NVIDIA ConnectX-8 SuperNIC for OCP 3.0 User Manual

Table of Contents

1 About This Manual

This User Manual describes NVIDIA® ConnectX®-8 SuperNICs for Open Compute Project Spec 3.0. It provides details as to the interfaces of the board, specifications, required software and firmware for operating the board, and relevant documentation.

1.1 Ordering Part Numbers

The table below provides the ordering part numbers (OPN) for the available ConnectX-8 SuperNICs for OCP Spec 3.0.

Ordering Part Numbers

Model	NVIDIA SKU	Port Type	Supported Speed	Form Factor	PCIe Express	Multi-Host and Socket Direct Extension Option	Crypto	Secure Boot	Bracket Type	Cooling Mechanism	Lifecycle
C8180P	**900-9X85E-00EX-MJ0**	Single-cage OSFP	InfiniBand: 800Gbs XDR IB (default) Ethernet: 2x400GbE	OCP3.0 TSFF	PCIe 6.0 x16	✓	✓	✓	Internal lock Bracket	Air-cooled	Prototype
C8180P Partner Cooled	**900-9X85E-00EX-MP0**	Single-cage OSFP	InfiniBand: 800Gbs XDR IB (default) Ethernet: 2x400GbE	OCP3.0 TSFF	PCIe 6.0 x16	✓	✓	✓	Internal lock Bracket	Partner-Cooled	Prototype

1.2 Technical Support

Customers who purchased NVIDIA products directly from NVIDIA are invited to contact us through the following methods:

- URL: https://www.nvidia.com > Support
- E-mail: enterprisesupport@nvidia.com

Customers who purchased NVIDIA Global Support Services, please see your contract for details regarding Technical Support.

Customers who purchased NVIDIA products through an NVIDIA-approved reseller should first seek assistance through their reseller.

1.3

Anchor | reldocs

Related Documentation

DOCA-Host	User Manual describing DOCA Host features, performance, band diagnostic, tools content and configuration.
DOCA Documentation	User Manual describing DOCA SDK platform development.
WinOF-2 for Windows Documentation	User Manual describing WinOF-2 features, performance, Ethernet diagnostic, tools content and configuration.
MFT User Manual (NVIDIA Firmware Tools)	User Manual describing the set of MFT firmware management tools for a single node.
InfiniBand Specifications	InfiniBand Architecture Specification Release 1.2.1, Vol 2 - Release 1.4, Vol 2 - Release 1.5 and Vol 1 - Release 1.7.
IEEE Ethernet Specifications	IEEE Std 802.3 Specification
PCI-SIG Specifications	Industry Standard PCI Express Base and Card Electromechanical Specifications.
LinkX Interconnect Solutions	LinkX cables and transceivers are designed to maximize the performance of High-Performance Computing networks, requiring high-bandwidth, low-latency connections between compute nodes and switch nodes. NVIDIA offers one of the industry's most complete line of 25, 100, 200, and 400GbE in Ethernet and EDR, HDR, NDR and XDR including Direct Attach Copper cables (DACs), copper splitter cables, Active Optical Cables (AOCs) and transceivers in a wide range of lengths from 0.5m to 10km. In addition to meeting Ethernet and IBTA standards, NVIDIA tests every product in an end-to-end environment ensuring a Bit Error Rate of less than 1E-15.
Open Compute Project 3.0 Specifications	https://www.opencompute.org/
NVIDIA ConnectX-8 Electrical and Thermal Specifications	You can access the "NVIDIA ConnectX-8 SuperNIC Product Specifications" document either by logging into NVOnline or by contacting your NVIDIA representative.

When discussing memory sizes, MB and MBytes are used in this document to mean size in MegaBytes. The use of Mb or Mbits (small b) indicates size in MegaBits. IB is used in this document to mean InfiniBand. In this document, PCIe is used to mean PCI Express.

1.4 Revision History

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

2 Introduction

2.1 Product Overview

The NVIDIA® ConnectX®-8 SuperNIC™ is optimized to supercharge hyperscale AI computing workloads. With support for both InfiniBand and Ethernet networking at up to 800 gigabits per second (Gb/s), ConnectX-8 SuperNIC delivers high-speed, efficient network connectivity, significantly enhancing system performance for AI factories and cloud data center environments. With support for both InfiniBand and Ethernet networking at up to 800 gigabits per second (Gb/s), ConnectX-8 SuperNIC delivers high-speed, efficient network connectivity, significantly enhancing system performance for AI factories and cloud data center environments.

ConnectX-8 SuperNICs are offered in two form factors and various flavors: stand-up PCIe and Open Compute Project (OCP) Spec 3.0 cards. This user manual covers the OCP 3.0 cards, for the low-profile PCIe stand-up cards hardware user manual, please refer to ConnectX-8 SuperNIC User Manual.

2.2 System Requirements

Item	Description
PCI Express Slot	**In PCIe x16 Configuration** PCIe Gen6 @ 64GT/s through x16 edge connector
System Power Supply	Refer to Specifications

Operating System	• In-box drivers for major operating systems:
	  • Linux: RHEL, Ubuntu
	  • Windows
	• DOCA Host
	• OpenFabrics Windows Distribution (WinOF-2)

Connectivity	• Interoperable with 25/100/200/400 Gb/s Ethernet switches and SDR/EDR/HDR100/HDR/NDR/XDR InfiniBand switches • Passive copper cable with ESD protection • Powered connectors for optical and active cable support

2.3 Package Contents

Category	Qty	Item
Cards	1	ConnectX-8 SuperNIC for OCP 3.0
Accessories	1	Internal Lock Bracket

2.4 Features and Benefits

⚠ Make sure to use a PCIe slot capable of supplying the required power and airflow to the ConnectX-8 SuperNICs as stated in the Specifications chapter.

⚠ This section describes hardware features and capabilities. Please refer to the relevant driver and firmware release notes for feature availability.

<table>
<tr><td>PCI Express (PCIe)</td><td>• PCIe x16 configurations:
PCIe Gen6 (64GT/s) through x16 edge connector</td></tr>
<tr><td>InfiniBand Architecture Specification v1.7 compliant</td><td>ConnectX-8 delivers low latency, high bandwidth, and computing efficiency for high-performance computing (HPC), artificial intelligence (AI), and hyperscale cloud data center applications. ConnectX-8 is InfiniBand Architecture Specification v1.7 compliant.
InfiniBand Network Protocols and Rates
<table>
<tr><th rowspan="2">Protocol</th><th rowspan="2">Standard</th><th colspan="3">Rate (Gb/s)</th><th rowspan="2">Encoding</th></tr>
<tr><th>4x
(4 lane)
Port</th><th>2x
(2 lane)
Port</th><th>1x
(1 lane)
Port</th></tr>
<tr><td>800G XDR</td><td>IBTA Vol1 1.7</td><td>--</td><td>425</td><td>212.5</td><td>PAM4</td></tr>
<tr><td>NDR</td><td>IBTA Vol2 1.5</td><td>425</td><td>212.5</td><td>106.25</td><td>PAM4</td></tr>
<tr><td>HDR / HDR100</td><td>IBTA Vol2 1.4</td><td>212.5</td><td>106.25</td><td>106.25</td><td>PAM4</td></tr>
</table></td></tr>
</table>

<table>
<tr><td>Up to 400 Gigabit Ethernet</td><td>ConnectX-8 SuperNICs comply with the following IEEE 802.3 standards:
400GbE / 200GbE / 100GbE / 25GbE / 10GbE
<table>
<tr><th>Protocol</th><th>Ethernet Network Rate</th></tr>
<tr><td>IEEE802.3ck</td><td>100/200/400Gb/s Gigabit Ethernet
(Include ETC enhancement)</td></tr>
<tr><td>IEEE802.3cd
IEEE802.3bs
IEEE802.3cm
IEEE802.3cn
IEEE802.3cu</td><td>50/100/200/400Gb/s Gigabit Ethernet
(Include ETC enhancement)</td></tr>
<tr><td>IEEE 802.3bj
IEEE 802.3bm</td><td>100 Gigabit Ethernet</td></tr>
<tr><td>IEEE 802.3by
Ethernet Technology Consortium</td><td>25/50 Gigabit Ethernet</td></tr>
<tr><td>IEEE 802.3ba</td><td>40 Gigabit Ethernet</td></tr>
<tr><td>IEEE 802.3ae</td><td>10 Gigabit Ethernet</td></tr>
<tr><td>IEEE 802.3cb</td><td>2.5/5 Gigabit Ethernet
(For 2.5: support only 2.5 x1000BASE-X)</td></tr>
<tr><td>IEEE 802.3ap</td><td>Based on auto-negotiation and KR startup</td></tr>
<tr><td>IEEE 802.3ad
IEEE 802.1AX</td><td>Link Aggregation</td></tr>
<tr><td>IEEE 802.1Q
IEEE 802.1P VLAN tags and priority</td><td></td></tr>
<tr><td>IEEE 802.1Qau (QCN)
Congestion Notification
IEEE 802.1Qaz (ETS)
EEE 802.1Qbb (PFC)
IEEE 802.1Qbg
IEEE 1588v2
IEEE 802.1AE (MACSec)
Jumbo frame support (9.6KB)</td><td></td></tr>
</table></td></tr>
<tr><td>Memory Components</td><td>• SPI - includes 512Mbit SPI Quad Flash device.
• FRU EEPROM - Stores the parameters and personality of the SuperNIC. The EEPROM capacity is 128Kbit. FRU I2C address is (0x50) and is accessible through the PCIe SMBus. (Note: Address 0x58 is reserved.)</td></tr>
<tr><td>Overlay Networks</td><td>In order to better scale their networks, data center operators often create overlay networks that carry traffic from individual virtual machines over logical tunnels in encapsulated formats such as NVGRE and VXLAN. While this solves network scalability issues, it hides the TCP packet from the hardware offloading engines, placing higher loads on the host CPU. ConnectX-8 effectively addresses this by providing advanced NVGRE and VXLAN hardware offloading engines that encapsulate and de-capsulate the overlay protocol.</td></tr>
<tr><td>Quality of Service (QoS)</td><td>Support for port-based Quality of Service enabling various application requirements for latency and SLA.</td></tr>
<tr><td>Hardware-based I/O Virtualization</td><td>ConnectX-8 provides dedicated adapter resources and guaranteed isolation and protection for virtual machines within the server.</td></tr>
</table>

SR-IOV	ConnectX-8 SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server.
High-Performance Accelerations	• Vector collective operations offload • MPI_Alltoall offloads • Rendezvous protocol offload
Secure Boot	The secure boot process assures the booting of authentic firmware/software that is intended to run on ConnectX-8. This is achieved using cryptographic primitives using asymmetric cryptography. ConnectX-8 supports several cryptographic functions in its HW Root-of-Trust (RoT) that has its key stored in on-chip FUSES.
Secure Firmware Update	The Secure firmware update feature enables a device to verify digital signatures of new firmware binaries to ensure that only officially approved versions can be installed from the host, the network, or a Board Management Controller (BMC). The firmware of devices with “secure firmware update” functionality (secure FW), restricts access to specific commands and registers that can be used to modify the firmware binary image on the flash, as well as commands that can jeopardize security in general. For further information, refer to the MFT User Manual.
Host Management	ConnectX-8 technology maintains support for host manageability through a BMC. ConnectX-8 PCIe stand-up adapter can be connected to a BMC using MCTP over SMBus or MCTP over PCIe protocols as if it is a standard NVIDIA PCIe stand-up SuperNIC. For configuring the adapter for the specific manageability solution in use by the server, please contact NVIDIA Support. • Protocols: PLDM, NCSI • Transport layer – RBT, MCTP over SMBus, and MCTP over PCIe • Physical layer: SMBus 2.0 / I2C interface for device control and configuration, PCIe • PLDM for Monitor and Control DSP0248 • PLDM for Firmware Update DSP026 • IEEE 1149.6 • Secured FW update • FW Recovery • NIC reset • Monitoring and control • Network port settings • Boot setting
RDMA and RDMA over Converged Ethernet (RoCE)	ConnectX-8, utilizing IBTA RDMA (Remote Data Memory Access) and RoCE (RDMA over Converged Ethernet) technology, delivers low latency and high-performance over InfiniBand and Ethernet networks. Leveraging datacenter bridging (DCB) capabilities as well as ConnectX-8 advanced congestion control hardware mechanisms, RoCE provides efficient low-latency RDMA services over Layer 2 and Layer 3 networks.
NVIDIA PeerDirect™	PeerDirect™ communication provides high-efficiency RDMA access by eliminating unnecessary internal data copies between components on the PCIe bus (for example, from GPU to CPU), and therefore significantly reduces application run time. ConnectX-8 advanced acceleration technology enables higher cluster efficiency and scalability to tens of thousands of nodes.
CPU Offload	Adapter functionality enables reduced CPU overhead allowing more available CPU for computation tasks. • Flexible match-action flow tables • Open VSwitch (OVS) offload using ASAP^{2}® • Tunneling encapsulation/decapsulation
Cryptography Accelerations	ConnectX-8 supports IPSec, MACSec, and PSP cryptography acceleration. Connectx-8 SuperNIC hardware-based accelerations offload the crypto operations and free up the CPU, reducing latency and enabling scalable crypto solutions.

NVIDIA Multi-Host	NVIDIA® Multi-Host technology enables next-generation Cloud, Web 2.0 and high-performance data centers to design and build new scale-out heterogeneous compute and storage racks with direct connectivity between multiple hosts and the centralized network controller. This enables direct data access with the lowest latency to significantly improve densities and maximizes data transfer rates. For more information, please visit NVIDIA Multi-Host Solutions.

3 Supported Interfaces

This section describes the ConnectX-8 supported interfaces. Each numbered interface that is referenced in the figures is described in the following table with a link to detailed information.

3.1 ConnectX-8 Layout and Interface Information

The below figures show the component side of the NVIDIA ConnectX-8 SuperNIC. Each numbered interface that is referenced in the figures is described in the following table with a link to detailed information.

> The below figures are for illustration purposes only and might not reflect the current revision of the SuperNIC.

SuperNIC with Heatsink	Partner-Cooled SuperNIC
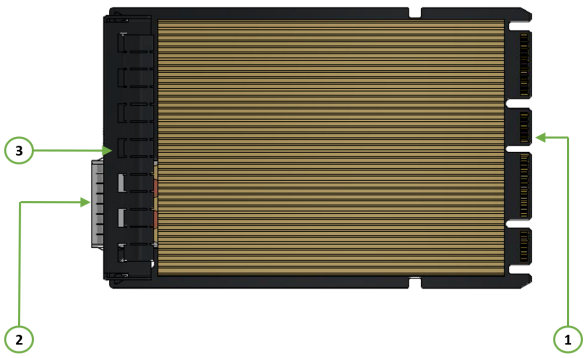	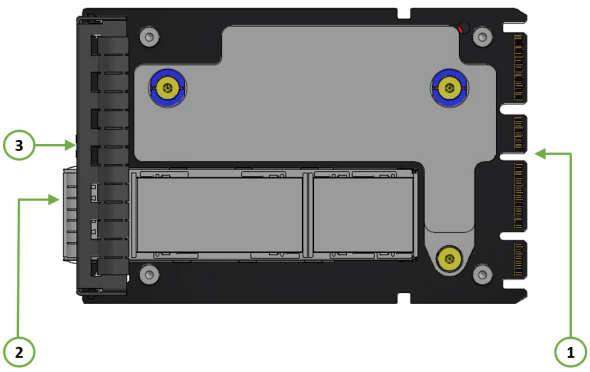

Item	Interface	Description
	ConnectX-8 IC	ConnectX-8 Integrated Circuit
2	PCI Express Interface	PCIe Gen 5.0 through x16 edge connector
3	Networking Interfaces	Network traffic is transmitted through the SuperNIC OSFP connectors. The networking connectors allow for the use of modules, optical and passive cable interconnect solutions
4	Networking Ports LEDs	Two I/O LEDs per port to indicate speed and link status
	FRU EEPROM	FRU EEPROM capacity 16KB
	SMBus Interface	Allows BMC connectivity using MCTP over SMBus or MCTP over PCIe protocols
	Voltage Regulators	Voltage supply pins that feed onboard regulators
	CPLD Interface	Controls the networking port logic LED (LED 1) and implements the OCP 3.0 host scan chain
	Heatsink	Dissipates the heat

3.2 Interfaces Detailed Description

3.2.1 ConnectX-8 IC

The ConnectX-8 family of IC devices delivers InfiniBand and Ethernet connectivity paired with best-in-class hardware capabilities that accelerate and secure cloud and data-center workloads.

3.2.1.1 NVIDIA Multi-Host™ Support

In addition to building exceptionally high bandwidth to the data center, the ConnectX-8 device enables leveraging this speed across the entire data center utilizing its NVIDIA Multi-Host feature.

Using its 16-lane PCI Express interface, a single ConnectX-8 device can provide 400GbE interconnect for up to four independent hosts without any performance degradation.

The figure below shows a ConnectX-8 device with NVIDIA Multi-Host connected to four separate hosts, each with a PCIe x4 interface, on one side to a switch on the other side.

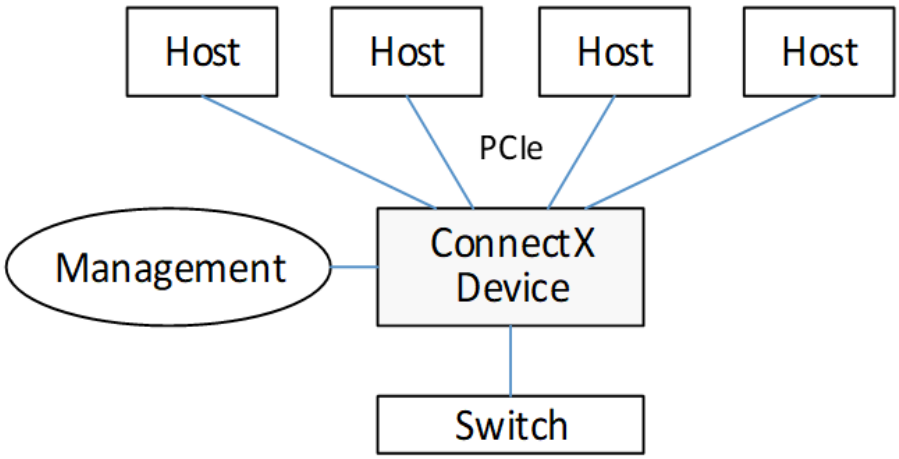

The below bifurcation are optional for the adapter's x16 PCIe interface:

- x1 PCIe x16, x1 PCIe x8, x1 PCIe x4
- x2 PCIe x8, x2 PCIe x4, x2 PCIe x2, x2 PCIe x1
- x4 PCIe x4, x4 PCIe x2, x4 PCIe x1

Multi-host capable cards also support Socket-Direct applications and work as regular Single-Host cards, depending on the type of server they are plugged into, assuming the server complies with the OCP 3.0 spec.

According to the OCP 3.0 spec, the SuperNIC advertises its capability through the PRSNTB[3:0]# pins. The server determines the configuration through the BIF[2:0]# pins, which it drives to the SuperNIC.
The NVIDIA OCP3.0 card has an internal logic that uses the BIF[2:0]# data and determines the

correct operating mode to boot at. The combination of the PRSNTB[3:0]# and BIF[2:0]# pins deterministically sets the PCIe lane width for a given combination of OCP 3.0 cards and baseboard. The logic and the decoding table can be found in the OCP 3.0 spec (Chapter 3.5 PCIe Bifurcation Mechanism).

For example:
the NVIDIA OCP 3.0 Multi-host adapter drives 0100 on PRSNTB[3:0]# to the server.

If Server Drivers	Adapter PCIe Mode
000	Single-Host Mode: **x1 PCIe x16**
001	Socket Direct Mode: **x2 PCIe x8**
010	Socket Direct Mode: **x4 PCIe x4**
101	Multi-Host Mode: **x2 PCIe x8**
110	Multi-Host Mode: **x4 PCIe x4**

3.2.2 PCI Express Interface

ConnectX-8 SuperNIC supports PCI Express Gen6 (5.0 and 4.0 compatible) through an x16 edge connector. The following lists host interface features:

- PCIe Gen6 or Gen5 (up to x16 PCIe lanes)
- NVIDIA Multi-Host™ (up to 4 hosts)
- MSI/MSI-X

3.2.3 Networking Interfaces

The SuperNIC includes special circuits to protect the SuperNIC/server from ESD shocks when plugging copper cables.

Ethernet and InfiniBand traffic is transmitted through the networking connectors (OSFP) on the SuperNIC.

Protocol	Specifications
Ethernet	The network ports comply with the IEEE 802.3 Ethernet standards in Features and Benefits.
InfiniBand	The network ports are compliant with the InfiniBand Architecture Specification, Release 1.7.

3.2.4 Networking Ports LEDs Specifications

There are two I/O LEDs per port to indicate port speed and link status.

- LED 1 is a bi-color LED (Yellow and Green)
- LED 2 is a single-color LED (Green)

Single-Port OSFP Internal Lock Bracket

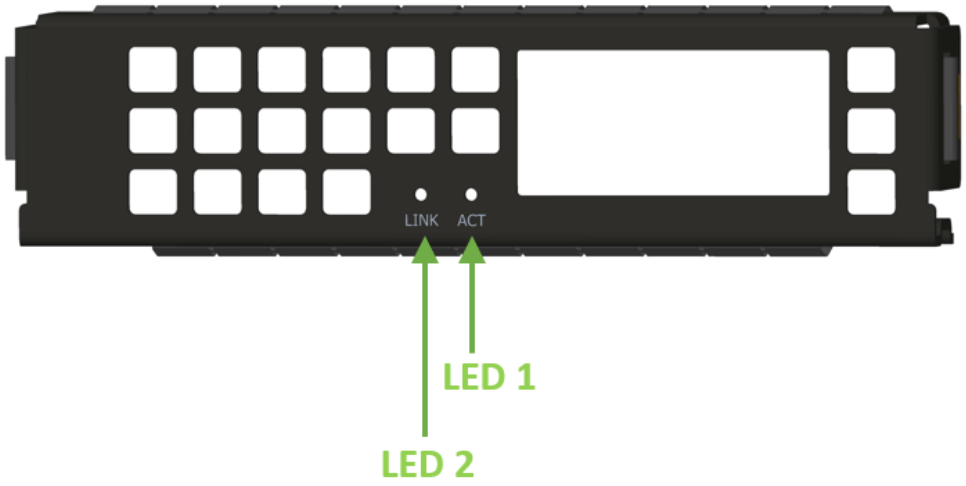

State	Bi-Color LED (Yellow/Green)	Single Color LED (Green)
Beacon command for locating the SuperNIC	1Hz blinking Yellow	OFF
Error	4Hz blinking Yellow Indicates an error with the link. The error can be one of the following: **Error Type** / **Description** / **LED Behavior** I^2C / I^2C access to the networking ports fails / Blinks until error is fixed Over-current / Over-current condition of the networking ports / Blinks until error is fixed	ON
Link Up	In full port speed: the Green LED is solid In less than full port speed: the Yellow LED is solid	ON
Physical Activity	The Green LED will blink.	Blinking

3.2.5 FRU EEPROM

FRU EEPROM allows the baseboard to identify different types of OCP 3.0 cards. FRU EEPROM is accessible through SMCLK and SMDATA. FRU EEPROM address is defined according to SLOT_ID0 and SLOT_ID1 and its capacity is 16KB.

3.2.6 SMBus Interface

ConnectX-8 technology maintains support for manageability through a BMC. ConnectX-8 OCP 3.0 adapter can be connected to a BMC using MCTP over SMBus or MCTP over PCIe protocols as if it is a standard NVIDIA OCP 3.0 adapter. For configuring the adapter for the specific manageability solution in use by the server, please contact NVIDIA Support.

3.2.7 Voltage Regulators

The SuperNIC incorporates a CPLD device that implements the OCP 3.0 host scan chain and controls the networking port logic LED (LED0). It draws its power supply from the 3.3V_EDGE and 12V_EDGE rails.

3.2.8 High-Precision Timing with DC-TCXO

The ConnectX-8 is equipped with a high-stability Digitally Controlled Temperature Compensated Oscillator (DC-TCXO). This oscillator significantly improves the accuracy and stability of timing-sensitive features, including IEEE 1588 Precision Time Protocol (PTP) and Synchronous Ethernet (SyncE).

The DC-TCXO provides the following specifications:

Parameter	Specification
Lifetime Accuracy	±1 ppm
Overall Temperature Stability	±100 ppb
Temperature Slope	±1 ppb/°C

4 Hardware Installation

Installation and initialization of ConnectX-8 SuperNIC require attention to the mechanical attributes, power specification, and precautions for electronic equipment.

4.1 Safety Warnings

Safety warnings are provided here in the English language. For safety warnings in other languages, refer to the SuperNIC Installation Safety Instructions.

Please observe all safety warnings to avoid injury and prevent damage to system components. Note that not all warnings are relevant to all models.

Note that not all warnings are relevant to all models.

	**General Installation Instructions** Read all installation instructions before connecting the equipment to the power source.
	**Jewelry Removal Warning** Before you install or remove equipment that is connected to power lines, remove jewelry such as bracelets, necklaces, rings, watches, and so on. Metal objects heat up when connected to power and ground and can meltdown, causing serious burns and/or welding the metal object to the terminals.
	**Over-temperature** This equipment should not be operated in an area with an ambient temperature exceeding the maximum recommended: 55°C (131°F). An airflow of 200LFM at this maximum ambient temperature is required for HCA cards and NICs. To guarantee proper airflow, allow at least 8cm (3 inches) of clearance around the ventilation openings.
	**During Lightning - Electrical Hazard** During periods of lightning activity, do not work on the equipment or connect or disconnect cables.
	**Copper Cable Connecting/Disconnecting** Some copper cables are heavy and not flexible, as such, they should be carefully attached to or detached from the connectors. Refer to the cable manufacturer for special warnings and instructions.
	**Equipment Installation** This equipment should be installed, replaced, or serviced only by trained and qualified personnel.
	**Equipment Disposal** The disposal of this equipment should be in accordance to all national laws and regulations.
	**Local and National Electrical Codes** This equipment should be installed in compliance with local and national electrical codes.
	**Hazardous Radiation Exposure** • Caution – Use of controls or adjustment or performance of procedures other than those specified herein may result in hazardous radiation exposure.For products with optical ports. • CLASS 1 LASER PRODUCT and reference to the most recent laser standards: IEC 60 825-1:1993 + A1:1997 + A2:2001 and EN 60825-1:1994+A1:1996+ A2:20

4.2 Installation Procedure Overview

The installation procedure of ConnectX-8 SuperNICs involves the following steps:

Step	Procedure	Direct Link
1	Check the system's requirements.	System Requirements
2	Follow the safety precautions	Safety Precautions
3	Unpack the package	Unpack the package
5	Follow the pre-installation checklist	Pre-Installation Checklist
6	(Optional) Replace the full-height mounting bracket with the supplied short bracket	Bracket Replacement Instructions
7	Install the ConnectX-8 SuperNIC in the chassis	ConnectX-8 SuperNIC Installation Instructions
8	Connect cables or modules to the cards	Cables and Modules
9	Identify the SuperNIC in the system	Identifying Your SuperNIC

4.3 System Requirements

To successfully integrate the ConnectX-8 SuperNIC into a chassis, ensure that the following system requirements are met:

Unless otherwise specified, NVIDIA products are designed to work in an environmentally controlled data center with low levels of gaseous and dust (particulate) contamination. The operating environment should meet severity level G1 as per ISA 71.04 for gaseous contamination and ISO 14644-1 class 8 for cleanliness level.

Hardware Requirements	• **Available PCIe Slot**:
	  • Ensure the system has an available **PCIe Gen6 x16** slot for OCP 3.0 for optimal performance.
	  • Ensure the chassis is compatible with the form factor of the ConnectX-8 SuperNIC - Tall Small Form Factor (TSFF) OCP 3.0.
	• **Power Supply**:
	  • Ensure the system's power supply unit (PSU) can provide sufficient power for the ConnectX-8 SuperNIC, typically requiring **75W** from the PCIe slot.
	• **Connectivity**:
	  • The ConnectX-8 SuperNIC supports Ethernet and InfiniBand network connections. Ensure the chassis is connected to the appropriate network infrastructure (switches, routers, etc.).

Environmental Requirements	ConnectX-8 SuperNICs are offered with two airflow patterns: from the heatsink to the network ports, and vice versa, as shown below.
	Please refer to the section for airflow numbers for each specific SuperNIC model.
	The ConnectX-8 SuperNICs are available in two airflow configurations: from the heatsink to the network ports, or from the network ports to the heatsink.
	• **Cooling and Ventilation**:
	  • Adequate airflow is necessary to keep the ConnectX-8 SuperNIC cool. Ensure the chassis has sufficient cooling mechanisms, including fan trays or directed airflow near the PCIe slots.
	  Refer to the Specifications section for the required airflow values.
	  • The card may operate at high thermal loads during peak performance; maintaining optimal ambient temperature is crucial to prevent overheating.
	• **Chassis Size and Slot Availability**:
	  • The chassis must have available PCIe slots of sufficient size (Tall Small Form Factor (TSFF) OCP 3.0)

Software Requirements	• **Operating System**:
	  • Ensure the operating system supports the ConnectX-8 SuperNIC, including DOCA Host for Linux and Windows Server
	  • versions.
	• **Firmware**:
	  • Ensure that the firmware for the ConnectX-8 SuperNIC is up-to-date. The latest firmware versions may include critical bug fixes, performance improvements, and support for new features.
	• **Software**:
	  • Ensure that the necessary network software and libraries are installed, such as:
	  • DOCA-Host for Linux
	  • Windows Server drivers for Windows environments

4.4 Safety Precautions Before Installing ConnectX-8 SuperNIC

The SuperNIC is being installed in a system that operates with voltages that can be lethal. Before opening the case of the system, observe the following precautions to avoid injury and prevent damage to system components.

4.4.1 Power Down and Disconnect

- Power Off the System: Ensure the system is powered down completely before beginning the installation process. This includes shutting down the operating system and turning off the power supply.
- Unplug the Power Cable: Physically disconnect the power cable from the system to prevent any electrical hazards.

4.4.2 Electrostatic Discharge (ESD) Prevention

- Use an Anti-Static Wrist Strap: Always wear an anti-static wrist strap to prevent electrostatic discharge (ESD) that could damage sensitive components. Attach the wrist strap to a grounded surface.

- Work on an Anti-Static Mat: If possible, use an anti-static mat to place the SuperNIC and other components on while working. This helps to further reduce the risk of static buildup.

4.4.3 Handle Components with Care

- Avoid Touching the Connectors: Handle the ConnectX-8 SuperNIC by its edges and avoid touching the connectors or PCB surface to prevent damage from oils, dirt, or static electricity.
- Use Proper Tools: Use appropriate tools, such as a screwdrivers, to remove any screws or fasteners. Do not force components into place.

4.4.4 Proper Grounding

- Ensure Proper Grounding: Make sure that the system and all components are properly grounded. This helps to minimize the risk of electrical shock or component damage.
- Verify Grounding of Work Area: Ensure the area where you are installing the SuperNIC is grounded and free from electrical hazards.

4.4.5 Environmental Safety

- Work in a Well-Ventilated Area: Ensure that the installation area is well-ventilated and free from excessive dust or humidity.
- Ensure Clean Hands: Before handling the ConnectX-8 SuperNIC, make sure your hands are clean and free from dirt or oils that could damage the components.

4.5 Bracket Replacement Instructions

The OCP 3.0 adapter card is shipped assembled either with a thumbscrew (pull-tab) bracket, an internal-lock bracket, or an ejector-latch bracket. If this form factor is suitable for your requirements, you can skip the remainder of this section and move to Installation Instructions. If you need to replace the assembled OCP 3.0 bracket with a different form-factor bracket, please follow the instructions in this section.

Due to risk of damaging the EMI gasket, it is not recommended to replace the bracket more than three times.

To replace the bracket you will need the following parts:

- The new bracket of the desired form factor.
- The screws supplied with the new bracket kit.
- The required torx tool type as specified in the instructions.

Removing the Existing Bracket

1. Using the torx tool type listed in the below table, remove the screws according to the instructions per OCP 3.0 bracket type.

Internal Lock Bracket	Pull-tab (Thumbscrew) Bracket	Ejector-Latch Bracket
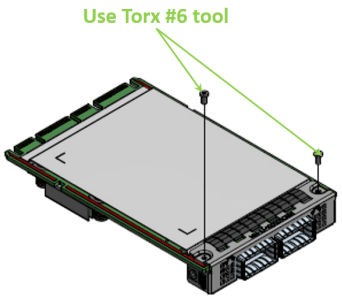	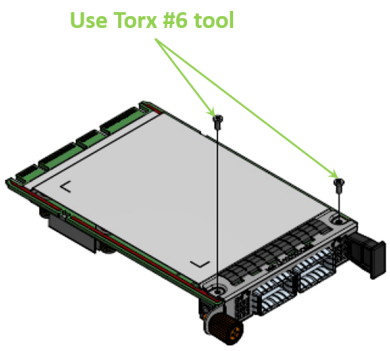	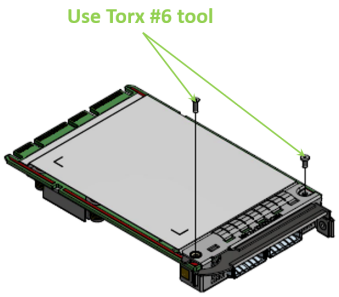

2. Gently separate the bracket from the OCP 3.0 card.

Internal Lock Bracket	Pull-tab Bracket	Ejector-Latch Bracket
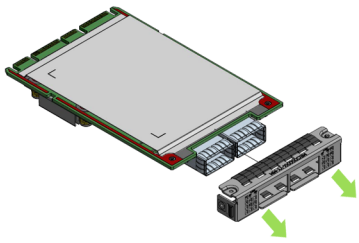	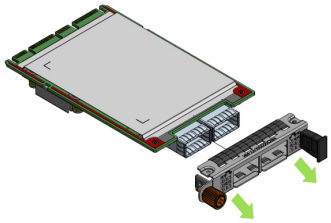	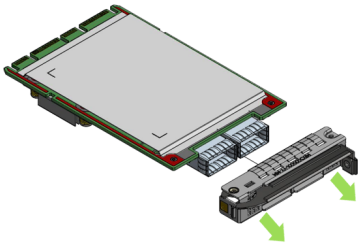

Be careful not to put stress on the LEDs on the adapter card.

3. Save the two screws.

Installing the New Bracket

1. Assemble the new bracket onto the card.

Internal Lock Bracket	Pull-tab Bracket	Ejector-Latch Bracket
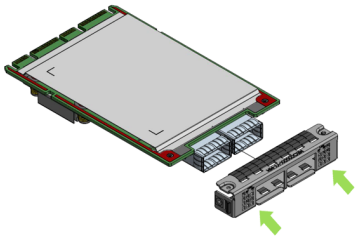	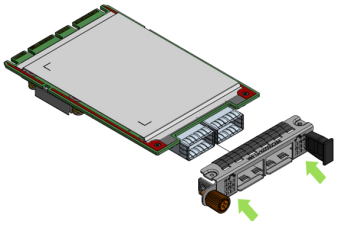	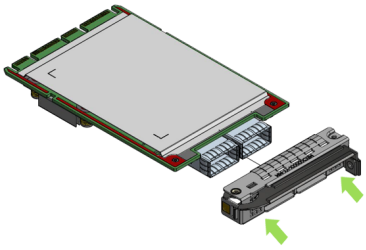

Do not force the bracket onto the adapter card.

2. Ensure that the insulator's front edge is beneath the bracket, as shown in the below figure.

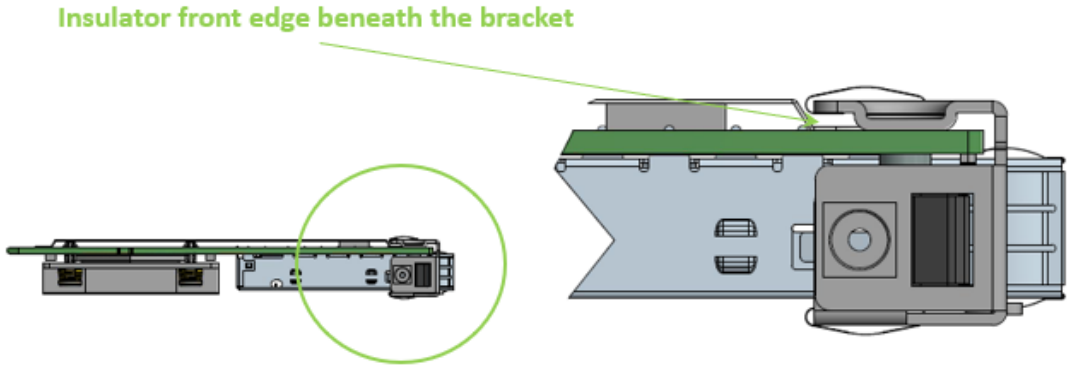

3. Screw on the OCP 3.0 bracket with the supplied screws that came with the new bracket kit. Use the specified torx tool type and apply the specified torque on the screws per bracket form factor.

Internal Lock Bracket	Pull-tab Bracket	Ejector-Latch Bracket
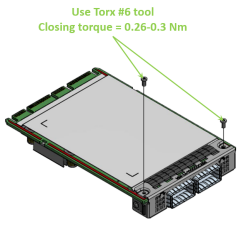	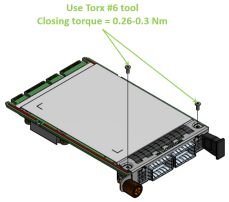	Note that one screw is flat-head 90° screw and the other is socket-head screw. Screw on the different types of screws as shown in the below illustration. **Use Torx #6 tool** **Closing torque = 0.14-0.18 Nm** 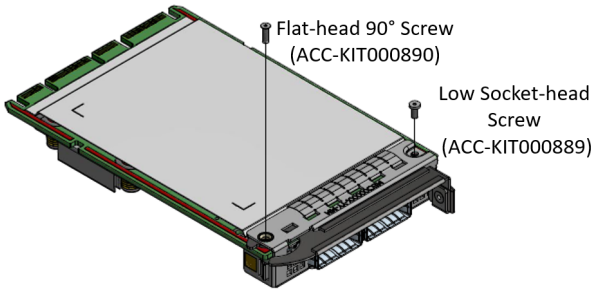

4.6 Preparation

- Power down the chassis: Ensure the chassis is completely powered off. Disconnect it from any power source and remove all cables.
- Discharge static electricity: Put on an anti-static wrist strap or regularly touch a grounded metal surface to discharge any static electricity.
- Open the chassis:
  - Using the appropriate screwdriver, remove the screws securing the chassis side panel or top cover (depending on the chassis model). Set the screws aside in a safe place.
  - Locate an available PCIe slot that is compatible with the ConnectX-8 SuperNIC. Refer to the chassis and motherboard manual to ensure the PCIe slot supports the NIC's specifications (PCIe Gen6 x16).

4.7 Install the ConnectX-8 SuperNIC

- Position the SuperNIC: Carefully align the gold connector of the NIC with the PCIe slot. Ensure the card is lined up with the slot and the retention bracket aligns with the chassis screw holes.
- Insert the SuperNIC: Gently but firmly push the NIC straight down into the PCIe slot until it is fully seated. Avoid excessive force.
- Secure the SuperNIC: Use a screw (often provided with the chassis or card) to secure the card's metal bracket to the chassis. This prevents the card from moving and ensures a stable connection.
- Once the SuperNIC is securely installed, reattach the chassis side panel or top cover. Fasten the screws that were removed earlier.

4.8 Connect Power and Cables

- Connect the power cable to the chassis.
- Install the network cables to the newly installed SuperNIC. For more information, refer to Networking Cable Installation.
- Ensure the connections are secure before powering on the chassis.

4.9 Power On and Verify the Installation

- Power on the chassis.
- Verify that the ConnectX-8 SuperNIC is properly recognized by the system. Depending on the operating system:
  - On Linux, run `lspci | grep Mellanox` to ensure the NIC is detected.
  - On Windows, check the Device Manager under "Network Adapters" for the presence of the new NIC.

4.10 Uninstallation Instructions

4.10.1 Preparation

- Power down the chassis:
  - Shut down the system properly, then disconnect the power supply and all cables from the chassis to ensure it's completely powered off.
  - Wait 30 seconds.
- Discharge static electricity:
  - Use an anti-static wrist strap or touch a grounded metal object regularly to discharge any static electricity that could damage the components.

4.10.2 Open the Chassis

- Using a screwdriver, unscrew and remove the chassis side panel (or top cover, depending on the chassis design). Keep the screws in a safe place for reassembly later.

- Identify the ConnectX-8 SuperNIC you wish to remove. It will be seated in one of the PCIe slots.

4.10.3 Disconnect the Networking Cable

To remove the networking cable, disengage the locks and slowly pull the connector away from the port receptacle. The LED indicator will turn off when the cable is unseated. Note which cables go where if they will need to be reconnected later.

4.10.4 Uninstall the ConnectX-8 SuperNIC

- Locate the screw that secures the SuperNIC's metal bracket to the chassis. Use a screwdriver to remove this screw, setting it aside safely.
- Holding the edges of the SuperNIC, gently pull it upwards out of the PCIe slot. Apply even pressure to avoid damaging the slot or the card. Do not twist or force it out.

4.10.5 (Optional) Inspect and Clean the Components

- If needed, inspect the PCIe slot and the NIC for dust or debris. Use compressed air to clean them, but avoid touching the connectors directly.

4.10.6 Close the Chassis

- Once the SuperNIC is removed and any further maintenance is complete, replace the chassis side panel and secure it with the screws you removed earlier.

4.10.7 Reconnect Power and Cables

- Reconnect the power cable and any necessary peripherals to the system.
- Ensure all connections are secure before powering the system back on.

4.10.8 Power On and Verify

- Power on the system and verify that it operates correctly without the NIC.
- Depending on your operating system, you may need to check if the drivers for the NIC have been disabled or removed. On Linux, you can use `lspci` to verify the card is no longer detected, and on Windows, you can check the Device Manager under "Network Adapters."

4.10.8.1 Cables and Modules

Networking Cable Installation and Uninstallation

Before connecting a cable to the SuperNIC, ensure that the bracket is fastened to the server chassis using a screw to prevent movement or unplugging of the SuperNIC when the cable is inserted or extracted.

1. All cables can be inserted or removed with the unit powered on.

2. To insert a cable, press the connector into the port receptacle until the connector is firmly seated.
    a. Support the weight of the cable before connecting the cable to the SuperNIC. Do this by using a cable holder or tying the cable to the rack.
    b. Determine the correct orientation of the connector to the SuperNIC before inserting the connector. Do not try and insert the connector upside down. This may damage the SuperNIC.
    c. Insert the connector into the SuperNIC. Be careful to insert the connector straight into the cage. Do not apply any torque, up or down, to the connector cage in the SuperNIC.
    d. Make sure that the connector locks in place.

    > ⚠ When installing cables make sure that the latches engage.

    > ⛔ Always install and remove cables by pushing or pulling the cable and connector in a straight line with the SuperNIC.

3. After inserting a cable into a port, the Green LED indicator will light when the physical connection is established (that is, when the unit is powered on and a cable is plugged into the port with the other end of the connector plugged into a functioning port). See LED Interface under the Interfaces section.
4. After plugging in a cable, lock the connector using the latching mechanism particular to the cable vendor. When data is being transferred the Green LED will blink. See LED Interface under the Interfaces section.
5. Care should be taken so as not to impede the air exhaust flow through the ventilation holes. Use cable lengths that allow for routing horizontally around to the side of the chassis before bending upward or downward in the rack.
6. To remove a cable, disengage the locks and slowly pull the connector away from the port receptacle. The LED indicator will turn off when the cable is unseated.

4.10.9 Identifying the SuperNIC in Your System

On Linux

Get the device location on the PCI bus by running `lspci` and locating lines with the string “Mellanox Technologies”:

```
84:00.0 Ethernet controller: Mellanox Technologies CX8 Family [ConnectX-8]
```

On Windows

1. Open Device Manager on the server. Click Start => Run, and then enter devmgmt.msc.
2. Expand System Devices and locate your ConnectX-8 SuperNIC.
3. Right-click the mouse on your SuperNIC's row and select Properties to display the SuperNIC properties window.
4. Click the Details tab and select Hardware IDs (Windows 2022/2019/2016/2012 R2) from the Property pull-down menu.

PCI Device (Example)

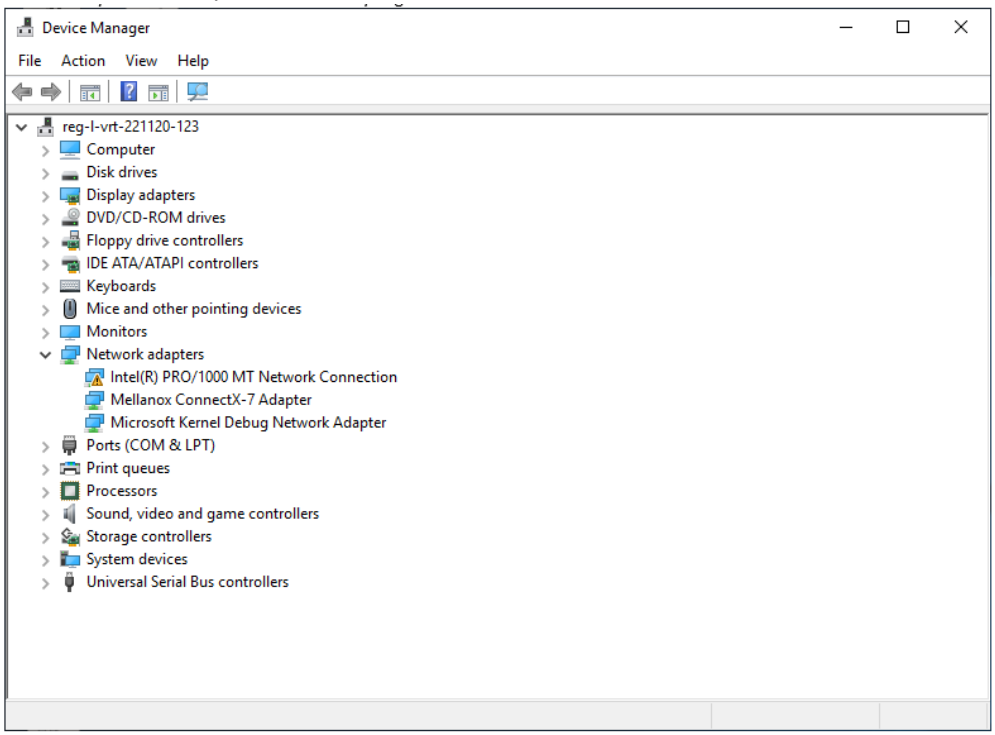

5. In the **Value** display box, check the fields VEN and DEV (fields are separated by '&'). In the display example above, notice the sub-string "PCI\VEN_15B3&DEV_1021": VEN is equal to 0x15B3 – this is the Vendor ID of Mellanox Technologies, and DEV is equal to 1023 (for ConnectX-8) – this is a valid NVIDIA PCI Device ID.

   ⚠ If the PCI device does not have a SuperNIC ID, return to Step 2 to check another device.

   ⚠ The list of NVIDIA PCI Device IDs can be found at the PCI ID repository.

5 General Cooling Recommendations for Partner-Cooled SuperNICs

IMPORTANT NOTES:

1. This section applies to 900-9X85E-00EX-MP0 SuperNIC.
2. The customer holds exclusive responsibility for the thermal design and for ensuring all board components never exceed their designated thermal operating limits.
3. The recommendations provided in this chapter are based on an air-cooled and liquid-cooled thermal design. For further information about NVIDIA's air-cooled thermal design, please refer to your NVIDIA representative.
4. If the customer's environment has low or no airflow (e.g., liquid cooling) where the SuperNIC is to be installed, the customer must ensure proper cooling for all the other board components that are not covered by the recommendations of this chapter or in the product's thermal model. The customer must guarantee those components do not exceed their thermal operating limits as well.
5. The partner-cooled thermal solution dimensions exceed the OCP 3.0 specifications. The customer must ensure that the liquid solution is assembled after the card is inserted in the server.

5.1 General Notes

- Recommended torque values for the bracket screws: M2: 0.14-0.18 Nm. Use tool with Torx #5.
- Recommended torque values for the customer cooling solution interface: M2.5: 0.44~0.46 Nm. Refer to the cooling mechanism documentation for the required torque.

5.2 Components that Require Cooling

It is the responsibility of the customer to ensure proper cooling and thermal regulation for all components on the board. The provided guidelines are minimal and are based on NVIDIA's air-cooling thermal solution.

There are two components that require cooling, the built-in TTP and the networking port cage.

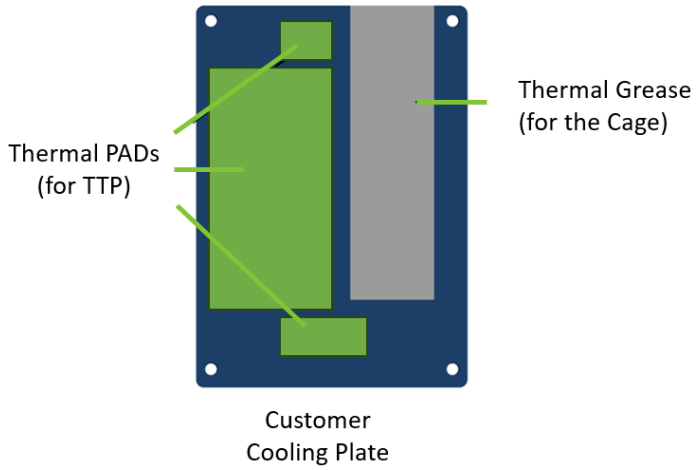

TTP

- The SuperNIC IC is in contact with the built-in TTP on the component side of the product.
- The maximum allowed force applied on the TTP should adhere to the following specifications:
  - 268 Newtons on the TTP (max thermal pads area:2077mm^2)
  - Force should be spread equally on the TTP top surface.
  - Thermal Pad minimum allowable conductivity is 2.5 W/mK

Networking Cage

- The thermal grease the comes in contact with the networking cage thermal bridge. Recommended thermal grease for the cage's thermal bridge is DOWSIL TC-5622 0.1+-0.05 mm thickness.

5.3 Suggested Server Connectivity Interfaces

The four screws should be assembled from the customer cooling solution, through the four standoffs on the ConnectX-8 SuperNIC.

The four standoffs type is M2.5x0.45 thread and the maximum thread depth 4.0 mm.

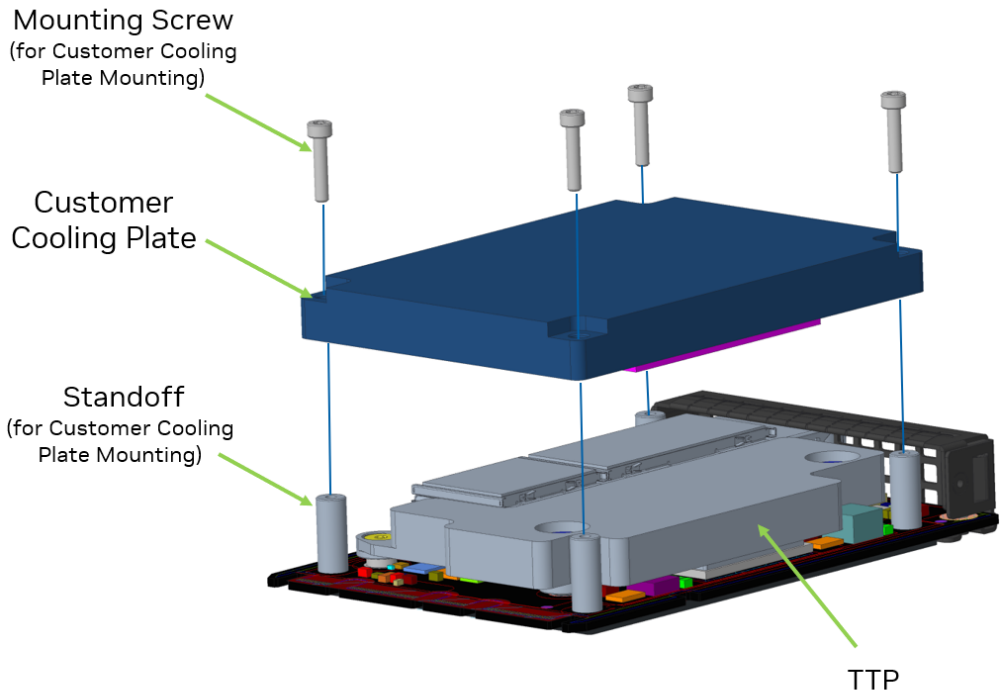

5.4 Threads and Holes

To ensure proper SuperNIC connectivity to the server and TTP, the below figure displays the SuperNIC threads and holes and provides its dimensions.

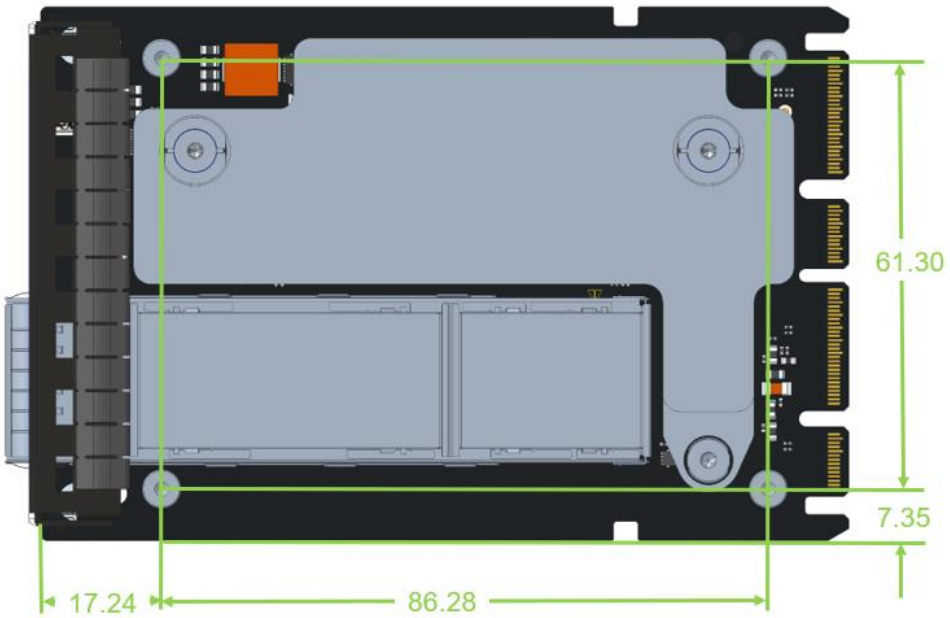
61.30
7.35
17.24
86.28

6 Driver Installation

Please refer to the relevant driver installation section.

- DOCA-Host for Linux Driver Installation
- Windows Driver Installation

6.1 DOCA-Host for Linux Driver Installation

This section describes how to install and test the DOCA-Host for Linux package on a single server with a ConnectX-8 SuperNIC installed.

6.1.1 Prerequisites

Requirements	Description
Platforms	A server platform with a ConnectX-8 SuperNIC installed.
Required Disk Space for Installation	1GB
Operating System	Linux operating system. For the list of supported operating system distributions and kernels, please refer to the DOCA-Host Release Notes.
Installer Privileges	The installation requires administrator (root) privileges on the target machine.

6.1.2 Downloading DOCA-Host

Verify that the system has a network adapter installed by running **lspci** command. The below table provides output examples per ConnectX-8 card configuration.

ConnectX-8 Card Configuration	Output Examples
Single-port Socket Direct Card/ MultiHost (2x PCIe x16) OPN: 900-9X81Q-00CN-ST0	`11:00.0 Infiniband controller: Mellanox Technologies CX8 Family [ConnectX-8]` `59:00.0 Infiniband controller: Mellanox Technologies CX8 Family [ConnectX-8]`

ConnectX-8 Card Configuration	Output Examples
Dual-port Socket Direct Card/ MultiHost (2x PCIe x16) OPN: 900-9X81Q-00CN-ST0	82:00.0 Ethernet controller: Mellanox Technologies CX8 Family [ConnectX-8] 82:00.1 Ethernet controller: Mellanox Technologies CX8 Family [ConnectX-8] 84:00.0 Ethernet controller: Mellanox Technologies CX8 Family [ConnectX-8] 84:00.1 Ethernet controller: Mellanox Technologies CX8 Family [ConnectX-8] In the output example above, the first two rows indicate that one card is installed in a PCI slot with PCI Bus address 05 (hexadecimal), PCI Device number 00, and PCI Function numbers 0 and 1. The other card is installed in a PCI slot with PCI Bus address 82 (hexadecimal), PCI Device number 00, and PCI Function numbers 0 and 1. Since the two PCIe cards are installed in two PCIe slots, each card gets a unique PCI Bus and Device number. Each of the PCIe x16 busses sees two network ports; in effect, the two physical ports of the ConnectX-7 Socket Direct adapter are viewed as four net devices by the system.
Single-port with x16 PCIe Down Stream Port Extension Option OPN: 900-9X81E-00EX-DT0	84:00.0 Ethernet controller: Mellanox Technologies CX8 Family [ConnectX-8]
Dual-port PCIe x16 Card OPN: 900-9X81Q-00CN-ST0	82:00.0 Ethernet controller: Mellanox Technologies CX8 Family [ConnectX-8] 82:00.1 Ethernet controller: Mellanox Technologies CX8 Family [ConnectX-8]
Single-port PCIe x16 Card OPN: 900-9X81E-00EX-ST0	84:00.0 Ethernet controller: Mellanox Technologies CX8 Family [ConnectX-8]

For instructions on downloading DOCA-Host, refer to NVIDIA DOCA Downloads.

6.1.3 Installing DOCA-Host

For Linux driver installation, please refer to NVIDIA DOCA Installation Guide for Linux.

6.2 Windows Driver Installation

For Windows, download and install the latest WinOF-2 for Windows software package available via the NVIDIA website at: WinOF-2 webpage. Follow the installation instructions included in the download package (also available from the download page).

The snapshots in the following sections are presented for illustration purposes only. The installation interface may slightly vary, depending on the operating system in use.

6.2.1 Software Requirements

For the full list of supported operating systems, please refer to the WinOF-2 Documentation -> Release Notes.

6.2.2 Downloading WinOF-2 Driver

To download the .exe file according to your Operating System, please follow the steps below:

1. Obtain the machine architecture.
   a. To go to the Start menu, position your mouse in the bottom-right corner of the Remote Desktop of your screen.
   b. Open a CMD console (Click Task Manager-->File --> Run new task and enter CMD).
   c. Enter the following command.

```
echo %PROCESSOR_ARCHITECTURE%
```

   On an x64 (64-bit) machine, the output will be “AMD64”.

2. Go to the WinOF-2 web page at: https://www.nvidia.com/en-us/networking/ > Products > Software > InfiniBand Drivers (Learn More) > Nvidia WinOF-2.
3. Download the .exe image according to the architecture of your machine (see Step 1).
   The name of the .exe is in the following format: MLNX_WinOF2-<version>_<arch>.exe.

   Installing the incorrect .exe file is prohibited. If you do so, an error message will be displayed.
   For example, if you install a 64-bit .exe on a 32-bit machine, the wizard will display the following (or a similar) error message: “The installation package is not supported by this processor type. Contact your vendor”

6.2.3 Installing WinOF-2 Driver

The snapshots in the following sections are for illustration purposes only. The installation interface may slightly vary, depending on the used operating system.

This section provides instructions for two types of installation procedures, and both require administrator privileges:

- Attended Installation
  An installation procedure that requires frequent user intervention.
- Unattended Installation
  An automated installation procedure that requires no user intervention.

6.2.3.1 Attended Installation

The following is an example of an installation session.

1. Double click the .exe and follow the GUI instructions to install MLNX_WinOF2.
2. [Optional] Manually configure your setup to contain the logs option (replace “LogFile” with the relevant directory).

```
MLNX_WinOF2_<revision_version>_All_Arch.exe /v"/l*vx [LogFile]"
```

3. [Optional] If you do not want to upgrade your firmware version (i.e., MT_SKIPFWUPGRD default value is False).

```
MLNX_WinOF2_<revision_version>_All_Arch.exe /v" MT_SKIPFWUPGRD=1"
```

4. [Optional] If you do not want to install the Rshim driver, run.

```
MLNX_WinOF2_<revision_version>_All_Arch.exe /v" MT_DISABLE_RSHIM_INSTALL=1"
```

> The Rshim driver installanion will fail if a prior Rshim driver is already installed. The following fail message will be displayed in the log:
>
> "ERROR!!! Installation failed due to following errors: MlxRshim drivers installation disabled and MlxRshim drivers Installed, Please remove the following oem inf files from driver store: <oem inf list>"

5. [Optional] If you want to skip the check for unsupported devices, run.

```
MLNX_WinOF2_<revision_version>_All_Arch.exe /v" SKIPUNSUPPORTEDDEVCHECK=1"
```

6. Click Next in the Welcome screen.

7. Read and accept the license agreement and click Next.

8. Select the target folder for the installation.

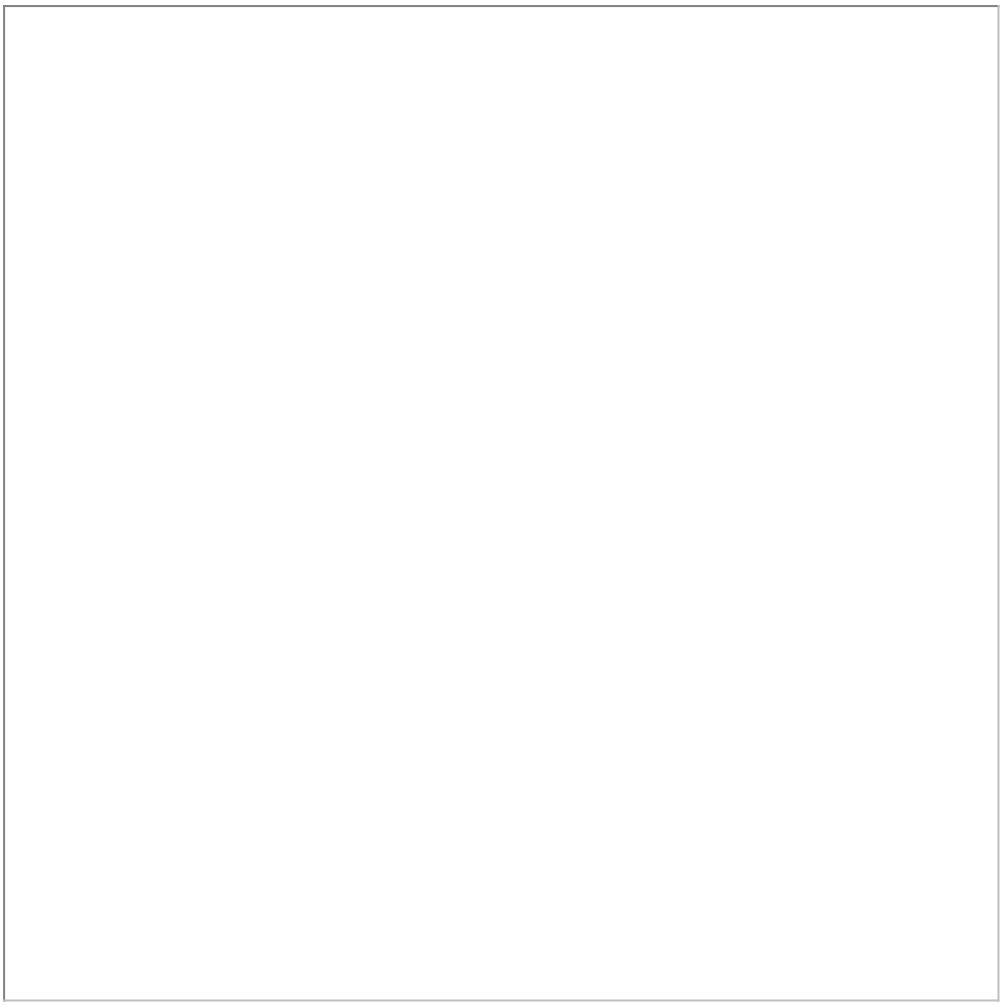

9. The firmware upgrade screen will be displayed in the following cases:
    - If the user has an OEM card. In this case, the firmware will not be displayed.
    - If the user has a standard NVIDIA® card with an older firmware version, the firmware will be updated accordingly. However, if the user has both an OEM card and a NVIDIA® card, only the NVIDIA® card will be updated.

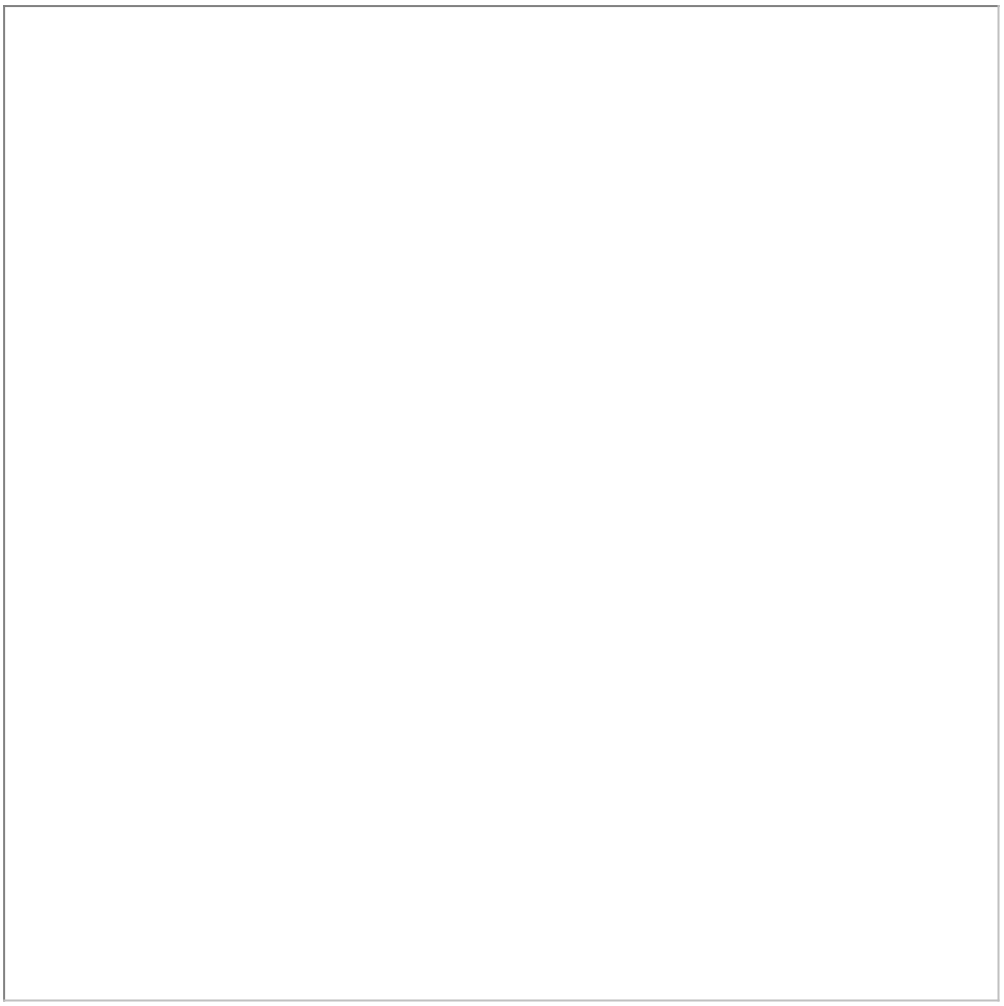

10. Select a Complete or Custom installation, follow Step a onward.

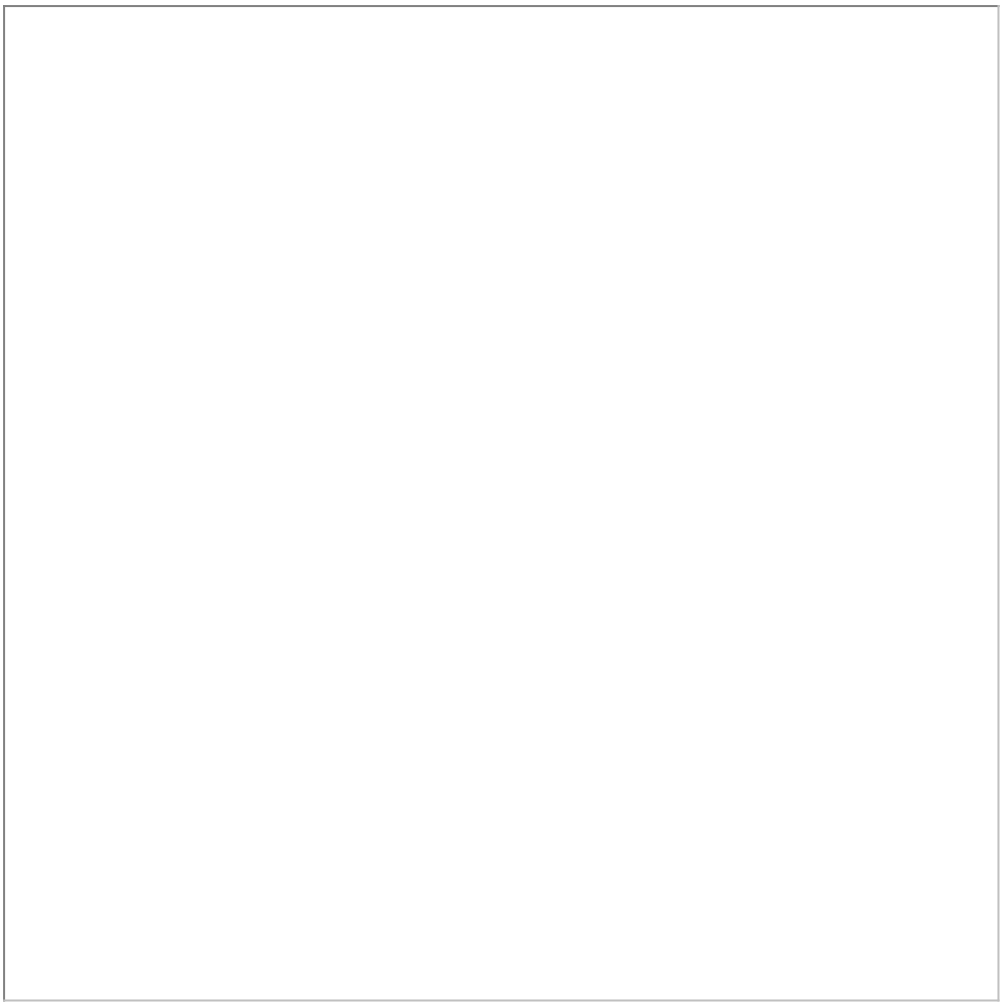

a. Select the desired feature to install:

- Performances tools - install the performance tools that are used to measure performance in user environment
- Documentation - contains the User Manual and Release Notes
- Management tools - installation tools used for management, such as mlxstat
- Diagnostic Tools - installation tools used for diagnostics, such as mlx5cmd

b. Click Next to install the desired tools.

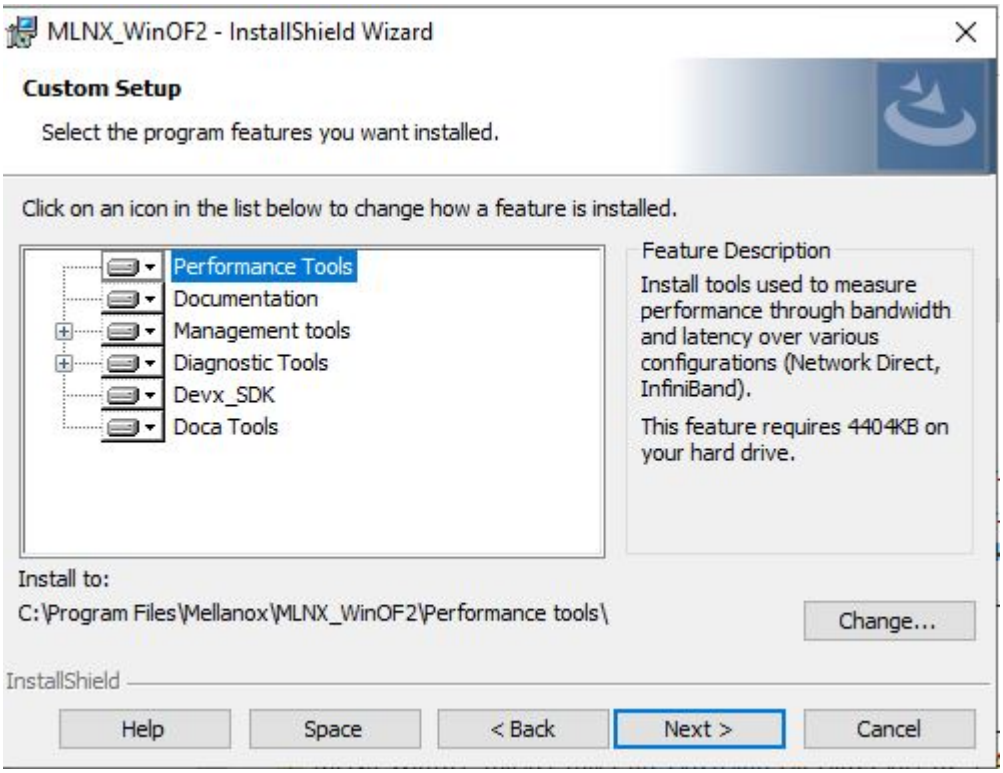

11. Click Install to start the installation.

12. In case firmware upgrade option was checked in Step 7, you will be notified if a firmware upgrade is required (see).

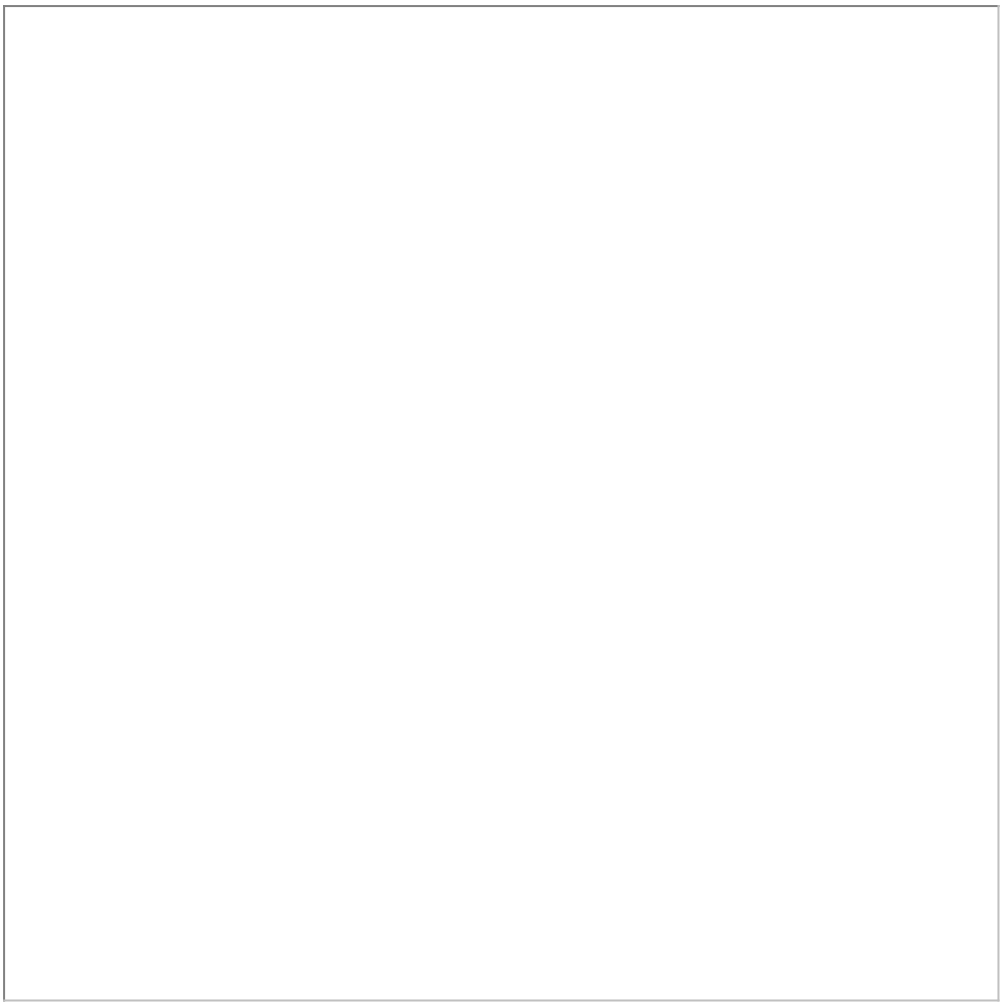

13. Click Finish to complete the installation.

6.2.3.2 Unattended Installation

> If no reboot options are specified, the installer restarts the computer whenever necessary without displaying any prompt or warning to the user.
> To control the reboots, use the */norestart* or */forcerestart* standard command-line options.

The following is an example of an unattended installation session.

1. Open a CMD console-> Click Start-> Task Manager File-> Run new task-> and enter CMD.
2. Install the driver. Run:

```
MLNX_WinOF2-[Driver/Version]_<revision_version>_All_-Arch.exe /S /v/qn
```

3. [Optional] Manually configure your setup to contain the logs option:

```
MLNX_WinOF2-[Driver/Version]_<revision_version>_All_-Arch.exe /S /v/qn /v"/l*vx [LogFile]"
```

4. [Optional] if you wish to control whether to install ND provider or not (i.e., *MT_NDPROPERTY default value is True*).

```
MLNX_WinOF2-[Driver/Version]_<revision_version>_All_Arch.exe /vMT_NDPROPERTY=1
```

5. [Optional] If you do not wish to upgrade your firmware version (i.e.,*MT_SKIPFWUPGRD default value is False*).

```
MLNX_WinOF2-[Driver/Version]_<revision_version>_All_Arch.exe /vMT_SKIPFWUPGRD=1
```

6. [Optional] If you do not want to install the Rshim driver, run.

```
MLNX_WinOF2_<revision_version>_All_Arch.exe /v" MT_DISABLE_RSHIM_INSTALL=1"
```

> The Rshim driver installanion will fail if a prior Rshim driver is already installed. The following fail message will be displayed in the log:
>
> "ERROR!!! Installation failed due to following errors: MlxRshim drivers installation disabled and MlxRshim drivers Installed, Please remove the following oem inf files from driver store: <oem inf list>"

7. [Optional] If you want to enable the default configuration for Rivermax, run.

```
MLNX_WinOF2_<revision_version>_All_Arch.exe /v"MT_RIVERMAX=1 /l*vx C:\Users\<user>\log.txt "
```

8. [Optional] If you want to skip the check for unsupported devices, run/

```
MLNX_WinOF2_<revision_version>_All_Arch.exe /v" SKIPUNSUPPORTEDDEVCHECK=1"
```

6.2.4 Firmware Upgrade

If the machine has a standard NVIDIA® card with an older firmware version, the firmware will be automatically updated as part of the NVIDIA® WinOF-2 package installation. For information on how to upgrade firmware manually, please refer to MFT User Manual.

If the machine has a DDA (pass through) facility, firmware update is supported only in the Host. Therefore, to update the firmware, the following must be performed:

1. Return the network adapters to the Host.
2. Update the firmware according to the steps in the MFT User Manual.
3. Attach the adapters back to VM with the DDA tools.

7 Firmware Upgrade

Each SuperNIC is shipped with the latest version of qualified firmware at the time of manufacturing. However, NVIDIA issues firmware updates occasionally that provide new features and bug fixes. To check that your card is programmed with the latest available firmware version, download the mlxup firmware update and query utility. The utility can query for available SuperNICS and indicate which SuperNIC require a firmware update. If the user confirms, mlxup upgrades the firmware using embedded images. The latest mlxup executable and documentation are available in mlxup - Update and Query Utility.

Firmware Update Example

```
[server1]# ./mlxup
Querying Mellanox devices firmware ...
Device Type:      ConnectX-8
Part Number:      900-9X85E-00EX-MC0
Description:      NVIDIA ConnectX-8 C8180P, OCP3.0 TSFF SuperNIC, 800Gbs XDR IB (default mode) / 2x400GbE, Single-
port OSFP, Multi-Host and Socket Direct capable, PCIe 6.0 x16, Crypto Enabled, Secure Boot Enabled, Thumbscrew
(Pull Tab) Bracket
PCI Device Name: 0b:00.0
Base MAC:         0000e41d2d5cf810
Versions:         Current        Available
  FW              28.33.0800     28.33.1000
  Status:           Update required


Device Type:      ConnectX-8
Part Number:      900-9X85E-00EX-MC0
Description:      NVIDIA ConnectX-8 C8180P, OCP3.0 TSFF SuperNIC, 800Gbs XDR IB (default mode) / 2x400GbE, Single-
port OSFP, Multi-Host and Socket Direct capable, PCIe 6.0 x16, Crypto Enabled, Secure Boot Enabled, Thumbscrew
(Pull Tab) Bracket
Base MAC:         0000e41d2d5cf810
Versions:         Current        Available
  FW              28.33.0800     28.33.1000
  Status:           Up to date

Perform FW update? [y/N]: y
Device #1: Up to date
Device #2: Updating FW ... Done

Restart needed for updates to take effect.
Log File: /var/log/mlxup/mlxup-yyyymmdd.log
```

8 SuperNIC Configurations

8.1 Setting High-Speed-Port Link Type

The following table lists the supported speeds, and the default networking port link type.

OPN	Data Transmission Rate	Default Protocol and Rate
900-9X85E-00EX-MJ0 900-9X85E-00EX-MP0	800Gbs XDR IB / 2x400GbE	800Gbs XDR IB

You can use the mlxconfig or the UEFI tools to configure the networking high-speed ports mode.

Use the UEFI tool to configure the SuperNIC before the operating system is up. User the mlxconfig tool to configure the SuperNIC once the operating system is up. According to your preference, use one of the below tools:

8.1.1 mlxconfig

The mlxconfig tool allows users to change device configurations without burning the firmware. The configuration is also kept after reset. By default, mlxconfig shows the configurations that will be loaded in the next boot. For more information and instructions, refer to Using mlxconfig to Set IB/ETH Parameters.

8.1.2 UEFI

PreBoot drivers initialize the SuperNIC, check the port protocol type – Ethernet or InfiniBand - and bring up the port. Then, it connects to a DHCP server to obtain its assigned IP address and network parameters and the source location of the kernel/OS from which to boot. The DHCP server instructs the PreBoot drivers to access the kernel/OS through a TFTP server, an iSCSI target, or some other service. For more information and instructions, refer to UEFI.

8.2 Port Configurations

ConnectX-8 SuperNICs offer a variety of network port configurations designed to meet the demands of different environments and deployments. This section outlines the available connectivity options for network ports per product model. While the configurations listed here are not exhaustive, they represent the tested and verified options. Customers seeking to support unlisted configurations should contact NVIDIA Support.

The Port Splitting feature allows a single physical networking module (OSFP) to be split into multiple network ports. This provides flexibility in optimizing port configurations for various network topology use cases.

Key Features

- Flexibility and Improved Traffic Segmentation: Port splitting enables network administrators to align device ports with their network infrastructure, enabling traffic segmentation and isolation.

- Increased Network Efficiency: By optimizing traffic flow and reducing congestion, port splitting improves overall network efficiency. Specific physical ports can be assigned to handle certain traffic types or applications, allowing control and performance tuning.
- Better Utilization of Resources: Instead of needing separate SuperNICs for each logical connection or network interface, port splitting allows a single physical SuperNIC to handle multiple functions. This reduces hardware costs while maintaining high network flexibility.

Use Cases

- Load Balancing: Distributing traffic across multiple physical ports can improve network load distribution and balance.
- Resiliency: Multiple physical ports can be used for redundancy, ensuring stable network connections even if one port encounters an issue.

Each configuration in this section includes a diagram and instructions on splitting the network port using the mlxconfig or the UEFI tools. The UEFI tool configures the device before the operating system is up, while mlxconfig configures the device once the operating system is up. Select the tool based on your preference; the configuration is saved across reboots.

mlxconfig: The mlxconfig tool allows users to change configurations without burning the firmware. The configuration is also kept after reset. By default, mlxconfig shows the configurations that will be loaded in the next boot. For more information and instructions, refer to mlxconfig.

UEFI: PreBoot drivers initialize the adapter device, check the port protocol type – Ethernet or InfiniBand - and bring up the port. Then it connects to a DHCP server to obtain its assigned IP address and network parameters and obtain the source location of the kernel/OS to boot from. The DHCP server instructs the PreBoot drivers to access the kernel/OS through a TFTP server, an iSCSI target, or some other service. For more information and instructions, refer to UEFI.

Before setting the desired configuration, take note of the following warnings:

- Any new configuration set by the customer overwrites the previous configuration.
- Misconfiguration can potentially damage the system.
- It is recommended to establish out-of-band connectivity to the ConnectX-8 SuperNIC before setting any configuration for the first time. This allows resetting the NVConfig parameters to their default values in case of misconfiguration.
- When switching from one non-default configuration to another, it is advised to first reset the device to its default configuration. This action deletes all current NVConfig settings. For the reset command and further details, refer to the mlxconfig – Changing Device Configuration Tool.

The following table summarizes the available port-splitting configurations.

Model	OPNs	Configuration	Configuration Description
C8180P	900-9X85E-00EX-MP0 900-9X85E-00EX-MJ0	Configuration #1 (Default)	One InfiniBand port supporting XDR 800Gb/s
		Configuration #2	Two Ethernet ports, each supporting 400GbE
		Configuration #3	Eight Ethernet ports, each supporting 100GbE

8.2.1 C8180P Port-Splitting Configurations

The C8180 ConnectX-8 SuperNIC, featuring a single-port OSFP module, supports up to 800Gb/s. For the reconfiguration to take effect, please perform power cycle.

8.2.1.1 C8180P Configuration #1 (Default): One InfiniBand Port, Supporting XDR 800Gb/s

As the default configuration, the device is configured with a single InfiniBand port, providing an InfiniBand XDR rate of 800Gb/s that operates as four times 200G/s. The configuration is illustrated in the diagram below.

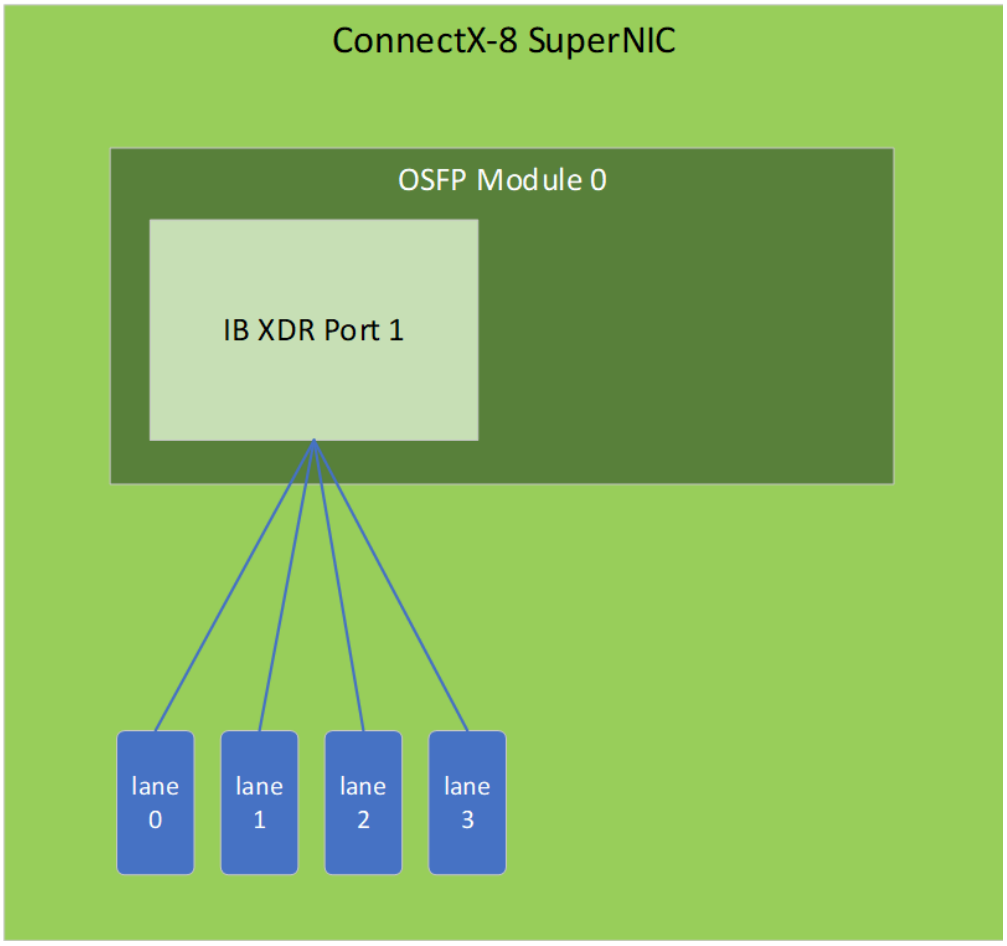

8.2.1.2 C8180P Configuration #2: Two Ethernet Ports, Each Supporting 400GbE

In this configuration, the device can be reconfigured to support two 400GbE Ethernet ports. The configuration is illustrated in the diagram below.

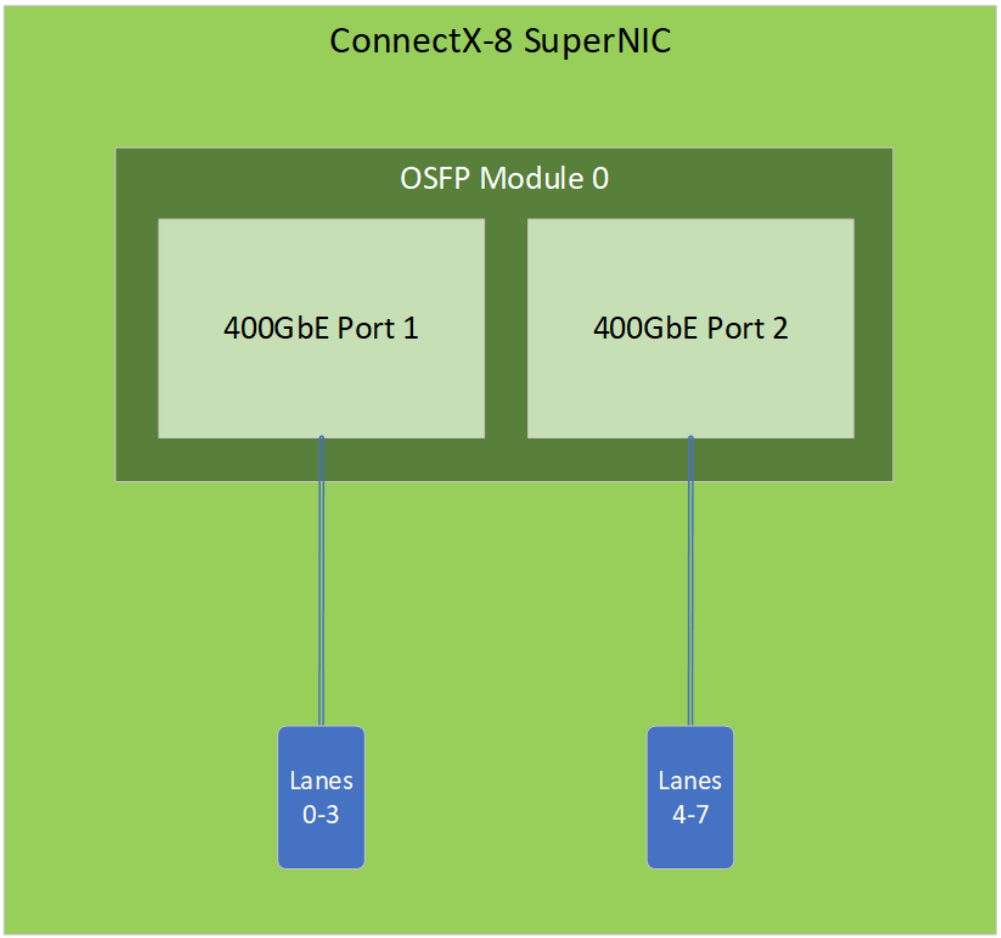

To apply this configuration, use the following mlxconfig commands:

```
mlxconfig -d <device> set LINK_TYPE_P1=2
mlxconfig -d <device> set NUM_OF_PLANES_P1 = 0
mlxconfig -d <device> set MODULE_SPLIT_M0[0..3]=1 MODULE_SPLIT_M0[4..7]=2 MODULE_SPLIT_M0[8..15]=FF
mlxconfig -d <device> set NUM_OF_PF = 2
```

8.2.1.3 C8180P Configuration #3: Eight Ethernet Ports, Each Supporting 100GbE

In this configuration, the device can be configured to support eight 100GbE Ethernet ports. The configuration is illustrated in the diagram below.

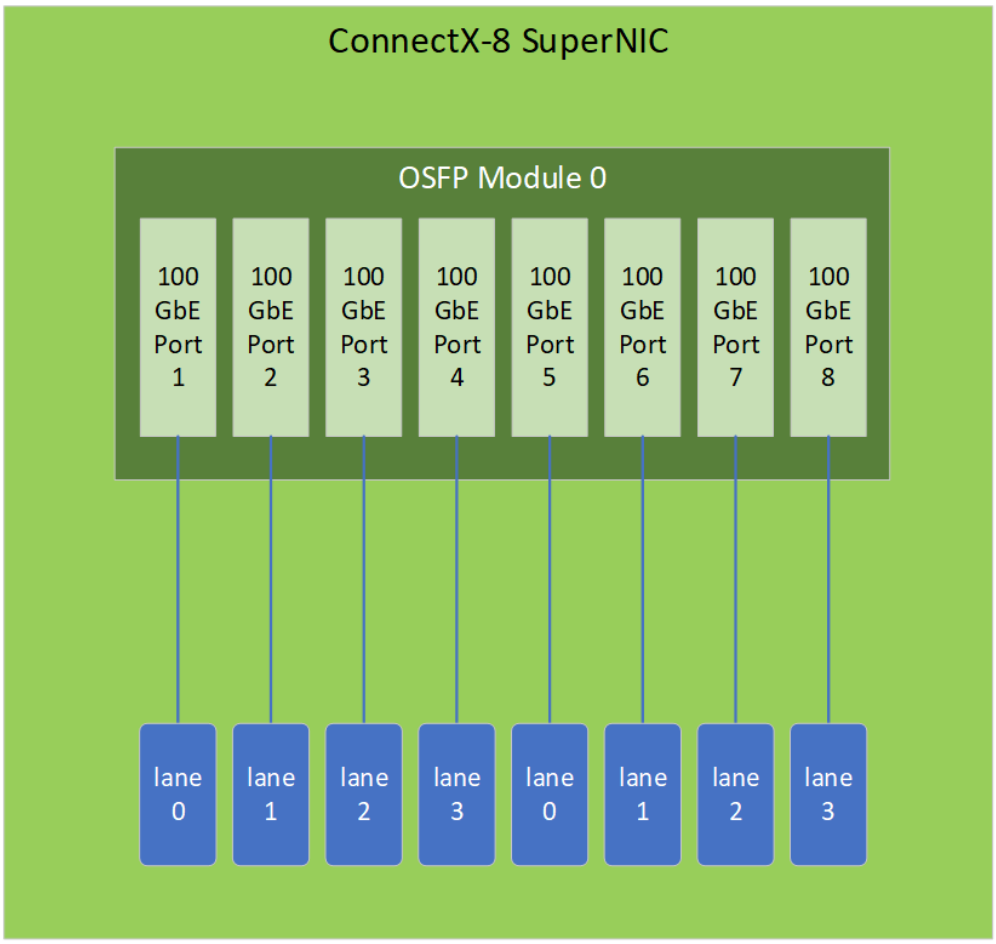

To enable this configuration, use the following mlxconfig commands:

```
mlxconfig -d <device> set LINK_TYPE_P1=2
mlxconfig -d <device> set NUM_OF_PLANES_P1 = 0
mlxconfig -d <device> set MODULE_SPLIT_M0[0]=1 MODULE_SPLIT_M0[1]=2 MODULE_SPLIT_M0[2]=3 MODULE_SPLIT_M0[3]=4
 MODULE_SPLIT_M0[4]=5 MODULE_SPLIT_M0[5]=6 MODULE_SPLIT_M0[6]=7 MODULE_SPLIT_M0[7]=8 MODULE_SPLIT_M0[8..15]=FF
mlxconfig -d <device> set NUM_OF_PF = 8
```

9 Troubleshooting

When working with the ConnectX-8 SuperNIC, you may encounter various issues related to connectivity, performance, or configuration. This section outlines common problems, their possible causes, and steps to resolve them.

9.1 General Troubleshooting

NIC Not Detected by the System	• Ensure that the SuperNIC is placed correctly • Make sure the SuperNIC slot and the SuperNIC are compatible Install the SuperNIC in a different PCI Express slot • Use the drivers that came with the SuperNIC or download the latest • Make sure your motherboard has the latest BIOS • Try to reboot the server
The SuperNIC no longer works	• Reseat the SuperNIC in its slot or a different slot, if necessary • Try using another cable • Reinstall the drivers for the network driver files may be damaged or deleted • Reboot the server
SuperNICs stopped working after installing another SuperNIC	• Try removing and re-installing all SuperNICs • Check that cables are connected properly • Make sure your motherboard has the latest BIOS
Link indicator light is off	• Try another port on the switch • Make sure the cable is securely attached • Check you are using the proper cables that do not exceed the recommended lengths • Verify that your switch and SuperNIC port are compatible
Link light is on, but with no communication established	• Check that the latest driver is loaded • Check that both the SuperNIC and its link are set to the same speed and duplex settings

9.2 DOCA-Host for Linux Driver

Environment Information	cat /etc/issue uname -a cat /proc/cupinfo | grep 'model name' | uniq ofed_info -s ifconfig -a ip link show ethtool <interface> ethtool -i <interface_of_Mellanox_port_num> ibdev2netdev
SuperNIC Detection	lspci | grep -i Mellanox
Mellanox Firmware Tool (MFT)	Download and install MFT: MFT Documentation Refer to the User Manual for installation instructions. Once installed, run: mst start mst status flint -d <mst_device> q
Ports Information	ibstat ibv_devinfo
Firmware Version Upgrade	To download the latest firmware version, refer to the NVIDIA Update and Query Utility.
Collect Log File	cat /var/log/messages dmesg >> system.log journalctl (Applicable on new operating systems) cat /var/log/syslog

9.3 Windows Troubleshooting

Environment Information	From the Windows desktop choose the Start menu and run: **msinfo32** To export system information to a text file, choose the Export option from the File menu. Assign a file name and save.
Mellanox Firmware Tool (MFT)	Download and install MFT: MFT Documentation Refer to the User Manual for installation instructions. Once installed, open a CMD window and run: WinMFT mst start mst status flint –d <mst_device> q
Ports Information	vstat
Firmware Version Upgrade	Download the latest firmware version using the PSID/board ID from here. flint –d <mst_device> –i <firmware_bin_file> b

Collect Log File	• Event log viewer
	• MST device logs:
	  • mst start
	  • mst status
	• flint –d <mst_device> dc > dump_configuration.log
	• mstdump <mst_device> dc > mstdump.log

10 Monitoring

10.1 Thermal Sensors

The adapter card incorporates the ConnectX IC, which operates in the range of temperatures between 0°C and 105°C.

Three thermal threshold definitions impact the overall system operation state and are designed to ensure timely application of server-level corrective actions, such as increased cooling, to keep the device temperature in the working range:

- Warning – 97°C (default) When the device crosses this threshold, an Over-Temperature Warning message is issued by the management SW. This may be used to increase the cooling of the system (e.g. - by increased Fans RPM). Note that this temperature threshold does not require nor lead to any action by hardware (such as adapter card shutdown).
- Critical – 102°C (default): When the device crosses this threshold, an Over-Temperature Critical message is issued by the management SW.
- Fatal – 110°C: Upon reaching this temperature, an Over-temperature Fatal message is issued. If the the device continues to heat-up, the FW will automatically shut down upon reaching 115°C threshold.

The card's thermal sensors can be read through the system's SMBus. The user can read these thermal sensors and adapt the system airflow following the readouts and the needs of the above-mentioned IC thermal requirements.

10.2 SuperNIC Heatsink

The heatsink is attached to the SuperNIC IC to dissipate the heat from the IC. It is attached by spring-loaded screws.
The SuperNIC IC has a thermal shutdown safety mechanism that automatically shuts down the SuperNICcard in cases of high-temperature events, improper thermal coupling, or heatsink removal.
For the required airflow (LFM) per OPN, please refer to the Specifications chapter.

11 Finding the GUID/MAC on the SuperNIC

Each SuperNIC is labeled with unique identifiers, including a serial number and specific protocol information; it includes a GUID (for InfiniBand protocols) and a MAC address (for Ethernet protocols, derived from the GUID).

The product revisions indicated on the labels in the following figures do not necessarily represent the latest revisions of the SuperNICs.

Board Label Example

12 Specifications

The ConnectX-8 SuperNIC is designed and validated for operation in data-center servers and other large environments that guarantee proper power supply and airflow conditions.

The SuperNIC is not intended for installation on a desktop or a workstation. Moreover, installing the adapter card in any system without proper power and airflow levels can impact the SuperNIC's functionality and potentially damage it. Failure to meet the environmental requirements listed in this user manual may void the warranty.

Please install the ConnectX-8 SuperNIC in a PCIe slot capable of supplying the required power and airflow.

12.1 C8180P SuperNICs Specifications

⚠ ConnectX-8 SuperNICs with OSFP form factor support RHS (Riding Heatsink) cage only

Physical	SuperNIC PCB Dimensions: Tall Small Form Factor (TSFF) OCP 3.0, 4.52 in. x 2.99 in (115.00mm x 76.00mm)		
	Bracket Type	900-9X85E-00EX-MP0 and 900-9X85E-00EX-MJ0: Internal lock Bracket	
	Cooling Mechanism	900-9X85E-00EX-MJ0: NVIDIA Heatsink	
		900-9X85E-00EX-MP0: Partner Cooled	
Interfaces	See Supported Interfaces		
	PCI Express Interface	Gen6 SERDES @ 64GT/s, x16 lanes (Gen 5 compatible)	
		Optional: Additional PCIe x16 Gen5 @ SERDES 32GT/s through the PCIe auxiliary card and MCIO harness	
	Networking Port: Single cage OSFP InfiniBand and Ethernet		
Data Rate	InfiniBand (Default)	XDR/NDR/HDR/HDR100/EDR/SDR	
	Ethernet	400/200/100 Gb/s Ethernet	
Protocol Support	InfiniBand: IBTA v1.7[a] Auto-Negotiation: XDR (4 lanes x 200Gb/s per) port, NDR (4 lanes x 100Gb/s per lane) port, NDR200 (2 lanes x 100Gb/s per lane) port, HDR (50Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane) port, EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/ 2X/4X SDR (2.5Gb/s per lane).		
	Ethernet: 400GAUI-4 C2M, 400GBASE-CR4, 200GAUI-2 C2M, 200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 C2M, 100GAUI-1 C2M, 100GBASE-CR4, 100GBASE-CR2, 100GBASE-CR1, 50GAUI-2 C2M, 50GAUI-1 C2M, 50GBASE-CR, 50GBASE-R2 , 40GBASE-CR4, 40GBASE-R2, 25GBASE-R, 10GBASE-R, 10GBASE-CX4, 1000BASE-CX, CAUI-4 C2M, 25GAUI C2M, XLAUI C2M , XLPPI, SFI		
Capabilities	Secure Boot Enabled, Crypto Enabled, Multi-Host and Socket Direct Capable		
Electrical and Thermal Specifications	Voltage: 12V, 3.3V, 3.3VAUX		
	TBD		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C[b]
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Regulatory	Safety: CB / cTUVus / CE		
	EMC: CE / FCC / VCCI / ICES / RCM / KC		

	RoHS: RoHS Compliant
Notes: a. The ConnectX-8 SuperNICs supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product. b. The non-operational storage temperature specifications apply to the product without its package.	

Cards Mechanical Drawings and Dimensions

All dimensions are in millimeters. The PCB mechanical tolerance is +/- 0.13mm.

C8180P SuperNICs	C8180P Partner-Cooled SuperNICs
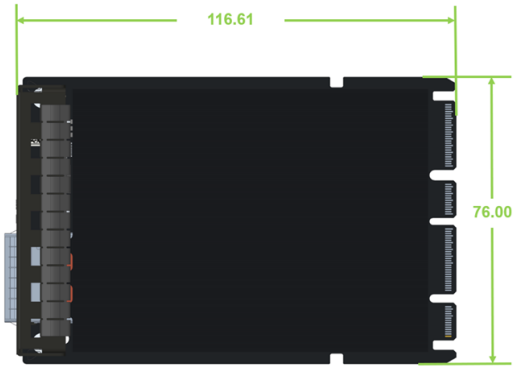 Dimensions include the Internal Lock Bracket	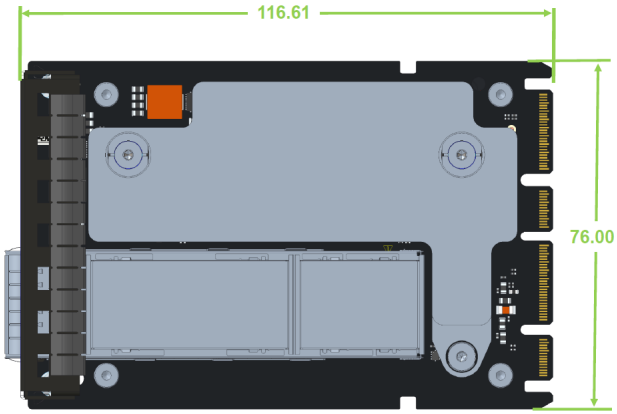 Dimensions include the Internal Lock Bracket

12.2 Brackets Mechanical Drawings and Dimensions

All dimensions are in millimeters. The PCB mechanical tolerance is +/- 0.13mm.

Single-Port OSFP Internal Lock Bracket

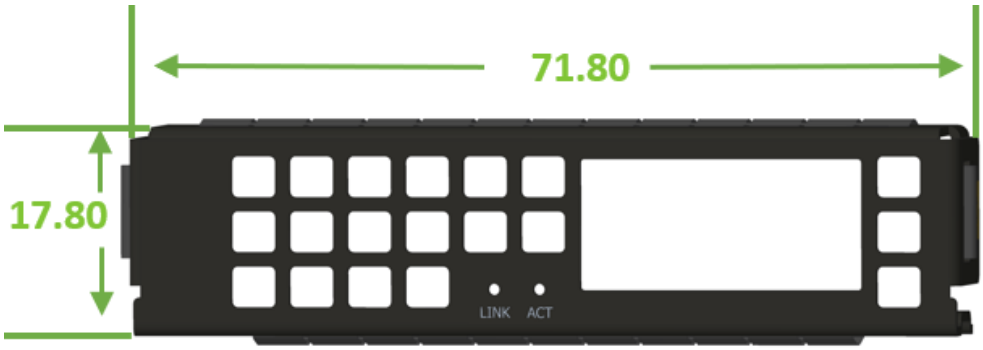

13 Document Revision History

Date	Comments/Changes
Feb. 2026	Updated bracket type across the document
Jul. 2025	Updated OPNs' marketing description
Apr. 2025	• Added 900-9X85E-00EX-MJ0 support across the document • Updated figures in Supported Interfaces • Added figures in Specifications
Feb. 2025	First release; Preliminary and Subject to Change

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