NVIDIA ConnectX-9 SuperNIC Firmware Release Notes v82.50.1002

Table of Contents

1 Release Notes Update History

Version	Date	Description
82.50.1002	August 2026	Initial release of this Release Notes version.

2 Overview

The NVIDIA® ConnectX®-9 SuperNIC™ unleashes next-generation gigascale AI with up to 800 Gb/s per port over InfiniBand and Ethernet, delivering ultra-fast, efficient connectivity that supercharges AI data centers and cloud platforms. Fully integrated with NVIDIA Spectrum-X™ Ethernet and Quantum-X800 networking, it drives up to 1.6 Tb/s of throughput to Rubin GPUs, while advanced programmable IO and intelligent congestion control maximize performance and efficiency.

2.1 Firmware Download

Please visit Firmware Downloads.

3 Firmware Compatible Products

These are the release notes for the NVIDIA® ConnectX®-9 SuperNIC firmware.

ConnectX-9 firmware supports the following speeds/protocols:0

- InfiniBand - XDR, NDR
- Ethernet - 800GAUI8 C2M, 800GAUI4 C2M, 400GAUI2 C2M, 400GAUI4 C2M, 200GAUI1 C2M, 200GAUI2 C2M, 100GAUI1 C2M, 100GAUI1_SFE C2M, 800GAUI8 CR8, 800GAUI4 CR4, 400GAUI2 CR2, 400GAUI4 CR4, 200GAUI1 CR1, 200GAUI2 CR2, 100GAUI1 CR1, 100GAUI1_SFE CR1
- PCI Express Gen6., supporting backwards compatibility for v5.0, v4.0, v3.0, v2.0 and v1.1

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

For additional supported speeds/protocols by the ConnectX-9 ASIC, please refer to the hardware User Manual.

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-9 Firmware	82.50.1002 / 82.49.1014 / 82.48.1000
DOCA-HOST	3.5.0 / 3.4.0 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.
WinOF-2	26.7.50000 / 26.4.50010 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.
MFT	4.37.0-154 / 4.36.0-147 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.
FlexBoot	3.9.101
UEFI	14.42.11

	Supported Version
NVOS	25.02.6000 onwards
Quantum-3 FW (part of NVOS)	35.2016.2080 onwards

4 Changes and New Features

For the list of changes and bug fixes in earlier versions, see Release Notes History section.

Feature/Change	Description
82.50.1002	

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

Precoded Data Reception Configuration	Added the ability to override host Rx and Tx precoding using mlxconfig. This can be configured via the `HOST_MNG_RX_PRECODING_P*` and `HOST_MNG_TX_PRECODING_P*` parameters. The supported values for host Tx/Rx precoding behavior are: 1. `- 0 (0x0): DEVICE_DEFAULT` 2. `- 1 (0x1): PRECODING_ENABLE` 3. `- 2 (0x2): PRECODING_DISABLE` Example usage: `mlxconfig -d set HOST_MNG_RX_PRECODING_P*=1 mlxconfig -d set HOST_MNG_TX_PRECODING_P*=1` ⚠ When connecting via a passive copper cable to a link partner configured in force mode, the local NIC must also be configured in force mode. Additionally, it is recommended to keep precoding set to "ON".
Module Precoding Control	Module precoding control allows precoding to be enabled or disabled through PHY TLVs on both the host and the module, aligning NIC and module behavior with host PRM precoding control.
Modify-header Action Dependency	Added the missing source dependency check (SRC) for modify-header parallel actions, preventing odd actions from reading a field modified by the preceding even action.
FEC Alignment	FEC alignment is supported for lane speeds up to 100G per lane.
Bug Fixes	See Bug Fixes section.

For the list of changes and bug fixes in earlier versions, see Release Notes History 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
82.48.1000	**Transition to 2023 Microsoft UEFI Certificate Authority**: 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."

Introduced in Version	Description

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

4.1.4 Discontinued Features

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

- 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-9 SuperNIC in this firmware version:

- Zero Touch Tuning (ZTT)
- Hot reset
- Segment on PCIe switch
- LAG Bonding with Q-Affinity
- ISSU
- The following protocols:
  - 25 / 50G NRZ
  - EDR and HDR

5 Bug Fixes in this Firmware Version

Internal Ref.	Issue
5232646	**Description:** Resolved an issue where PLDM update from the auxiliary card side could fail on devices in multi-host configurations.
	Keywords: PLDM
	Detected in version: 82.49.1014
	Fixed in Release: 82.50.1002
4959871 / 4876707	**Description:** Increased TLV list sizes: • Fast list: 2K → 3K • Regular list: 24K → 36K • Shadow list: 16K → 40K The new sizes take effect only during new `NV_DATA` initialization scenarios, including `NV_DATA` erase, `CONF_4` TLV write, or `NV_DATA` garbage collection.
	Keywords: TLV list
	Detected in version: 82.48.1000
	Fixed in Release: 82.50.1002
5045949 / NVbug 6206019	**Description:** Resolved an issue where NSM set/get device mode settings v2 in PCIe device mode used only 3 bytes for user parameters. PCIe device mode now requires an additional 5 bytes of zero-filled padding.
	Keywords: PCIe
	Detected in version: 82.48.1000
	Fixed in Release: 82.50.1002
5041711	**Description:** Resolved an issue where, after toggling a port, the peer could incorrectly remain in the PHY-up state.
	Keywords: Port toggle
	Detected in version: 82.49.1014
	Fixed in Release: 82.50.1002
5012336 / NVbug 6140150	**Description:** Resolved a potential issue that could result in lower-than-expected bidirectional RDMA bandwidth in PCIe switch topologies.
	Keywords: RDMA bandwidth
	Detected in version: 82.49.1014
	Fixed in Release: 82.50.1002

Internal Ref.	Issue
4851684	**Description:** Resolved an issue where RDMA workloads using null memory regions with MKEY_BY_NAME enabled could experience up to a 30% throughput drop. The internal null/dump-fill key was incorrectly routed through the slower signature datapath. Workloads not using MKEY_BY_NAME or null memory regions were unaffected.
	Keywords: RDMA
	Detected in version: 82.49.1014
	Fixed in Release: 82.50.1002
4851684	**Description:** Resolved an issue where, under high-stress LRO traffic, an internal FIFO could enter a hardware deadlock state, preventing the NIC from receiving additional traffic. The fix detects and releases the deadlock if it occurs, restoring the NIC to a functional state.
	Keywords: LRO traffic
	Detected in version: 82.49.1014
	Fixed in Release: 82.50.1002
5174504	**Description:** Resolved an issue where intensive parallel QP INIT2RTR / DESTROY operations could incur high operation latency and reduce connection establishment rates when PCC was enabled.
	Keywords: PCC, INIT2RTR, DOCA PCC
	Detected in version: 82.49.1014
	Fixed in Release: 82.50.1002
5166648	**Description:** Resolved an issue involving an uninitialized stack read and an inverted EtherType check in the RX parser.
	Keywords: RX parser
	Detected in version: 82.48.1000
	Fixed in Release: 82.50.1002
4937640	**Description:** Resolved an issue where a PPCC mlxreg read operation could succeed when executed on a PF that is not the port owner, which could have security implications.
	Keywords: PCC, PPCC
	Detected in version: 82.49.1014
	Fixed in Release: 82.50.1002

Internal Ref.	Issue
5154717 / 5152076 / 5154806 / NVbug 6422398	**Description:** Resolved an issue where, when a transceiver module held its I2C management bus line stuck (SDA or SCL), the firmware detected the stuck bus and reset the module. After the reset, the module remained in low-power mode and its datapath was not re-established, preventing the link from coming up. On a shared module, all associated ports could drop simultaneously.
	Keywords: Cables
	Detected in version: 82.48.1000
	Fixed in Release: 82.50.1002
5155445	**Description:** Resolved an issue with SyncE frequency adjustment overflow when using a third-party secondary vendor.
	Keywords: SyncE frequency
	Detected in version: 82.49.1014
	Fixed in Release: 82.50.1002
5093695	**Description:** Resolved an issue where using mlxconfig to configure all modules with the 0xff module split value could reset the modules to the default configuration.
	Keywords: mlxconfig
	Detected in version: 82.48.1000
	Fixed in Release: 82.50.1002
5090639	**Description:** Resolved an issue in single-domain mode where an FLR on a GVMI did not clear the per-domain usage bits in `SP_core_diag_multi_domain->pcie_link_latency_max_read` and `SP_core_diag_multi_domain->pcie_link_latency_min_read`. As a result, when the next GVMI took ownership and attempted to configure these counters, it could receive `DIAG_DATA_LIST_WRITE_STATUS_NO_RESOURCES`.
	Keywords: PCIe Link Latency Arch counters
	Detected in version: 82.49.1014
	Fixed in Release: 82.50.1002
5089348	**Description:** Resolved an issue where the MFT database was not updated to configure profile 12 for the Flex Parser.
	Keywords: MFT, Flex Parser
	Detected in version: 82.48.1000
	Fixed in Release: 82.50.1002

Internal Ref.	Issue
5107808	**Description:** Resolved an issue in event coalescing where, when two events were coalesced, the selected event information depended on the cr value: if cr was 0, the earlier event information was used; if cr was 1, the later event information was used. The main affected fields were sn and timestamp.
	Keywords: TLP traffic
	Detected in version: 82.49.1014
	Fixed in Release: 82.50.1002
5077326 / 4902171	**Description:** Resulved an issue where a slow module response to the NCSI GetModuleSerialData command could incorrectly trigger a dead IRISC watchdog assert. The firmware now handles this waiting state correctly, so slow module responses no longer cause a false watchdog hang indication.
	Keywords: NCSI
	Detected in version: 82.48.1000
	Fixed in Release: 82.50.1002
4947006 / 4980956 / 4982511	**Description:** Resolved two rare race conditions in the PCC steering extension when the feature is enabled through path migration, empty RTT, or a custom header. Under rare timing conditions during QP creation or destruction, these issues could cause traffic disruption or resource leaks.
	Keywords: PCC steering extension
	Detected in version: 82.49.1014
	Fixed in Release: 82.50.1002

6 Known Issues

ConnectX-9 has the same feature set and limitations as the ConnectX-8 SuperNIC. For the list of ConnectX-8 known issues, refer to: https://docs.nvidia.com/networking/software/adapter-firmware/index.html#connectx-8.

The limitations listed below apply to ConnectX-9 only.

Internal Ref.	Issue
5234743 / 5145460	**Description:** Query Port Telemetry v2 (CMD 0x14) reports tag 0x20 (rx_los_mask), which should not be populated.
	Workaround: N/A
	Keywords: Query Port Telemetry v2
	Detected in version: 82.50.1002
5234750	**Description:** The Plug Voltage sensor does not read the value from the second module.
	Workaround: N/A
	Keywords: Plug Voltage sensor
	Detected in version: 82.50.1002
5226500	**Description:** For STC PACKET_OP, immediate ICRC recalculation is supported only when the reparse bit is enabled. If ICRC recalculation is requested without enabling reparse, the resulting ICRC may differ from the expected value.
	Workaround: Enable STC PACKET_OP reparse whenever ICRC recalculation is requested.
	Keywords: STC PACKET_OP, ICRC recalculation
	Detected in version: 82.50.1002
5230758 / 4553130 / 5231785	**Description:** For LPO cables, link-up may fail after the port on the peer device is toggled.
	Workaround: Toggle the port on the SuperNIC side.
	Keywords: LPO cables
	Detected in version: 82.50.1002
4901395 / 4919262	**Description:** When PXE is enabled for both Legacy (FlexBoot) and UEFI drivers, the UEFI driver supports only a single TFTP session after the DHCP handshake.
	Workaround: For scenarios that require multiple concurrent TFTP sessions, disable PXE filters via the HCA HII menu.
	Keywords: PXE, Filter enabled, UEFI
	Detected in version: 82.48.1000
5072222	**Description:** In DirectNIC and SocketDirect topologies, PCIe TLP Analyzer visibility is limited to the host's management domain. Traffic passing through an unmanaged host link, such as a DirectNIC connection to a GPU, is not visible to the local tracer.
	Workaround: N/A
	Keywords: PCIe TLP Analyzer

Internal Ref.	Issue
	Detected in version: 82.49.1014
4893639	**Description:** Direct Memory Access (DMA) protection is not supported in the ipxe.efi driver.
	Workaround: N/A
	Keywords: DMA
	Detected in version: 82.49.1014
4835832	**Description:** In rare cases, certain module types may experience link-up failures.

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

	Keywords: Cables
	Detected in version: 82.48.1000
4809493 / 4804815 / 4884845	**Description:** Running lane margin tests may disrupt traffic, cause connection drops, or degrade performance.
	Workaround: 1. Disable the drivers before starting the Lane Margin test. 2. Set an error threshold of 2. 3. If the connection drops or performance degrades, re-run the process.
	Keywords: Lane Margin
	Detected in version: 82.48.1000
4806964 / 4829338 / 4843305 / 4857738	**Description:** When bringing up 800G_8x on ConnectX-9 with Spectrum-4 switch systems, switch firmware version xx.201.3916-020 or later is required; earlier versions may fail to link up at 800G_8x.
	Workaround: N/A
	Keywords: Linkup, cables, Spectrum-4, 800G_8x
	Detected in version: 82.48.1000

7 PreBoot Drivers (FlexBoot/UEFI)

7.1 FlexBoot Changes and New Features

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

7.2 UEFI Changes and Major New Features

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

8 Validated and Supported Cables and Switches

8.1 Validated and Supported Cables and Modules

8.1.1 Cables Lifecycle Legend

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

8.1.2 XDR / 800GbE / 1600GbE Cables

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Marketing Description	LifeCycle Phase
OSFP	XDR	1600GE	980-9IAM1-00X001	NVIDIA Linear Active Copper Cable, 1600Gbps to 1600Gbps, OSFP, 1.1m, RHS to RHS, NIC to NIC	P-Rel
OSFP	XDR	1600GE	980-9IAM6-00X001	NVIDIA Active copper cable, 1600Gbps to 1600Gbps, OSFP, 1.1m, RHS to RHS	Prototype
OSFP	XDR	1600GE	980-9IAM6-00X002	NVIDIA Active copper cable, 1600Gbps to 1600Gbps, OSFP, 2m, RHS to RHS	Prototype
OSFP	XDR	1600GE	980-9IAM6-00X003	NVIDIA Active copper cable, 1600Gbps to 1600Gbps, OSFP, 3m, RHS to RHS	Prototype
OSFP to 2xOSFP	XDR	1600GE	980-9IAO5-00X001	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 2x800Gb/s, OSFP to 2xOSFP, 1m, IHS to RHS	P-Rel
OSFP to 2xOSFP	XDR	1600GE	980-9IAO5-00X002	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 2x800Gb/s, OSFP to 2xOSFP, 2m, IHS to RHS	P-Rel

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Marketing Description	LifeCycle Phase
OSFP to 2xOSFP	XDR	1600GE	980-9IAO5-00X01A	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 2x800Gb/s, OSFP to 2xOSFP, 1.5m, IHS to RHS	P-Rel
OSFP to 2xOSFP	XDR	1600GE	980-9IAO6-00X001	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 2x800Gb/s, OSFP to 2xOSFP, 1m, RHS to RHS	Prototype
OSFP to 2xOSFP	XDR	1600GE	980-9IAO6-00X002	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 2x800Gb/s, OSFP to 2xOSFP, 2m, RHS to RHS	Prototype
OSFP to 4xOSFP	XDR	1600GE	980-9IAQ5-00X001	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 4x400Gb/s, OSFP to 4xOSFP, 1m, IHS to RHS	Prototype
OSFP to 4xOSFP	XDR	1600GE	980-9IAQ5-00X002	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 4x400Gb/s, OSFP to 4xOSFP, 2m, IHS to RHS	Prototype
OSFP to 4xOSFP	XDR	1600GE	980-9IAQ5-00X01A	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 4x400Gb/s, OSFP to 4xOSFP, 1.5m, IHS to RHS	Prototype
OSFP to 4xOSFP	XDR	1600GE	980-9IAQ6-00X001	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 4x400Gb/s, OSFP to 4xOSFP, 1m, RHS to RHS	Prototype
OSFP to 4xOSFP	XDR	1600GE	980-9IAQ6-00X002	NVIDIA Linear Active copper splitter cable, 1600Gb/s to 4x400Gb/s, OSFP to 4xOSFP, 2m, RHS to RHS	Prototype
OSFP	XDR	NA	980-9IAT0-00XM00	NVIDIA single port transceiver, DR4, 800Gbps, OSFP, MPO(APC), 1310nm SMF, up to 500m, RHS, FRO, XDR Only	P-Rel
OSFP	XDR	NA	980-9IAU0-00XM00	NVIDIA twin port transceiver, 2xDR4, 1600Gbps, OSFP, 2xMPO(APC), 1310nm SMF, up to 500m, RHS, FRO, Gen2, XDR Only	P-Rel
OSFP	XDR	1600GE	980-9IAU1-00XM00	NVIDIA twin port transceiver, 2xDR4, 1600Gbps, OSFP, 2xMPO(APC), 1310nm SMF, up to 500m, RHS, FRO, Gen2, XDR + Ethernet	P-Rel
OSFP	XDR	800GE	980-9IAY0-00XM00	NVIDIA single port transceiver, DR4, 800Gbps, OSFP, MPO(APC), 1310nm SMF, up to 500m, RHS, FRO, Gen2, XDR + Ethernet	Prototype

8.1.3 NDR / 400GbE / 800GbE Cables

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	Lifecycle Phase
OSFP	NA	1600GE	980-9IAJ0-00XM00	N/A	NVIDIA twin port transceiver, 2xDR4, 1600Gbps, OSFP, 2xMPO(APC), 1310nm SMF, up to 500m, RHS, TRO, Ethernet Only	P-Rel
OSFP	NA	800GE	980-9IAK0-00XM00	N/A	NVIDIA single port transceiver, DR4, 800Gbps, OSFP, MPO(APC), 1310nm SMF, up to 500m, RHS, TRO, Ethernet Only	Early BOM
OSFP	NDR	NA	980-9I600-00N003	MCA4J80-N003-FLT	Active copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 3m, flat to flat	MP
OSFP	NDR	NA	980-9I601-00N003	MCA4J80-N003-FTF	NVIDIA Active copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 3m, flat to fin	MP
OSFP	NDR	800GE	980-9IA0G-00N001	MCP4Y10-N001-FLT	NVIDIA passive Copper cable, 800(2x400)Gbps, OSFP to OSFP, 1m, flat to flat	LTB [MP]
OSFP	NDR	800GE	980-9IA0J-00N002	MCP4Y10-N002-FLT	NVIDIA passive Copper cable, 800(2x400)Gbps, OSFP to OSFP, 2m, flat to flat	LTB [MP]
OSFP	NDR	800GE	980-9IA0L-00N00A	MCP4Y10-N00A-FLT	NVIDIA passive Copper cable, 800(2x400)Gbps, OSFP to OSFP, 0.5m, flat to flat	LTB [MP]
OSFP	NDR	800GE	980-9IA0R-00N01A	MCP4Y10-N01A-FLT	NVIDIA passive Copper cable, 800(2x400)Gbps, OSFP to OSFP, 1.5m, flat to flat	LTB [MP]
OSFP to 2xOSFP	NDR	800GE	980-9I432-00N001	MCP7Y00-N001	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP,1m, fin to flat	P-Rel
OSFP to 2xOSFP	NDR	800GE	980-9I433-00N001	MCP7Y00-N001-FLT	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP,1m, flat to flat	LTB [P-Rel]
OSFP to 2xOSFP	NDR	800GE	980-9I924-00N002	MCP7Y00-N002	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 2m, fin to flat	P-Rel
OSFP to 2xOSFP	NDR	800GE	980-9I925-00N002	MCP7Y00-N002-FLT	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 2m, flat to flat	LTB [P-Rel]
OSFP to 2xOSFP	NDR	800GE	980-9I92N-00N003	MCP7Y00-N003	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 2x400Gbps, OSFP to 2xOSFP, 3m, fin to flat	P-Rel

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	Lifecycle Phase
OSFP to 4xOSFP	NDR	800GE	980-9I75E-00N001	MCP7Y50-N001	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 1m, fin to flat	P-Rel
OSFP to 4xOSFP	NDR	800GE	980-9I75F-00N001	MCP7Y50-N001-FLT	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 1m, flat to flat	LTB [P-Rel]
OSFP to 4xOSFP	NDR	800GE	980-9I46G-00N002	MCP7Y50-N002	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 2m, fin to flat	P-Rel
OSFP to 4xOSFP	NDR	800GE	980-9I46H-00N002	MCP7Y50-N002-FLT	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 2m, flat to flat	LTB [P-Rel]
OSFP to 4xOSFP	NDR	800GE	980-9I46T-00N003	MCP7Y50-N003	NVIDIA passive copper splitter cable, 800(2x400)Gbps to 4x200Gbps, OSFP to 4xOSFP, 3m, fin to flat	P-Rel
QSFP112	NDR	400GE	980-9I693-00NS00	MMA1Z00-NS400	NVIDIA single port transceiver, 400Gbps, QSFP112, MPO12 APC, 850nm MMF, up to 50m, flat top	P-Rel
QSFP112	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
OSFP	NDR	800GE	980-9I51A-00NS00	MMA4Z00-NS-FLT	NVIDIA twin port transceiver, 800(2x400)Gbps, OSFP, 2xMPO12 APC, 850nm MMF, up to 50m, flat top	MP
OSFP	NDR	400GE	980-9I51S-00NS00	MMA4Z00-NS400	NVIDIA single port transceiver, 400Gbps, OSFP, MPO12 APC, 850nm MMF, up to 50m, flat top	MP
OSFP	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
QSFP112	NDR	400GE	980-9I068-00NM00	MMS1X00-NS400	NVIDIA single port transceiver, 400Gbps, QSFP112, MPO, 1310nm SMF, up to 500m, flat top	P-Rel
OSFP	NDR	800GE	980-9I301-00NM00	MMS4X00-NM-FLT	NVIDIA twin port transceiver, 800(2x400)Gbps, OSFP, 2xMPO12 APC, 1310nm SMF, up to 500m, flat top	P-Rel
OSFP	NDR	800GE	980-9I302-00NM00	MMS4X00-NM-HGX	NVIDIA twin port transceiver, 800(2x400)Gbps, OSFP, 2xMPO12 APC, 1310nm SMF, up to 500m, flat top - For HGX Rubin NVL8	P-Rel

Form Factor	IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	Lifecycle Phase
OSFP	NDR	400GE	980-9I31N-00NM00	MMS4X00-NS400	NVIDIA single port transceiver, 400Gbps, OSFP, MPO12 APC, 1310nm SMF, up to 100m, flat top	MP

8.2 Validated and Supported Optical Fibers Cables

MFP7E	x0-	Nxxx
NVIDIA passive fiber cable, MPO12 APC	**Fiber Type** 10 = MMF 20 = **1:2** splitter MMF 30 = SMF 40 = **1:2** splitter SMF	**Fiber Length** MMF: 3, 5, 7, 10, 15, 20, 30, 35, 40, 50m SMF: 3, 5, 7, 10, 15, 20, 30, 40, 50, 60, 70

8.3 Tested Switches

8.3.1 XDR / 800GbE Switches

Speed	NVIDIA SKU	Legacy OPN	Description	LifeCycle Phase
XDR	920-9B34F-00RX-FS0	Q3200-RA	Quantum-3 based Two-Adjoining XDR InfiniBand Switches, Q3200-RA, 2U, with 36 XDR Ports over 18 OSFP cages per Switch, 4 Power Supplies (Power Cords Not Included), Standard Depth, Managed, C2P Airflow, Rail Kit	P-Rel
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	P-Rel
800GbE	920-9N42F-00RI-7C0	SN5600	NVIDIA Spectrum-4 based 800GbE 2U Open Ethernet switch with Cumulus Linux Authentication, 64 OSFP ports and 1 SFP28 port, 2 power supplies (3KW, AC), x86 CPU, Secure-boot, standard depth, C2P airflow, Tool-less Rail Kit	LTB
800GbE	920-9N42F-00RI-3C1	SN5610	Nvidia Spectrum-4 based 800GbE 2U Open Ethernet switch with Cumulus Linux Authentication, 64 OSFP ports and 2 SFP28 ports, 4 AC PSUs, Secure-boot, standard depth, Connector-to-Power Airflow, Tool-less Rail Kit	P-Rel
0GbE	920-9N52F-00RB-7C0	SN5640	NVIDIA Spectrum-5 based 400GbE 2U Open Ethernet switch with Cumulus Linux Authentication, 64 OSFP and 2 SFP28 ports, 512 MACs, 4 AC Power-Supplies, AMD CPU, Secure-boot, Connector-to-Power Airflow.	P-Rel

9 Release Notes History

9.1 Changes and New Feature History

Feature/Change	Description
82.49.1014	
PCC/ZTR-RTT Congestion-Control Histogram Collection	Added support for Congestion-Control histogram collection in the PCC/ZTR-RTT algorithm. After enabling this capability, customers can read RATE and RTT histogram counters for PCC-managed flows.
Precoded Data Reception Configuration	Added automatic configuration for Tx precoding at 200G/lane based on the transceiver's Command Data Block (CDB). The default setting for Rx precoding remains ON.
PCC TopoAware PLB (probe_metadata Behavior)	**PCC_CONFIG.probe_metadata** is now supported as a global setting that applies uniformly across all probe slots. Updates made via CREATE/MODIFY (regardless of probe_type_slot) propagate to all slots.
ZTR_RTTCC Tunable Probe Timeout	Added a new parameter to the ZTR_RTTCC algorithm to define the probe-packet timeout threshold.
PTP/PPS/SyncE over I2C DC-TCXO	Firmware now enables full timing functionality on ConnectX-9: PTP (Precision Time Protocol, IEEE 1588) time synchronization, PPS (Pulse Per Second) input/output, and SyncE (Synchronous Ethernet) frequency synchronization on platforms that use an external DC-TCXO. This is done by adding I2C-based control of the DC-TCXO so the NIC can steer/discipline its reference clock for accurate and stable timing.
Link up/Down dmesg Messages	When powering down the OSFP port, the link up/down dmesg messages were previously sent only to the primary ASIC, they are now sent to the secondary ASIC as well.
PSP Transport Packet Reformat	Added support for a new packet reformat type for PSP transport packets. This reformat removes the PSP trailer and the UDP + PSP transport headers, and updates the IP Protocol field based on the PSP next_header value.
Persistent CRDT Token	CRDT token behavior has changed and is now persistent. Previously, debug firmware was allowed only if a CRDT token was applied temporarily, meaning it was erased after reset, or if another debug firmware was already running. Once a user moved to production firmware, a new token was required to switch back to debug firmware. With the new behavior, the token is saved to flash, meaning it remains readable via MDSR even after reset. Once installed, the token allows the user to move between debug and production firmware until it is explicitly revoked using the MDSR set command.
Extended Recovery Support for XDR Optics	Added support for extended recovery on XDR optics when the peer device also supports the capability. This enhancement improves link recovery behavior by allowing both sides of the connection to use the extended recovery mechanism, helping increase resiliency in supported XDR optic configurations.
Bug Fixes	See Bug Fixes section.

For the list of changes and bug fixes in earlier versions, see Release Notes History section.

Feature/Change	Description
82.48.1000	
General	This is the initial firmware release for the NVIDIA® ConnectX®-9 SuperNIC. ConnectX-9 provides the same feature set as the ConnectX-8 adapter card. For a complete list of ConnectX-8 firmware features, refer to the ConnectX-8 Firmware Release Notes. The features described below are new and are provided in addition to the existing ConnectX-8 feature set.
Lane Margin Test (LMT)	Lane Margin Test (LMT) is a PCIe diagnostic tool that performs controlled per-lane margining in the voltage and timing domains to evaluate link robustness. It helps identify weak or unstable lanes and supports signal-integrity validation, platform bring-up, and long-term reliability analysis. **Notes:** • NVIDIA recommends setting the error threshold to **2**. LMT is a destructive test, link drops or performance degradation may occur; if this happens, rerun the test. • NVIDIA also recommends disabling drivers before starting the Lane Margin Test.
T10 Data Integrity (DIF)	This capability provides block-level data protection by appending a Data Integrity Field (DIF) to each block, including checksums and metadata. The firmware automatically verifies these fields during reads, writes, and data transfers, helping to detect and prevent data corruption, ensure end-to-end data integrity, and increase reliability for critical storage workloads. By enabling T10 DIF, systems can achieve more robust data protection, reducing the risk of silent errors and improving overall storage confidence.
Cyclic Redundancy Check (CRC)	CRC provides a fast and reliable method to detect data corruption by generating a checksum for each block of data. The firmware automatically verifies CRC values during reads, writes, and transfers, helping to detect and prevent silent data errors, ensure end-to-end data integrity, and increase overall system reliability. By leveraging CRC, storage systems can maintain high confidence in critical workloads and minimize the risk of data loss.
Transport Layer Security (TLS) Handshake	The TLS handshake establishes a secure communication session by authenticating endpoints, negotiating cryptographic algorithms, and securely exchanging encryption keys before data transfer begins. This process ensures confidentiality of data in transit, protection against tampering and man-in-the-middle attacks, and trusted identity verification, enabling secure and reliable communication for sensitive and high-value workloads.
NVMe over TCP Acceleration	This capability offloads and accelerates NVMe/TCP data path processing to hardware, reducing CPU overhead and latency while increasing throughput. By streamlining NVMe command processing and TCP handling, it enables higher IOPS, lower and more predictable latency, and improved host CPU efficiency, making it ideal for scalable, high-performance, and cloud-based storage deployments.
DPA Process Limit Update	The system-wide limit for DPA processes has been reduced to **30**. This total includes both user processes across all GVMIs and internal ProgCC processes. The `max_dpa_processes` value reported to the user is calculated as: `max_dpa_processes=30-number_of_progcc_processes`

Feature/Change	Description
82.48.1000	
Host Rate Limiting Support Above 255 Gbps	Host rate limiting has been extended to support bandwidth values above 255 Gbps. To remove the previous cap, a new `max_bw_value_msb` field was added to `est_global`, providing additional MSB bits to represent higher bandwidth values. With this enhancement, firmware and host tooling can correctly configure and report rate limits beyond 255 Gbps on high-speed links.
PLDM PDR Repository Change Event Support	PLDM now supports the PDR Repository Change event type, enabling notification to the BMC when PDRs change. With this flow, the BMC can detect cable insertion/removal events. Refer to DSP0248 for details.
BMC Write-Protection Check on Firmware Update Failure	Added a validation step during firmware updates to detect whether the BMC is asserting write protection, helping diagnose and prevent update failures.
CMB SysRAM Window Partitioning for PSC Isolation	CMB configuration enforces secure access to System RAM (SysRAM) by defining memory windows and partitions that isolate the Platform Security Controller (PSC) from other system components.
Physical Access Monitor and Protection	The Physical Monitor feature helps protect the platform against threats from attackers with physical access to the system. It monitors for physical-access events and enforces safeguards that prevent or restrict sensitive operations, reducing the risk of exposing secure information or compromising security-critical functionality.
Parallel Save/Load Support for VF Migration	Added support for running save and load operations in parallel, enabling multiple contexts (e.g., multiple VFs) to be checkpointed and restored concurrently instead of serially. This reduces overall migration time and improves scalability in environments that need to migrate or recover many VFs at once.
NVGRE VSID Modify-Header Support	Extended packet modify-header operations to support `set` and `copy` actions on the NVGRE VSID (Virtual Subnet Identifier). A new field, `TUNNEL_HDR_DW_2` (0x84), enables dynamic VSID modification, adding header rewrite support for NVGRE tunnel traffic in addition to existing filtering capabilities.
mlxlink	mlxlink `show_links` now reports the full PCIe identifier (domain/segment + BDF), improving device-to-link mapping and avoiding ambiguous/duplicate BDF entries on multi-domain systems.

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.49.1014	• License • 3rd Party Notice • 3rd Party Unify Notice • HCA Firmware EULA
DOCA-Host	3.4.0	• License • 3rd Party Notice • 3rd Party Unify Notice
MFT FreeBSD	4.36.0-147	• License • 3rd Party Notice • 3rd Party Unify Notice
MFT Linux		• License • 3rd Party Notice • 3rd Party Unify Notice
MFT VMware		• License • 3rd Party Notice • 3rd Party Unify Notice
MFT Windows		• License • 3rd Party Notice • 3rd Party Unify Notice

Notice

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

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

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

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

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

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

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

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

NVIDIA