in the PPCNT register to track multicast and unicast packets transmitted and received. The counters include: <ul style="list-style-type: none"> <li>• <code>port_multicast_xmit_pkts_high</code></li> <li>• <code>port_multicast_xmit_pkts_low</code></li> <li>• <code>port_multicast_rcv_pkts_high</code></li> <li>• <code>port_multicast_rcv_pkts_low</code></li> <li>• <code>port_unicast_xmit_pkts_high</code></li> <li>• <code>port_unicast_xmit_pkts_low</code></li> <li>• <code>port_unicast_rcv_pkts_high</code></li> <li>• <code>port_unicast_rcv_pkts_low</code></li> </ul>

Feature/Change	Description
<p><b>Safely Identify DPU/SmartNICs in a Machine and PCIe Slot</b></p>	<p>A new access register is introduced that accepts a type, length, and R/W command.</p> <ul style="list-style-type: none"> <li>• <b>Write operation:</b> 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</li> <li>• <b>Read operation:</b> If a buffer exists, its data is copied out. If not, the access register returns a size of 0 or an explicit error</li> </ul> <p>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 <math>(base + (type \ll 8))</math>. 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.</p> <p>This feature is designed for limited use cases and does not address multi-host scenarios or broader ICMC utilization implications.</p>
<p><b>Latency Histogram Counter</b></p>	<p>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.</p>
<p><b>Incoming NC-SI Messages Validation for the payload_len Field</b></p>	<p>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.</p>
<p><b>RSS with Crypto Offload</b></p>	<p>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).</p>
<p><b>SPDM</b></p>	<p>Updated SPDM measurements report to version 1.1.</p>
<p><b>Bug Fixes</b></p>	<p>See <i>Bug Fixes in this Firmware Version</i> section.</p>

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

Feature/Change	Description
	40.44.1036
<b>Static Split 8x100G ConnectX-8 to Spectrum-4 with SM Modules</b>	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.
<b>DOCA Telemetry</b>	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.
<b>PCIe Switch fwreset</b>	Added support for a new synchronized flow, including a tool and driver, to perform a fwreset on setups with a PCIe switch configuration.
<b>PTP</b>	Unified PTP is now supported across different VFs on the same PF.
<b>Dual-Mode Temperature Compensated Crystal Oscillator (DC-TCXO) and Synchronous Ethernet (SyncE) Source</b>	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.
<b>DPA Application Signing</b>	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.
<b>Data-Path Accelerator (DPA)</b>	The DPA hardware version is now exposed as a new capability, labeled "dpa_platform_version."

Feature/Change	Description
<b>Block SMP Traffic</b>	Added a new NV config (SM_DISABLE, default 0) which, when enabled, blocks SMP traffic that does not originate from the SM.
<b>Dynamic Long Cables</b>	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.
<b>Bug Fixes</b>	See <i>Bug Fixes in this Firmware Version</i> section.

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

Feature/Change	Description
40.44.0208	
<b>General</b>	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 <a href="#">ConnectX-7 Firmware Release Notes</a> . The features described here are new features in addition to the ConnectX-7 set.
<b>Link Speed</b>	NVIDIA® ConnectX®-8 SuperNIC supports 800Gb/s or XDR IB or 2 x 400GbE link speeds. <b>Note:</b> 800GbE link speed is not supported on a single port.
<b>Planarized Topology Network</b>	ConnectX®-8 SuperNIC uses planarized topology network to reach Extended Data Rate (XDR) performance.

Feature/Change	Description
<p><b>Direct NIC-GPU Datapath</b></p>	<p>To read/write data directly from the GPU and to overcome grace CPU PCIe bandwidth issue a direct NIC-GPU datapath is required.</p> <p>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.</p> <p>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.</p> <p>It supports only the following fields:</p> <pre>a, rw, rr, lw, lr, relaxed_ordering_write, relaxed_ordering_read, mkey[7:0], length64, pd, start_addr, len</pre> <p>. All other fields are reserved.</p>
<p><b>Congestion Control</b></p>	<p>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.</p>
<p><b>Multiple Encapsulation/Decapsulation Operation on a Packet</b></p>	<p>This capability enables the encapsulation table to be opened on both the FDB and the NIC tables together.</p>
<p><b>Crypto Algorithms</b></p>	<p>Extended the role-based authentication to cover all crypto algorithms. Now the <code>TLS. IPsec. MACsec. GCM, mem2mem</code>, and <code>NISP</code> work when</p> <pre>nv_crypto_conf.crypto_policy = CRYPTO_POLICY_FIPS_LEVEL_2</pre> <p>, meaning all cryptographic engines can also work in wrapped mode and not only in plaintext mode.</p>

Feature/Change	Description
<b>RoCE: Adaptive Timer</b>	Enabled ADP timer to allow the user to configure RC or DC qp_timeout values lower than 16.
<b>Multiple-Window in DPA Mode</b>	Multi-window capability is now supported in DPA mode.
<b>Doorbell Less QP</b>	The new capability enables the user to send a queue without a doorbell record. To create a doorbell less QP/SP, set <code>send_dbr_mode = 1</code> in qp/sq ctx as defined in the PRM.
<b>Packet's Flow Label Fields</b>	The <code>flow_label</code> fields can be set, added or copied from the packet.
<b>ODP Event</b>	The following prefetch fields are available ODP event: <code>pre_demand_fault_pages</code> , <code>post_demand_fault_pages</code>
<b>Jump from NIC_TX to FDB_TX</b>	The user can jump from <code>NIC_TX</code> to <code>FDB_TX</code> table and bypass the ACL table using the <code>'table_type_valid'</code> and <code>'table_type'</code> fields available in the steering action (STC) "Jump To Flow" table.

## Bug Fixes History

### Note

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 .



---

# 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.47.1088	<ul style="list-style-type: none"><li>• <a href="#">License</a></li><li>• <a href="#">3rd Party Notice</a></li><li>• <a href="#">3rd Party Unify Notice</a></li><li>• <a href="#">HCA Firmware EULA</a></li></ul>
DOCA-Host	3.2.1	<ul style="list-style-type: none"><li>• <a href="#">License</a></li><li>• <a href="#">3rd Party Notice</a></li><li>• <a href="#">3rd Party Unify Notice</a></li></ul>

Product	Version	Legal Notices and 3rd Party Licenses
MFT FreeBSD	4.34.1-10	<ul style="list-style-type: none"> <li>• <a href="#">License</a></li> <li>• <a href="#">3rd Party Notice</a></li> <li>• <a href="#">3rd Party Unify Notice</a></li> </ul>
MFT Linux		<ul style="list-style-type: none"> <li>• <a href="#">License</a></li> <li>• <a href="#">3rd Party Notice</a></li> <li>• <a href="#">3rd Party Unify Notice</a></li> </ul>
MFT VMware		<ul style="list-style-type: none"> <li>• <a href="#">License</a></li> <li>• <a href="#">3rd Party Notice</a></li> <li>• <a href="#">3rd Party Unify Notice</a></li> </ul>
MFT Windows		<ul style="list-style-type: none"> <li>• <a href="#">License</a></li> <li>• <a href="#">3rd Party Notice</a></li> <li>• <a href="#">3rd Party Unify Notice</a></li> </ul>

## 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. NVIDIA Corporation (“NVIDIA”) makes no 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 and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© Copyright 2025, NVIDIA. PDF Generated on 01/13/2026