NVIDIA ConnectX-8 SuperNIC Firmware Release Notes v40.44.1036

Table of Contents

1 Release Notes Update History

Version	Date	Description
40.44.1036	February 2025	Initial release of this Release Notes version.

2 Overview

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

2.1 Firmware Download

Please visit Firmware Downloads.

3 Firmware Compatible Products

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

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

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

[2]. Speed that supports PAM4 mode only.

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

3.1 Supported Devices

Orderable Part Number (OPN)	PSID	Description
900-9X81E-00EX-ST0	MT_0000001167	NVIDIA ConnectX-8 C8180 HHHL SuperNIC; 800Gbs XDR IB (default mode) / 2x400GbE; Single-cage OSFP; PCIe 6 x16 with x16 PCIe Socket Direct/ Multi Host Extension option; Crypto Enabled; Secure Boot Enabled
900-9X81Q-00CN-ST0	MT_0000001222	NVIDIA ConnectX-8 C8240 HHHL SuperNIC; 400GbE (default mode) / 400Gb/s IB; Dual-port QSFP112; PCIe 6 x16 with x16 PCIe Socket Direct / Multi Host Extension option; Crypto Enabled; Secure Boot Enabled

Driver Software, Tools and Switch Firmware

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

	Supported Version
ConnectX-8 Firmware	40.44.1036 / 40.44.0212 / 40.44.0208
DOCA-HOST	2.10.0 / 2.9.1 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.
WinOF-2	25.1.50020 / 24.10.50010 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.
MFT	4.31.0-149 / 4.30.2-20 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.
mstflint	4.31.0-149 / 4.30.2-20 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.

	Supported Version
FlexBoot	3.7.500
UEFI	14.37.14
MLNX-OS	3.11.1004 onwards
NVOS	25.02.1934 onwards
Cumulus	5.11.0.0026 onwards
NVIDIA Quantum-3 Firmware	35.2014.2038 onwards
NVIDIA Quantum-2 Firmware	31.2014.2082 onwards

4 Changes and New Features

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

Feature/Change	Description
40.44.1036	
Static Split 8x100G ConnectX-8 to Spectrum-4 with SM Modules	A static split of 8x100G channels from a ConnectX-8 SuperNIC to a Spectrum-4 switch allows the system to use Single Mode (SM) optical modules for high-speed data transmission across a long-distance fibber link. This setup is typically used in high-performance networks where there is a need for high throughput (e.g., 800G in total bandwidth) with low latency, such as in data centers or high-performance computing environments.
DOCA Telemetry	DOCA Telemetry enables users to monitor and collect data related to the performance, health, and behavior of systems or applications running on DOCA. To optimize for a faster sampling period, it is recommended to configure all PCIe-related Diagnostic Data IDs sequentially, one after another to prevent a prolonged sampling period.
PCIe Switch fwreset	Added support for a new synchronized flow, including a tool and driver, to perform a fwreset on setups with a PCIe switch configuration.
PTP	Unified PTP is now supported across different VFs on the same PF.
Dual-Mode Temperature Compensated Crystal Oscillator (DC-TCXO) and Synchronous Ethernet (SyncE) Source	DC-TCXO is used now as the source of timing for SyncE, providing an accurate and stable clock for the synchronized operation of network devices that rely on Ethernet for timing.
DPA Application Signing	Allows DOCA applications signed with OEM/NVIDIA certificate private keys to be loaded onto the DPA engine, after the OEM/NVIDIA root certificates are installed on the NIC.
Data-Path Accelerator (DPA)	The DPA hardware version is now exposed as a new capability, labeled "dpa_platform_version."
Block SMP Traffic	Added a new NV config (SM_DISABLE, default 0) which, when enabled, blocks SMP traffic that does not originate from the SM.
Dynamic Long Cables	Added the ability to set cable length as a parameter in the PFCC access register. The cable length is used in the calculation of RX lossless buffer parameters, including size, Xoff, and Xon thresholds.
Bug Fixes	See *Bug Fixes in this Firmware Version* section.

4.1 Customer Affecting Changes

4.1.1 Changes in This Release

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

Introduced in Version	Description
40.44.1036	In ConnectX-8, the DIAG_COUNTER interface has been changed from the following set of commands: • SET_DIAGNOSTIC_PARAMS • QUERY_DIAGNOSTIC_PARAMS • QUERY_DIAGNOSTIC_COUNTERS • ICMD_SET_DIAGNOSTIC_PARAMS • ICMD_QUERY_DIAGNOSTIC_PARAMS • ICMD_QUERY_DIAGNOSTIC_COUNTERS to: • DIAG_DATA_OWNERSHIP • DIAG_DATA_PARAMS_CONTEXT • DIAG_DATA_ID_LIST • DIAG_DATA_QUERY The old interface will now return zero values when queried.

4.1.2 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
N/A	N/A

4.1.4 Discontinued Features

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

- Ethernet:
  - T10 Data Integrity Field (DIF)

    - CRC
    - Transport Layer Security (TLS) handshake
    - NVMe over TCP acceleration
- InfiniBand:
    - FDR and lower speeds

> ⚠ For inquiries regarding mitigation, please contact NVIDIA Support.

4.2 Declared Unsupported Features

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

- Zero Touch Tuning (ZTT)
- NC-SI Path-Through
- Lane margin
- Hot reset
- DRS
- LAG Bonding with Q-Affinity
- ISSU

5 Bug Fixes in this Firmware Version

Internal Ref.	Issue
4087432	**Description:** Increased the RX lossless buffer size to delay the transmission of Pause/PFC frames during NIC congestion.
	Keywords: RX lossless buffer size
	Detected in version: 40.44.0208
	Fixed in Release: 40.44.1036
4230051	**Description:** Fixed an issue with configuring the Log FIFO in the operational state.
	Keywords: Log FIFO
	Detected in version: 40.44.0208
	Fixed in Release: 40.44.1036
4249985	**Description:** Fixed an issue where error messages were not sent to the host if the NIC was an EP behind the embedded switch.
	Keywords: DSP switch, error message
	Detected in version: 40.44.0208
	Fixed in Release: 40.44.1036
4209411	**Description:** When querying flow counters with a large number of counters, the process took longer than the dead iris timestamp. The timestamp has been extended to address this issue.
	Keywords: Flow counters
	Detected in version: 40.44.0208
	Fixed in Release: 40.44.1036
4176679 / 4119723	**Description:** Fixed an issue that prevented bandwidth from reaching its line rate when sending RoCE traffic using 1 or 2 QPs over 100GbE or 400GbE link speed with Congestion Control enabled.
	Keywords: RoCE, Congestion Control, 400GbE, performance
	Detected in version: 40.44.0208
	Fixed in Release: 40.44.1036

6 Known Issues

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

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

VF+SF Network Function Limitations in Switchdev Mode

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

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

The below are limitations related to ConnectX-8 only.

Internal Ref.	Issue
4163634	**Description:** When connecting a Quantum-3 switch system (with a port split into 8 ports) to a ConnectX-8 single port SuperNIC, the link will not be established.
	Workaround: Configure the Quantum-3 switch system port to be split into 2 or 4 ports, or set the ConnectX-8 to operate in multiplane mode.
	Keywords: Port split, Quantum-3
	Detected in version: 40.44.0212
4230775	**Description:** Due to a known issue, telemetry rate must be set to lower than 3 minutes.
	Workaround: N/A
	Keywords: Telemetry rate
	Detected in version: 40.44.0212
4161849	**Description:** The Green LED remains solid and does not blink when running traffic.
	Workaround: N/A
	Keywords: Green LED
	Detected in version: 40.44.0212
4230775	**Description:**
	Workaround: N/A
	Keywords:
	Detected in version: 40.44.0212
-	**Description:** Although ConnectX-8 SuperNIC is defined to work at Gen6 x16 (default) or Gen5 x32, in firmware v40.44.0204, the default configuration is Gen5 x32. Changing between the modes is done by an NVConfig command.
	Workaround: N/A
	Keywords: Gen6, Gen5, PCIe

Internal Ref.	Issue
	Detected in version: 40.44.0208
4038325 4031430 4038341 4046105	**Description:** Connecting to systems with NRZ speeds of 1,10, 25, 40, 50, or 100Gb/s is not supported in the current release.
	Workaround: N/A
	Keywords: NRZ, Connectivity
	Detected in version: 40.44.0208
4158184	**Description:** The Lane Error Status may occasionally appear in Configuration space. It can be safely ignored as it does not have any impact on device performance. Users are encouraged to monitor their systems, but this condition does not warrant any immediate action unless other issues arise.
	Workaround: N/A
	Keywords: Lane Error Status
	Detected in version: 40.44.0208
4208960	**Description:** A packet may be parsed incorrectly, if a driver uses the `header_length_field_mask` when creating a `PARSE_GRAPH_NODE` object, and the mask value is not composed of continuous bits or does not commence at the least significant bit.
	Workaround: Insert the `header_length_field_mask` with continuous bits and commence at the least significant bit.
	Keywords: PARSE GRAPH NODE, Flex Parser
	Detected in version: 40.44.0208
4201405	**Description:** Upgrading to firmware 40.44.0xxx from any previous Engineering Sample version requires power cycling the driver and not just resetting it (using mlxfwreset).
	Workaround: N/A
	Keywords: Upgrade, power cycle, reset
	Detected in version: 40.44.0208

7 PreBoot Drivers (FlexBoot/UEFI)

7.1 FlexBoot Changes and New Features

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

7.2 UEFI Changes and Major New Features

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

8 Validated and Supported Cables and Switches

8.1 Validated and Supported Cables and Modules

8.1.1 Cables Lifecycle Legend

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

8.1.2 XDR / 800GbE 1600GbE Cables

IB Data Rate	Eth Data Rate	NVIDIA SKU	Marketing Description	LifeCycle Phase
XDR	800GE	980-9IAT0-00XM00	NVIDIA single port transceiver for ConnectX-8 Mezz Card, 800Gbps,OSFP, MPO, 1310nm SMF, EML, up to 500m, RHS	EVT

8.1.3 NDR / 400GbE / 800GbE Cables

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	LifeCycle Phase
NDR	NA	980-9I81B-00N004	MCA7J65-N004	NVIDIA Active copper splitter cable, IB twin port NDR 800Gb/s to 2x400Gb/s, OSFP to 2xQSFP112, 4m	P-Rel
NDR	NA	980-9I81C-00N005	MCA7J65-N005	NVIDIA Active copper splitter cable, IB twin port NDR 800Gb/s to 2x400Gb/s, OSFP to 2xQSFP112, 5m	P-Rel
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	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	NA	980-9I693-00NS00	MMA1Z00-NS400	NVIDIA single port transceiver, 400Gbps,NDR, QSFP112, MPO12 APC, 850nm MMF, up to 50m, flat top	P-Rel
NDR	NA	980-9I51S-00NS00	MMA4Z00-NS400	NVIDIA single port transceiver, 400Gbps,NDR, 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	Early BOM
NDR	NA	980-9I31N-00NM00	MMS4X00-NS400	NVIDIA single port transceiver, 400Gbps,NDR, OSFP, MPO12 APC, 1310nm SMF, up to 100m, flat top	MP

8.1.4 HDR / 200GbE Cables

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	LifeCycle Phase
HDR	200GE	980-9I549-00H002	MCP1650-H002E26	Nvidia Passive Copper cable,up to 200Gbps, QSFP56 to QSFP56, 2m	HVM
HDR	200GE	980-9I54A-00H00A	MCP1650-H00AE30	Nvidia Passive Copper cable, up to 200Gbps, QSFP56 to QSFP56, 0.5m	HVM
HDR	200GE	980-9I46K-00H001	MCP7Y60-H001	NVIDIA passive copper splitter cable, 400(2x200)Gbps to 2x200Gbps, OSFP to 2xQSFP56, 1m, fin to flat	MP
HDR	200GE	980-9I46L-00H002	MCP7Y60-H002	NVIDIA passive copper splitter cable, 400(2x200)Gbps to 2x200Gbps, OSFP to 2xQSFP56, 2m, fin to flat	MP

8.1.5 EDR / 100GbE Cables

IB Data Rate	Eth Data Rate	NVIDIA SKU	Legacy P/N	Marketing Description	LifeCycle Phase
EDR	NA	980-9I62Q-00E001	MCP1600-E001E30	Mellanox Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 1m, Black, 30AWG	HVM
EDR	NA	980-9I62Z-00E005	MCP1600-E005E26	Mellanox Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 5m, Black, 26AWG	HVM

8.1.6 Supported 3rd Party Cables and Modules

Speed	Cable OPN	Description
800GbE	T-RS8CNT-NMT	Innolight 800G DR8 OSFP RHS, dual MPO-12/APC optical connectors

8.1.7 Tested Switches

8.1.8 XDR / 800GbE Switches

Speed	NVIDIA SKU	Legacy OPN	Description	LifeCycle Phase
XDR	920-9B34F-00RX-FS0	Q3200-RA	Quantum-3 based Two-Adjoining XDR InfiniBand Switches, Q3200-RA, 2U, with 36 XDR Ports over 18 OSFP cages per Switch, 4 Power Supplies (Power Cords Not Included), Standard Depth, Managed, C2P Airflow, Rail Kit	Prototype
XDR	920-9B36F-00RX-8S0	Q3400-RA	NVIDIA Quantum-3 based XDR InfiniBand Switch, Q3400-RA, 4U, 144 XDR Ports over 72 OSFP Cages, 8 Power Supplies (Power Cords Not Included), Standard Depth, Managed, C2P Airflow, Rail Kit	Prototype

8.1.9 NDR / 400GbE Switches

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

8.1.10 HDR / 200GbE Switches

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

9 Release Notes History

9.1 Changes and New Feature History

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

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

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

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

9.2 Bug Fixes History

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

Internal Ref.	Issue
4161303	**Description:** PCI ARCH counters are not supported. Other ARCH counters are supported, but not fully tested.
	Keywords: Counters
	Detected in version: 40.44.0208
	Fixed in Release: 40.44.0212
4202233	**Description:** Address Translation Service (ATS) is at Beta level. Enabling ATS from mlxconfig and stopping the driver can result at a call trace in dmesg.

Internal Ref.	Issue
	Keywords: ATS
	Detected in version: 40.44.0208
	Fixed in Release: 40.44.0212
4220173	**Description:** In firmware version 40.44.0208, the `'max_parse_graph_header_length_base_value'` field allows setting higher values than the recommended.
	Keywords: PARSE_GRAPH_NODE Capabilities Layout
	Detected in version: 40.44.0208
	Fixed in Release: 40.44.0212

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.44.1030	• HCA Firmware EULA • 3rd Party Unify Notice • License
OFED Drivers	25.01	• License • 3rd Party Notice
MFT FreeBSD	4.31.0-147	• 3rd Party Notice • License
MFT Linux		• 3rd Party Notice • License
MFT VMware		• 3rd Party Notice • License
MFT Windows		• 3rd Party Notice • License

Notice

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

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

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

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

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

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

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

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

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation and/