



NVIDIA ConnectX-7 Adapter Cards User Manual

Table of Contents

About This Manual	6
Intended Audience	7
Ordering Part Numbers.....	8
ConnectX-7 PCIe x16 Stand-up Adapter Cards.....	8
ConnectX-7 for Telecommunication Applications	9
ConnectX-7 Socket Direct Ready Cards for Dual-Slot Servers	10
Technical Support	13
Related Documentation	14
Document Conventions.....	15
Revision History	16
Introduction.....	17
Product Overview	17
PCIe x16 Stand-up Adapter Cards.....	17
Socket Direct Ready Cards.....	17
System Requirements	18
Package Contents	18
Features and Benefits	19
Supported Interfaces	23
ConnectX-7 Layout and Interface Information	23
Interfaces Detailed Description.....	24
ConnectX-7 IC	24
PCI Express Interface.....	24
Networking Interfaces.....	24
Networking Ports LEDs Specifications.....	25
Cablance CA-II Plus Connectors.....	26
SMBus Interface	27
Voltage Regulators	27
Hardware Installation	28
Safety Warnings	28
Installation Procedure Overview.....	29
System Requirements	29
Hardware Requirements	29

Airflow Requirements	30
Software Requirements	30
Safety Precautions	30
Pre-Installation Checklist	30
Bracket Replacement Instructions	31
Installation Instructions	31
Cables and Modules.....	32
Identifying the Card in Your System.....	32
ConnectX-7 PCIe x16 Installation Instructions.....	34
Installing the Card	34
Uninstalling the Card.....	35
ConnectX-7 Socket Direct (2x PCIe x16) Installation Instructions	37
Installing the Cards	38
Uninstalling the Cards.....	43
Driver Installation.....	46
Linux Driver Installation	46
Prerequisites.....	46
Downloading MLNX_OFED.....	46
Installing MLNX_OFED	47
Driver Load Upon System Boot.....	52
Additional Installation Procedures	54
Performance Tuning	60
Windows Driver Installation	60
Software Requirements	61
Downloading WinOF-2 Driver.....	61
Installing WinOF-2 Driver	61
Firmware Upgrade	69
VMware Driver Installation	69
Hardware and Software Requirements	69
Installing NATIVE ESXi Driver for VMware vSphere.....	69
Removing Earlier NVIDIA Drivers.....	70
Firmware Programming	70
Updating Adapter Firmware.....	71
Setting High-Speed-Port Link Type	72

mlxconfig.....	72
UEFI	72
Troubleshooting	73
General Troubleshooting.....	73
Linux Troubleshooting	73
Windows Troubleshooting.....	74
Specifications.....	75
MCX75310AAC-NEAT / MCX75310AAS-NEAT Specifications.....	75
MCX75310AAS-HEAT Specifications	76
MCX755106AC-HEAT / MCX755106AS-HEAT Specifications	77
MCX715105AS-WEAT Specifications.....	78
MCX75510AAS-HEAT Specifications	79
MCX75510AAS-NEAT Specifications	80
MCX713106AC-CEAT and MCX713106AS-CEAT Specifications	82
MCX713106AC-VEAT and MCX713106AS-VEAT Specifications	83
MCX713104AC-ADAT and MCX713104AS-ADAT Specifications	84
MCX713114TC-GEAT and MCX713114GC-GEAT Specifications	85
Cards Mechanical Drawings and Dimensions	85
Brackets Mechanical Drawings and Dimensions	86
Monitoring.....	88
Thermal Sensors.....	88
Adapter Card Heatsink	88
Finding the GUID/MAC on the Adapter Card	89
PCIe Auxiliary Card Kit	90
PCIe Auxiliary Card Package Contents	91
Channel Insertion Loss	91
Cabline CA-II Plus Harness Pinouts	91
Component Side.....	91
Print Side	100
PCIe Auxiliary Kit Technical Specifications.....	109
Technical Specifications	109
PCIe Auxiliary Card Mechanical Drawings and Dimensions.....	110
Bracket Mechanical Drawings and Dimensions	110
Cabline CA-II Plus Harnesses Mechanical Drawing	111

Document Revision History 112

About This Manual

This User Manual describes NVIDIA® ConnectX®-7 InfiniBand and Ethernet adapter cards. It provides details as to the interfaces of the board, specifications, required software and firmware for operating the board, and relevant documentation.

Intended Audience

This manual is intended for the installer and user of these cards. The manual assumes basic familiarity with InfiniBand and Ethernet network and architecture specifications.

Ordering Part Numbers

The table below lists the ordering part numbers (OPNs) for the available ConnectX-7 stand-up cards. For the Open Compute Project (OCP 3.0) cards, please refer to [NVIDIA ConnectX-7 Adapter Cards for OCP 3.0 User Manual](#).

ConnectX-7 PCIe x16 Stand-up Adapter Cards

NVIDIA SKU	Legacy OPN	Form Factor	Data Transmission Rate	No. of Ports and Type	PCIe Support	Secure Boot	Crypto	Bracket Type	Lifecycle
900-9X76-6-003N-SQ0	MCX753 10AAS-NEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	InfiniBand: NDR 400Gb/s (Default speed) Ethernet: 400GbE	Single-port OSFP	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	✓	-	Tall Bracket	Mass Production
900-9X76-6-003N-SR0	MCX753 10AAC-NEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	InfiniBand: NDR 400Gb/s (Default speed) Ethernet: 400GbE	Single-port OSFP	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	✓	✓	Tall Bracket	Engineering Samples
900-9X76-6-003N-ST0	MCX753 10AAS-HEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	InfiniBand: NDR200 200Gb/s (Default speed) Ethernet: 200GbE	Single-port OSFP	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	✓	-	Tall Bracket	Mass Production
900-9X7-AH-0086-SQ0	MCX713 106AC-CEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	Ethernet: 100GbE	Dual-port QSFP112	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	✓	✓	Tall Bracket	Mass Production
900-9X7-AH-0076-ST0	MCX713 106AS-CEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	Ethernet: 100GbE	Dual-port QSFP112	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	✓	-	Tall Bracket	Mass Production

NVIDIA SKU	Legacy OPN	Form Factor	Data Transmission Rate	No. of Ports and Type	PCIe Support	Secure Boot	Crypto	Bracket Type	Lifecycle
900-9X7 AO-00C3-STZ	MCX713 104AC-ADAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	Ethernet: 50/25GbE	Quad-port SFP56	PCIe x16 Gen 4.0 @ SERDES 16GT/s	✓	✓	Tall Bracket	Engineering Samples
900-9X7 AO-0003-ST0	MCX713 104AS-ADAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	Ethernet: 50/25GbE	Quad-port SFP56	PCIe x16 Gen 4.0 @ SERDES 16GT/s	✓	-	Tall Bracket	Engineering Samples

ConnectX-7 for Telecommunication Applications

NVIDIA SKU	Legacy OPN	Form Factor	Data Transmission Rate	No. of Ports and Type	PCIe Support	Secure Boot	Crypto	Timing Capabilities	Bracket Type	Lifecycle
900-9X7 AH-004 N-CT0	MCX713 114TC-GEAT	PCIe Full Height, Half Length 4.53 in. x 6.6 in. (115.15 mm x 167.65 mm)	Ethernet: 50/25GbE	Quad-port SFP56	PCIe x16 Gen 4.0 @ SERDES 16GT/s	✓	✓	PPS In / Out, SMAs, SyncE & PTP GM Support	Tall Bracket	Engineering Samples

ConnectX-7 Socket Direct Ready Cards for Dual-Slot Servers

NVIDIA SKU	Legacy OPN	Form Factor	Data Transmission Rate	No. of Ports and Type	PCIe Support	Socket Direct Ready - PCIe Extension Option	Secure Boot	Crypto	Bracket Type	Lifecycle
900-9X7AH-0039-STZ	MCX715105AS-WEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	InfiniBand: NDR 400Gb/s Ethernet: 400GbE (Default Speed)	Single-port QSFP112	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	Optional: PCIe x16 Gen 4.0 @ SERDES 16GT/s	✓	-	Tall Bracket	Engineering Samples
900-9X721-003N-DT1	MCX75510AAS-NEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	InfiniBand: NDR 400Gb/s	Single-port OSFP	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	Optional: PCIe x16 Gen 4.0 @ SERDES 16GT/s	✓	-	Tall Bracket	Mass Production
900-9X721-003N-DT1	MCX75510AAS-HEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	InfiniBand: NDR200 200Gb/s	Single-port OSFP	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	Optional: PCIe x16 Gen 4.0 @ SERDES 16GT/s	✓	-	Tall Bracket	Mass Production

NVIDIA SKU	Legacy OPN	Form Factor	Data Transmission Rate	No. of Ports and Type	PCIe Support	Socket Direct Ready - PCIe Extension Option	Secure Boot	Crypto	Bracket Type	Lifecycle
900-9X7AH-0078-DTZ	MCX755 106AS-HEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	InfiniBand: NDR200 200Gb/s Ethernet: 200GbE (Default Speed)	Dual-port QSFP112	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	Optional: PCIe x16 Gen 4.0 @ SERDES 16GT/s	✓	-	Tall Bracket	Mass Production
900-9X7AH-0079-DTZ	MCX755 106AC-HEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	InfiniBand: NDR200 200Gb/s Ethernet: 200GbE (Default Speed)	Dual-port QSFP112	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	Optional: PCIe x16 Gen 4.0 @ SERDES 16GT/s	✓	✓	Tall Bracket	Mass Production

Legacy (EOL) Ordering Part Numbers

NVIDIA SKU	Legacy OPN	Form Factor	Data Transmission Rate	No. of Ports and Type	PCIe Support	Secure Boot	Crypto	Timing Capabilities	Bracket Type	Lifecycle
900-9X7AH-0088-STO	MCX713 106AC-VEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	Ethernet : 200GbE	Dual-port QSFP112	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	✓	✓	-	Tall Bracket	Last Time Buy
900-9X7AH-0078-STO	MCX713 106AS-VEAT	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	Ethernet : 200GbE	Dual-port QSFP112	PCIe x16 Gen 4.0/5.0 @ SERDES 16GT/s/32GT/s	✓	-	-	Tall Bracket	Last Time Buy

NVIDIA SKU	Legacy OPN	Form Factor	Data Transmission Rate	No. of Ports and Type	PCIe Support	Secure Boot	Crypto	Timing Capabilities	Bracket Type	Lifecycle
900-9X7 AH-004 N-GT0	MCX713 114GC-GEAT	PCIe Full Height, Half Length 4.53 in. x 6.6 in. (115.15 mm x 167.65 mm)	Ethernet : 50/25GbE	Quad-port SFP56	PCIe x16 Gen 4.0 @ SERDES 16GT/s	✓	✓	Enhanced-SyncE & PTP Grand Master support and GNSS/PPS Out	Tall Bracket	Last Time Buy

For more information, please refer to [PCIe Auxiliary Card Kit](#).

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.

Related Documentation

<i>MLNX_OFED for Linux User Manual and Release Notes</i>	User Manual describing OFED features, performance, band diagnostic, tools content and configuration. See MLNX_OFED for Linux Documentation .
<i>WinOF-2 for Windows User Manual and Release Notes</i>	User Manual describing WinOF-2 features, performance, Ethernet diagnostic, tools content and configuration. See WinOF-2 for Windows Documentation .
<i>NVIDIA VMware for Ethernet User Manual</i>	User Manual and release notes describing the various components of the NVIDIA ConnectX® NATIVE ESXi stack. See VMware® ESXi Drivers Documentation .
<i>NVIDIA Firmware Utility (mlxup) User Manual and Release Notes</i>	NVIDIA firmware update and query utility used to update the firmware. Refer to Firmware Utility (mlxup) Documentation .
<i>NVIDIA Firmware Tools (MFT) User Manual</i>	User Manual describing the set of MFT firmware management tools for a single node. See MFT User Manual .
InfiniBand Architecture Specification Release 1.2.1, Vol 2 - Release 1.4, and Vol 2 - Release 1.5	InfiniBand Specifications
<i>IEEE Std 802.3 Specification</i>	IEEE Ethernet Specifications
<i>PCI Express 5.0 Specifications</i>	Industry Standard PCI Express Base and Card Electromechanical Specifications. Refer to PCI-SIG 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 10, 25, 40, 50, 100, 200, and 400GbE in Ethernet and EDR, HDR, and NDR, 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. Read more at LinkX Cables and Transceivers .
NVIDIA ConnectX-7 Electrical and Thermal Specifications	Electrical and Thermal specifications are provided in the "NVIDIA ConnectX-7 Adapter Cards Electrical and Thermal Specifications" document available at NVOnline following login.

Document Conventions

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.

Revision History

A list of the changes made to this document is provided in [Document Revision History](#).

Introduction

Product Overview

The NVIDIA ConnectX-7 family of network adapters supports both the InfiniBand and Ethernet protocols. It enables a wide range of smart, scalable, and feature-rich networking solutions that address traditional enterprise needs up to the world's most demanding AI, scientific computing, and hyperscale cloud data center workloads.

ConnectX-7 network adapters 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 PCIe stand-up cards, for the OCP 3.0 cards, please refer to [NVIDIA ConnectX-7 Cards for OCP Spec 3.0 User Manual](#).

 Make sure to use a PCIe slot capable of supplying the required power and airflow to the ConnectX-7, as stated in the [Specifications](#) chapter.

PCIe x16 Stand-up Adapter Cards

ConnectX-7 HCAs are available in various configurations; Single-port 400Gb/s or 200Gb/s, with octal small form-factor pluggable (QSFP) connectors or Dual-port 100 or 200Gb/s with quad small form-factor pluggable (QSFP112) connectors on PCIe standup half-height, half-length (HHHL) form factor, with options for NVIDIA Socket Direct. Also available, Dual-port 50/25 GbE with quad small form-factor pluggable (SFP56) connectors on PCIe standup full-height, half-length (FHHL) form factor, with timing capabilities.

ConnectX-7 cards can either support both InfiniBand and Ethernet, or Ethernet only, as described in the below table. The inclusive list of OPNs is available [here](#).

Supported Protocols	Port Type	Supported Speed
Ethernet Only Card	Dual-port QSFP112	• 100GbE
	Quad-port SFP56	• 50/25GbE
InfiniBand and Ethernet Cards	Single-port OSFP	• NDR 400Gb/s and 400GbE • NDR200 200Gb/s and 200GbE

Socket Direct Ready Cards

The Socket Direct technology offers improved performance to dual-socket servers by enabling direct access from each CPU in a dual-socket server to the network through its dedicated PCIe interface.

NVIDIA offers ConnectX-7 Socket Direct adapter cards, which enable 400Gb/s or 200Gb/s connectivity, and also for servers with PCIe Gen 4.0 capability. The adapter's 32-lane PCIe bus is split into two 16-lane buses, with one bus accessible through a PCIe x16 edge connector and the other bus through an x16 Auxiliary PCIe Connection card. The two cards should be installed into two PCIe x16 slots and connected using two Cabline SA-II Plus harnesses.

To use this card in the Socket-Direct configuration, please order the additional PCIe Auxiliary Card kit according to the desired harness length. Cards that support socket direct can function as separate x16 PCIe cards.

Socket Direct cards can support both InfiniBand and Ethernet, or InfiniBand only, as described below.

Supported Protocols	Port Type	Supported Speed
InfiniBand Only	Single-port OSFP	- NDR 400Gb/s - NDR200 200Gb/s
InfiniBand and Ethernet	Dual-port QSFP112	NDR200 200Gb/s and 200GbE
	Single-port QSFP112	NDR 400Gb/s and 400GbE

For more information on the passive PCIe Auxiliary kit, please refer to [PCIe Auxiliary Card Kit](#).

System Requirements

Item	Description
PCI Express slot	In PCIe x16 Configuration PCIe Gen 5.0 (32GT/s) through x16 edge connector. In Socket Direct Configuration (2x PCIe x16) - PCIe Gen 4.0/5.0 SERDES @16/32GT/s through edge connector - PCIe Gen 4.0 SERDES @16GT/s through PCIe Auxiliary Connection Card
System Power Supply	Refer to Specifications
Operating System	- In-box drivers for major operating systems: - Linux: RHEL, Ubuntu - Windows - Virtualization and containers - VMware ESXi (SR-IOV) - Kubernetes - OpenFabrics Enterprise Distribution (OFED) - OpenFabrics Windows Distribution (WinOF-2)
Connectivity	- Interoperable with 1/10/25/40/50/100/200/400 Gb/s Ethernet switches and SDR/DDR/EDR/HDR100/HDR/NDR200/NDR InfiniBand switches - Passive copper cable with ESD protection - Powered connectors for optical and active cable support

Package Contents

Category	Qty	Item
Cards	1	ConnectX-7 adapter card
Accessories	1	Adapter card short bracket
	1	Adapter card tall bracket (shipped assembled on the card)

Features and Benefits

⚠ Make sure to use a PCIe slot capable of supplying the required power and airflow to the ConnectX-7 cards 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.

PCI Express (PCIe)	According to the OPN you have purchased, the card uses the following PCIe express interfaces: • PCIe x16 configurations: PCIe Gen 4.0/5.0 (16GT/s / 32GT/s) through x16 edge connector. • 2x PCIe x16 configurations (Socket-Direct): PCIe Gen 4.0/5.0 (SERDES @ 16GT/s / 32GT/s) through x16 edge connector PCIe Gen 4.0 SERDES @ 16GT/s through PCIe Auxiliary Connection Card																														
InfiniBand Architecture Specification v1.5 compliant	ConnectX-7 delivers low latency, high bandwidth, and computing efficiency for high-performance computing (HPC), artificial intelligence (AI), and hyperscale cloud data center applications. ConnectX-7 is InfiniBand Architecture Specification v1.5 compliant. InfiniBand Network Protocols and Rates: <table> <tr> <th rowspan="2">Protocol</th><th rowspan="2">Standard</th><th colspan="2">Rate (Gb/s)</th><th rowspan="2">Comments</th></tr> <tr> <th>4x Port (4 Lanes)</th><th>2x Ports (2 Lanes)</th></tr> <tr> <td>NDR/NDR200</td><td>IBTA Vol2 1.5</td><td>425</td><td>212.5</td><td>PAM4 256b/257b encoding and RS-FEC</td></tr> <tr> <td>HDR/HDR100</td><td>IBTA Vol2 1.4</td><td>212.5</td><td>106.25</td><td>PAM4 256b/257b encoding and RS-FEC</td></tr> <tr> <td>EDR</td><td>IBTA Vol2 1.3.1</td><td>103.125</td><td>51.5625</td><td>NRZ 64b/66b encoding</td></tr> <tr> <td>FDR</td><td>IBTA Vol2 1.2</td><td>56.25</td><td>N/A</td><td>NRZ 64b/66b encoding</td></tr> </table>				Protocol	Standard	Rate (Gb/s)		Comments	4x Port (4 Lanes)	2x Ports (2 Lanes)	NDR/NDR200	IBTA Vol2 1.5	425	212.5	PAM4 256b/257b encoding and RS-FEC	HDR/HDR100	IBTA Vol2 1.4	212.5	106.25	PAM4 256b/257b encoding and RS-FEC	EDR	IBTA Vol2 1.3.1	103.125	51.5625	NRZ 64b/66b encoding	FDR	IBTA Vol2 1.2	56.25	N/A	NRZ 64b/66b encoding
Protocol	Standard	Rate (Gb/s)		Comments																											
		4x Port (4 Lanes)	2x Ports (2 Lanes)																												
NDR/NDR200	IBTA Vol2 1.5	425	212.5	PAM4 256b/257b encoding and RS-FEC																											
HDR/HDR100	IBTA Vol2 1.4	212.5	106.25	PAM4 256b/257b encoding and RS-FEC																											
EDR	IBTA Vol2 1.3.1	103.125	51.5625	NRZ 64b/66b encoding																											
FDR	IBTA Vol2 1.2	56.25	N/A	NRZ 64b/66b encoding																											

Up to 400 Gigabit Ethernet	ConnectX-7 adapter cards comply with the following IEEE 802.3 standards: 400GbE / 200GbE / 100GbE / 50GbE / 40GbE / 25GbE / 10GbE	
	Protocol	MAC Rate
	IEEE802.3ck	100/200/400Gb/s Gigabit Ethernet (Include ETC enhancement)
	IEEE802.3cd IEEE802.3bs IEEE802.3cm IEEE802.3cn IEEE802.3cu	50/100/200/400Gb/s Gigabit Ethernet (Include ETC enhancement)
	IEEE 802.3bj IEEE 802.3bm	100 Gigabit Ethernet
	IEEE 802.3by Ethernet Technology Consortium	25/50 Gigabit Ethernet
	IEEE 802.3ba	40 Gigabit Ethernet
	IEEE 802.3ae	10 Gigabit Ethernet
	IEEE 802.3cb	2.5/5 Gigabit Ethernet (For 2.5: support only 2.5 x1000BASE-X)
	IEEE 802.3ap	Based on auto-negotiation and KR startup
	IEEE 802.3ad IEEE 802.1AX	Link Aggregation
	IEEE 802.1Q IEEE 802.1P VLAN tags and priority	
	IEEE 802.1Qau (QCN) Congestion Notification IEEE 802.1Qaz (ETS) IEEE 802.1Qbb (PFC) IEEE 802.1Qbg IEEE 1588v2 IEEE 802.1AE (MACSec) Jumbo frame support (9.6KB)	
Memory Components	- SPI - includes 256Mbit SPI Quad Flash device. - FRU EEPROM - Stores the parameters and personality of the card. The EEPROM capacity is 128Kbit. FRU I2C address is (0x50) and is accessible through the PCIe SMBus. (Note: Address 0x58 is reserved.)	
Overlay Networks	In order to better scale their networks, datacenter 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-7 effectively addresses this by providing advanced NVGRE and VXLAN hardware offloading engines that encapsulate and de-capsulate the overlay protocol.	
Quality of Service (QoS)	Support for port-based Quality of Service enabling various application requirements for latency and SLA.	

Hardware-based I/O Virtualization	ConnectX-7 provides dedicated adapter resources and guaranteed isolation and protection for virtual machines within the server.
Storage Acceleration	A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols can leverage: • RDMA for high-performance storage access • NVMe over Fabric offloads for the target machine • NVMe over TCP acceleration
SR-IOV	ConnectX-7 SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server.
High-Performance Accelerations	• Collective operations offloads • Vector collective operations offloads • MPI tag matching • MPI_Alltoall offloads • Rendezvous protocol offload
RDMA Message Rate	330-370 million messages per second.
Secure Boot	The secure boot process assures booting of authentic firmware/software that is intended to run on ConnectX-7. This is achieved using cryptographic primitives using asymmetric cryptography. ConnectX-7 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.
Advanced storage capabilities	Block-level encryption and checksum offloads.




Host Management	ConnectX-7 technology maintains support for host manageability through a BMC. ConnectX-7 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 adapter card. 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
Accurate timing	NVIDIA offers a full IEEE 1588v2 PTP software solution, as well as time-sensitive related features called “5T”. NVIDIA PTP and 5T software solutions are designed to meet the most demanding PTP profiles. ConnectX-7 incorporates an integrated Hardware Clock (PHC) that allows ConnectX-7 to achieve sub 20u Sec accuracy and also offers many timing-related functions such as time-triggered scheduling or time-based SND accelerations (time-based ASAP²). Furthermore, 5T technology enables the software application to transmit fronthaul (ORAN) compatible in high bandwidth. The PTP part supports the subordinate clock, master clock, and boundary clock. ConnectX-7 PTP solution allows you to run any PTP stack on your host. With respect to testing and measurements, selected NVIDIA adapters allow you to use the PPS-out signal from the onboard SMA connector, ConnectX-7 also allows measuring PTP in scale, with a PPS-In signal. The PTP HW clock on the Network adapter will be sampled on each PPS-In signal, and the timestamp will be sent to the SW.
RDMA and RDMA over Converged Ethernet (RoCE)	ConnectX-7, 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-7 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-7 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 ASAP²® • Tunneling encapsulation/decapsulation

Supported Interfaces

This section describes the ConnectX-7 supported interfaces. 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 adapter card.

ConnectX-7 Layout and Interface Information

Single-Port QSFP112 Adapter Cards	Dual-Port QSFP112 Adapter Cards
TBD	
Quad-Port SFP56 Cards with PPS IN/OUT	Quad-port SFP56 Cards
TBD	
Single-Port OSFP Adapter Cards	
	

Item	Interface	Description
1	ConnectX-7 IC	ConnectX-7 Integrated Circuit
2	PCI Express Interface	PCIe Gen 4.0/5.0 through x16 edge connector
3	Networking Interfaces	Network traffic is transmitted through the adapter card networking 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
5	Cabline CA-II Plus Connectors	In Socket-Direct ready cards, two Cabline CA-II plus connectors are populated to allow connectivity to an additional PCIe x16 Auxiliary card. Applies to OPNs: MCX755106AC-HEAT, MCX755106AS-HEAT, MCX75510AAS-HEAT, MCX75510AAS-NEAT.

Interfaces Detailed Description

ConnectX-7 IC

The ConnectX-7 family of adapter IC devices delivers InfiniBand and Ethernet connectivity paired with best-in-class hardware capabilities that accelerate and secure cloud and data-center workloads.

PCI Express Interface

ConnectX-7 adapter cards support PCI Express Gen 5.0 (4.0 and 3.0 compatible) through x16 edge connector. The following lists PCIe interface features:

- PCIe Gen 5.0 compliant, 4.0, 3.0, 2.0 and 1.1 compatible
- 2.5, 5.0, 8.0, 16.0 and 32GT/s link rate x16/x32 (Socket-Direct configuration)
- Support for PCIe bifurcation: Auto-negotiates to x32, x16, x8, x4, x2, or x1
- NVIDIA Multi-Host™ supports connection of up to 4x hosts
- Transaction layer packet (TLP) processing hints (TPH)
- PCIe switch Downstream Port Containment (DPC)
- Advanced error reporting (AER)
- Access Control Service (ACS) for peer-to-peer secure communication
- Process Address Space ID (PASID)
- Address translation services (ATS)
- Support for MSI/MSI-X mechanisms
- Support for SR-IOV

Networking Interfaces

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

Protocol	Specifications
Ethernet	The network ports comply with the IEEE 802.3 Ethernet standards listed in Features and Benefits . Ethernet traffic is transmitted through the networking connectors on the adapter card.
InfiniBand	The network ports are compliant with the <i>InfiniBand Architecture Specification, Release 1.5</i> . InfiniBand traffic is transmitted through the cards' networking connectors.

Networking Ports LEDs Specifications

For the networking ports LEDs description, follow the below table depending on the ConnectX-7 SKU you have purchased.

SKU	LEDs Scheme
900-9X7AO-0003-ST0 900-9X7AO-00C3-STZ	Scheme 1: One Bi-Color LED
All cards	Scheme 2: Two LEDs

Scheme 1: One Bi-Color LED

There is one bi-color (Yellow and Green) I/O LED per port to indicate port speed and link status.

There is one bi-color (Yellow and Green) I/O LED per port to indicate port speed and link status.


State	Bi-Color LED (Yellow/Green)		
Beacon command for locating the adapter card	1Hz blinking Yellow		
Error	4Hz blinking Yellow Indicates an error with the link. The error can be one of the following:		
	Error Type	Description	LED Behavior
	I ² C	I ² C 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
Physical Activity	The Green LED will blink.		
Link Up	The Green LED will be solid.		
Physical Up (IB Only)	The Yellow LED will be solid.		

Scheme 2: Two LEDs

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

- LED1 is a bi-color LED (Yellow and Green)

- LED2 is a single-color LED (Green)



State	Bi-Color LED (Yellow/Green)	Single Color LED (Green)									
Beacon command for locating the adapter card	1Hz blinking Yellow	OFF									
Error	4Hz blinking Yellow Indicates an error with the link. The error can be one of the following: <table border="1"> <thead> <tr> <th>Error Type</th><th>Description</th><th>LED Behavior</th></tr> </thead> <tbody> <tr> <td>I²C</td><td>I²C access to the networking ports fails</td><td>Blinks until error is fixed</td></tr> <tr> <td>Over-current</td><td>Over-current condition of the networking ports</td><td>Blinks until error is fixed</td></tr> </tbody> </table>	Error Type	Description	LED Behavior	I ² C	I ² C 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
Error Type	Description	LED Behavior									
I ² C	I ² C 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									
Physical Activity	The Green LED will blink.	Blinking									
Link Up	In full port speed: the Green LED is solid In less than full port speed: the Yellow LED is solid	ON									

Cabline CA-II Plus Connectors

⚠ Applies to OPNs: MCX755106AC-HEAT, MCX755106AS-HEAT, MCX75510AAS-HEAT, MCX75510AAS-NEAT.

The Cabline CA-II connectors on the Socket-Direct ready cards enable connectivity to an additional Auxiliary PCIe x16 Connection card through the Cabline CA-II harnesses.

SMBus Interface

ConnectX-7 technology maintains support for manageability through a BMC. ConnectX-7 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 adapter. For configuring the adapter for the specific manageability solution in use by the server, please contact NVIDIA Support.

Voltage Regulators

The voltage regulator power is derived from the PCI Express edge connector 12V supply pins. These voltage supply pins feed on-board regulators that provide the necessary power to the various components on the card.

Hardware Installation

Installation and initialization of ConnectX-7 adapter cards require attention to the mechanical attributes, power specification, and precautions for electronic equipment.

Safety Warnings

❗ Safety warnings are provided here in the English language. For safety warnings in other languages, refer to the [Adapter 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

Installation Procedure Overview

The installation procedure of ConnectX-7 adapter cards involves the following steps:

Step	Procedure	Direct Link
1	Check the system's hardware and software requirements.	System Requirements
2	Pay attention to the airflow consideration within the host system	Airflow Requirements
3	Follow the safety precautions	Safety Precautions
4	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-7 PCIe x16 adapter card in the system	ConnectX-7 PCIe x16 Adapter Cards Installation Instructions
	Install the ConnectX-7 2x PCIe x16 Socket Direct adapter card in the system	ConnectX-7 Socket Direct (2x PCIe x16) Installation Instructions
8	Connect cables or modules to the card	Cables and Modules
9	Identify ConnectX-7 in the system	Identifying Your Card

System Requirements

Hardware Requirements

⚠ 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.

⚠ For proper operation and performance, please make sure to use a PCIe slot with a corresponding bus width that can supply sufficient power to your card. Refer to the [Specifications](#) section of the manual for more power requirements.


⚠ Please make sure to install the ConnectX-7 cards in a PCIe slot that is capable of supplying the required power as stated in [Specifications](#).

ConnectX-7 Configuration	Hardware Requirements
PCIe x16	A system with a PCI Express x16 slot is required for installing the card.
Socket Direct 2x PCIe x16 (dual-slot server)	A system with two PCIe x16 slots is required for installing the cards.

Airflow Requirements

ConnectX-7 adapter cards are offered with two airflow patterns: from the heatsink to the network ports, and vice versa, as shown below.

Please refer to the [Specifications](#) section for airflow numbers for each specific card model.



⚠ All cards in the system should be planned with the same airflow direction.

Software Requirements

- See [System Requirements](#) section under the Introduction section.
- Software Stacks - NVIDIA® OpenFabrics Enterprise Distribution for Linux (MLNX_OFED), WinOF-2 for Windows, and VMware. See the [Driver Installation](#) section.

Safety Precautions

The adapter 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.

- Remove any metallic objects from your hands and wrists.
- Make sure to use only insulated tools.
- Verify that the system is powered off and is unplugged.
- It is strongly recommended to use an ESD strap or other antistatic devices.

Pre-Installation Checklist

- Unpack the ConnectX-7 Card; Unpack and remove the ConnectX-7 card. Check against the package contents list that all the parts have been sent. Check the parts for visible damage that may have occurred during shipping. Please note that the cards must be placed on an antistatic surface. For package contents please refer to [Package Contents](#).

⚠ Please note that if the card is removed hastily from the antistatic bag, the plastic ziplock may harm the EMI fingers on the networking connector. Carefully remove the card from the antistatic bag to avoid damaging the EMI fingers.

- Shut down your system if active; Turn off the power to the system, and disconnect the power cord. Refer to the system documentation for instructions. Before you install the ConnectX-7 card, make sure that the system is disconnected from power.
- (Optional) Check the mounting bracket on the ConnectX-7 or PCIe Auxiliary Connection Card; If required for your system, replace the full-height mounting bracket that is shipped mounted on the card with the supplied low-profile bracket. Refer to [Bracket Replacement Instructions](#).

Bracket Replacement Instructions

The ConnectX-7 card and PCIe Auxiliary Connection card are usually shipped with an assembled high-profile 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 high-profile bracket with the short bracket that is included in the shipping box, please follow the instructions in this section.

- ⚠ During the bracket replacement procedure, do not pull, bend, or damage the EMI fingers cage. It is recommended to limit bracket replacements to three times.

To replace the bracket you will need the following parts:

- The new brackets of the proper height
- The 2 screws saved from the removal of the bracket

Removing the Existing Bracket

1. Using a torque driver, remove the two screws holding the bracket in place.
2. Separate the bracket from the ConnectX-7 card.

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

3. Save the two screws.

Installing the New Bracket

1. Place the bracket onto the card until the screw holes line up.

- ⚠ Do not force the bracket onto the adapter card.

2. Screw on the bracket using the screws saved from the bracket removal procedure above.

- ⚠ Use a torque driver to apply up to 2 lbs-in torque on the screws.

Installation Instructions

This section provides detailed instructions on how to install your adapter card in a system. Choose the installation instructions according to the ConnectX-7 configuration you would like to use.

OPNs	Installation Instructions
All ConnectX-7 cards	ConnectX-7 (PCIe x16) Adapter Card

OPNs	Installation Instructions
MCX755106AC-HEAT MCX755106AS-HEAT MCX75510AAS-HEAT MCX75510AAS-NEAT	ConnectX-7 Socket Direct (2x PCIe x16) Adapter Card

Cables and Modules

Cable Installation

Before connecting a cable to the adapter card, ensure that the bracket is fastened to the server chassis using a screw to prevent movement or unplugging of the card 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 adapter card. Do this by using a cable holder or tying the cable to the rack.
 - b. Determine the correct orientation of the connector to the card before inserting the connector. Do not try and insert the connector upside down. This may damage the adapter card.
 - c. Insert the connector into the adapter card. Be careful to insert the connector straight into the cage. Do not apply any torque, up or down, to the connector cage in the adapter card.
 - 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 card.

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.

Identifying the Card in Your System

On Linux


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

ConnectX-7 Card Configuration	Output Example
Single-port Socket Direct Card (2x PCIe x16)	<pre>[root@mftqa-009 ~]# lspci grep mellanox -i a3:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7] e3:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7]</pre>
Dual-port Socket Direct Card (2x PCIe x16)	<pre>[root@mftqa-009 ~]# lspci grep mellanox -i 05:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7] 05:00.1 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7] 82:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7] 82:00.1 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7]</pre> 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 PCIe x16 Card	<pre>[root@mftqa-009 ~]# lspci grep mellanox -ia 3:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7]</pre>

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-7 adapter card.
3. Right-click the mouse on your adapter's row and select Properties to display the adapter card 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 1021 (for ConnectX-7) - this is a valid NVIDIA PCI Device ID.

⚠ If the PCI device does not have an NVIDIA adapter ID, return to Step 2 to check another device.

⚠ The list of NVIDIA PCI Device IDs can be found at the [PCI ID repository](#).

ConnectX-7 PCIe x16 Installation Instructions


Installing the Card

⚠ This section applies to all cards.
In case you would like to use the Socket Direct configuration (PCIe x32) that is available in MCX75510AAS-HEAT, MCX75510AAS-NEAT and MCX755106AS-HEAT, please refer to [ConnectX-7 Socket Direct \(2x PCIe x16\) Adapter Card Installation instructions](#).

⚠ Please make sure to install the ConnectX-7 cards in a PCIe slot that is capable of supplying the required power and airflow as stated in [Specifications](#).

⚠ The below images are for illustration purposes only.


- Connect the adapter Card in an available PCI Express x16 slot in the chassis.
Step 1: Locate an available PCI Express x16 slot and insert the adapter card to the chassis.



Step 2: Applying even pressure at both corners of the card, insert the adapter card in a PCI Express slot until firmly seated.

⚠ Do not use excessive force when seating the card, as this may damage the chassis.

- Secure the adapter card to the chassis.
Secure the bracket to the chassis with the bracket screw.



Uninstalling the Card

Safety Precautions


The adapter is installed in a system that operates with voltages that can be lethal. Before uninstalling the adapter card, please observe the following precautions to avoid injury and prevent damage to system components.

1. Remove any metallic objects from your hands and wrists.
2. It is strongly recommended to use an ESD strap or other antistatic devices.
3. Turn off the system and disconnect the power cord from the server.


Card Removal

 Please note that the following images are for illustration purposes only.

1. Verify that the system is powered off and unplugged.
2. Wait 30 seconds.
3. To remove the card, disengage the retention mechanisms on the bracket (clips or screws).



4. Holding the adapter card from its center, gently pull the ConnectX-7 card out of the PCI Express slot.



ConnectX-7 Socket Direct (2x PCIe x16) Installation Instructions

⚠ This section applies to the following adapter cards when used as Socket Direct cards in dual-socket servers.

- MCX755106AS-HEAT
- MCX755106AC-HEAT
- MCX75510AAS-NEAT
- MCX75510AAS-HEAT
- MCX715105AS-WEAT

⚠ The below images are for illustration purposes only.

The hardware installation section uses the terminology of white and black harnesses to differentiate between the two supplied cables. Due to supply chain variations, some cards may be provided with two black harnesses instead. To clarify the difference between these two harnesses, one black harness was marked with a “WHITE” label and the other with a “BLACK” label.

The Cabline harness marked with the “WHITE” label should be connected to the connector on the ConnectX-7 and PCIe card engraved with “White Cable,” while the one marked with the “BLACK” label should be connected to the connector on the ConnectX-7 and PCIe card engraved with “Black Cable”.

⚠ The harnesses' minimal bending radius is 10[mm].

Installing the Cards

⚠ The installation instructions include steps that involve a retention clip to be used while connecting the Cabline harnesses to the cards. Please note that this is an optional accessory.


⚠ Please make sure to install the ConnectX-7 cards in a PCIe slot capable of supplying the required power and airflow as stated in the [Specifications](#).

➤ Connect the adapter card with the Auxiliary connection card using the supplied Cabline CA-II Plus harnesses.


Step 1: Slide the black and white Cabline CA-II Plus harnesses through the retention clip while ensuring the clip opening is facing the plugs.



Step 2: Plug the Cabline CA-II Plus harnesses into the ConnectX-7 adapter card while paying attention to the color coding. As indicated on both sides of the card, plug the black harness into the component side and the white harness into the print side.




Step 2: Verify the plugs are locked.



Step 3: Slide the retention clip latches through the cutouts on the PCB. The latches should face the annotation on the PCB.



Step 4: Clamp the retention clip. Verify both latches are firmly locked.




Verify that both latches are firmly snapped




Step 5: Slide the Cabline CA-II Plus harnesses through the retention clip. Make sure that the clip opening is facing the plugs.



Step 6: Plug the Cabline CA-II Plus harnesses into the PCIe Auxiliary Card. As indicated on both sides of the Auxiliary connection card, plug the black harness into the component side and the white harness into the print side.





Step 7: Verify the plugs are locked.



Step 8: Slide the retention clip through the cutouts on the PCB. Ensure latches are facing "Black Cable" annotation, as seen in the picture below.



Step 9: Clamp the retention clip. Verify both latches are firmly locked.




Verify that both latches are firmly snapped

➤ Connect the ConnectX-7 adapter and PCIe Auxiliary Connection cards in available PCI Express x16 slots in the chassis.


Step 1: Locate two available PCI Express x16 slots.

Step 2: Applying even pressure at both corners of the cards, insert the adapter card in the PCI Express slots until firmly seated.




⚠ Do not use excessive force when seating the cards, as this may damage the system or the cards.

Step 3: Applying even pressure at both corners of the cards, insert the Auxiliary Connection card in the PCI Express slots until firmly seated.



- Secure the ConnectX-7 adapter and PCIe Auxiliary Connection Cards to the chassis.
Step 1: Secure the brackets to the chassis with the bracket screws.



Uninstalling the Cards

Safety Precautions


The adapter is installed in a system that operates with voltages that can be lethal. Before uninstalling the adapter card, please observe the following precautions to avoid injury and prevent damage to system components.

1. Remove any metallic objects from your hands and wrists.
2. Using an ESD strap or other antistatic devices is strongly recommended.
3. Turn off the system and disconnect the power cord from the server.


Card Removal

 Please note that the following images are for illustration purposes only.

1. Verify that the system is powered off and unplugged.
2. Wait 30 seconds.
3. To remove the card, disengage the retention mechanisms on the brackets (clips or screws).



4. Holding the adapter card from its center, gently pull the ConnectX-7 and Auxiliary Connections cards out of the PCI Express slot.



Driver Installation

Please refer to the relevant driver installation section.

- [Linux Driver Installation](#)
- [Windows Driver Installation](#)
- [VMware Driver Installation](#)

Linux Driver Installation

This section describes how to install and test the MLNX_OFED for Linux package on a single server with a ConnectX-7 adapter card installed.

Prerequisites

Requirements	Description
Platforms	A server platform with a ConnectX-7 InfiniBand/Ethernet adapter card 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 <i>MLNX_OFED Release Notes</i> .
Installer Privileges	The installation requires administrator (root) privileges on the target machine.

Downloading MLNX_OFED

1. Verify that the system has a network adapter installed by running `lspci` command. The below table provides output examples per ConnectX-7 card configuration.

ConnectX-7 Card Configuration	Output Examples
Single-port Socket Direct Card (2x PCIe x16)	<pre>[root@mftga-009 ~]# lspci grep mellanox -i a3:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7] e3:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7]</pre>

ConnectX-7 Card Configuration	Output Examples
Dual-port Socket Direct Card (2x PCIe x16)	<pre>[root@mftqa-009 ~]# lspci grep mellanox -i 05:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7] 05:00.1 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7] 82:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7] 82:00.1 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7]</pre> 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 number 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 number 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 PCIe x16 Card	<pre>[root@mftqa-009 ~]# lspci grep mellanox -ia 3:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7]</pre>
Dual-port PCIe x16 Card	<pre>[root@mftqa-009 ~]# lspci grep mellanox -ia 86:00.0 Network controller: Mellanox Technologies MT2910 Family [ConnectX-7] 86:00.1 Network controller: Mellanox Technologies MT2910 Family [ConnectX-7]</pre>

2. Download the ISO image to your host.

The image's name has the format `MLNX_OFED_LINUX-<ver>-<OS label><CPU arch>.iso`.

You can download and install the latest OpenFabrics Enterprise Distribution (OFED) software package available via the NVIDIA web site at nvidia.com/en-us/networking → Products → Software → InfiniBand Drivers → [NVIDIA MLNX OFED](#)

- Scroll down to the Download wizard, and click the Download tab.
- Choose your relevant package depending on your host operating system.
- Click the desired ISO/tgz package.
- To obtain the download link, accept the End User License Agreement (EULA).

3. Use the Hash utility to confirm the file integrity of your ISO image. Run the following command and compare the result to the value provided on the download page.

```
SHA256 MLNX_OFED_LINUX-<ver>-<OS label>.iso
```

Installing MLNX_OFED

Installation Script

The installation script, `mlnxofedinstall`, performs the following:

- Discovers the currently installed kernel

- Uninstalls any software stacks that are part of the standard operating system distribution or another vendor's commercial stack
- Installs the MLNX_OFED_LINUX binary RPMs (if they are available for the current kernel)
- Identifies the currently installed InfiniBand and Ethernet network adapters and automatically upgrades the firmware

Note: To perform a firmware upgrade using customized firmware binaries, a path can be provided to the folder that contains the firmware binary files, by running `--fw-image-dir`. Using this option, the firmware version embedded in the MLNX_OFED package will be ignored.

Example:

```
./mlnxofedinstall --fw-image-dir /tmp/my_fw_bin_files
```

 If the driver detects unsupported cards on the system, it will abort the installation procedure. To avoid this, make sure to add `--skip-unsupported-devices-check` flag during installation.

Usage

```
./mnt/mlnxofedinstall [OPTIONS]
```

The installation script removes all previously installed OFED packages and re-installs from scratch. You will be prompted to acknowledge the deletion of the old packages.

 Pre-existing configuration files will be saved with the extension `".conf.rpmsave"`.

- If you need to install OFED on an entire (homogeneous) cluster, a common strategy is to mount the ISO image on one of the cluster nodes and then copy it to a shared file system such as NFS. To install on all the cluster nodes, use cluster-aware tools (such as `pdsh`).
- If your kernel version does not match with any of the offered pre-built RPMs, you can add your kernel version by using the `"mlnx_add_kernel_support.sh"` script located inside the MLNX_OFED package.

 On Redhat and SLES distributions with errata kernel installed there is no need to use the `mlnx_add_kernel_support.sh` script. The regular installation can be performed and weak-updates mechanism will create symbolic links to the MLNX_OFED kernel modules.

 If you regenerate kernel modules for a custom kernel (using `--add-kernel-support`), the packages installation will not involve automatic regeneration of the `initramfs`. In some cases, such as a system with a root filesystem mounted over a ConnectX card, not regenerating the `initramfs` may even cause the system to fail to reboot.

In such cases, the installer will recommend running the following command to update the initramfs:

```
dracut -f
```

On some OSs, `dracut -f` might result in the following error message which can be safely ignore.

```
libkmod: kmod_module_new_from_path: kmod_module 'mdev' already exists with different path
```

The “`mlnx_add_kernel_support.sh`” script can be executed directly from the `mlnxofedinstall` script. For further information, please see '--add-kernel-support' option below.

 On Ubuntu and Debian distributions drivers installation use Dynamic Kernel Module Support (DKMS) framework. Thus, the drivers' compilation will take place on the host during MLNX_OFED installation. Therefore, using “`mlnx_add_kernel_support.sh`” is irrelevant on Ubuntu and Debian distributions.

Example: The following command will create a MLNX_OFED_LINUX ISO image for RedHat 7.3 under the /tmp directory.

```
# ./MLNX_OFED_LINUX-x.x-x-rhel7.3-x86_64/mlnx_add_kernel_support.sh -m /tmp/MLNX_OFED_LINUX-x.x-x-rhel7.3-x86_64/ --make-tgz
Note: This program will create MLNX_OFED_LINUX TGZ for rhel7.3 under /tmp directory.
All Mellanox, OEM, OFED, or Distribution IB packages will be removed.
Do you want to continue? [y/N]:y
See log file /tmp/mlnx_ofed_iso.21642.log

Building OFED RPMs. Please wait...
Removing OFED RPMs...
Created /tmp/MLNX_OFED_LINUX-x.x-x-rhel7.3-x86_64-ext.tgz
```

- The script adds the following lines to `/etc/security/limits.conf` for the userspace components such as MPI:
 - * soft memlock unlimited
 - * hard memlock unlimited
 - These settings set the amount of memory that can be pinned by a userspace application to unlimited. If desired, tune the value unlimited to a specific amount of RAM.

For your machine to be part of the InfiniBand/VPI fabric, a Subnet Manager must be running on one of the fabric nodes. At this point, OFED for Linux has already installed the OpenSM Subnet Manager on your machine.

For the list of installation options, run:

```
./mlnxofedinstall --h
```

Installation Procedure

This section describes the installation procedure of MLNX_OFED on NVIDIA adapter cards.

- a. Log in to the installation machine as root.

b. Mount the ISO image on your machine.

```
host1# mount -o ro,loop MLNX_OFED_LINUX-<ver>-<OS label>-<CPU arch>.iso /mnt
```

c. Run the installation script.

```
/mnt/mlnxofedinstall
Logs dir: /tmp/MLNX_OFED_LINUX-x.x-x.logs
This program will install the MLNX_OFED_LINUX package on your machine.
Note that all other Mellanox, OEM, OFED, RDMA or Distribution IB packages will be removed.
Those packages are removed due to conflicts with MLNX_OFED_LINUX, do not reinstall them.
Starting MLNX_OFED_LINUX-x.x-x installation ...
.....
Installation finished successfully.
Attempting to perform Firmware update...
Querying Mellanox devices firmware ...
```

 For unattended installation, use the --force installation option while running the MLNX_OFED installation script:

```
/mnt/mlnxofedinstall --force
```

 MLNX_OFED for Ubuntu should be installed with the following flags in chroot environment:

```
./mlnxofedinstall --without-dkms --add-kernel-support --kernel <kernel
version in chroot> --without-fw-update --force
```

For example:

```
./mlnxofedinstall --without-dkms --add-kernel-support --kernel 3.13.0-85-
generic --without-fw-update --force
```

Note that the path to kernel sources (--kernel-sources) should be added if the sources are not in their default location.

 In case your machine has the latest firmware, no firmware update will occur and the installation script will print at the end of installation a message similar to the following:

Device #1:

```
Device Type: ConnectX-X
Part Number: MCXXXX-XXX
PSID: MT_<version>
PCI Device Name: 0b:00.0
Base MAC: 0000e41d2d5cf810
Versions: Current Available
FW XX.XX.XXXX
Status: Up to date
```

 In case your machine has an unsupported network adapter device, no firmware update will occur and one of the error messages below will be printed. Please contact your hardware vendor for help with firmware updates.

Error message #1:

Device #1:

```
Device Type: ConnectX-X
Part Number: MCXXXX-XXX
```

PSID: MT_<version>
 PCI Device Name: 0b:00.0
 Base MAC: 0000e41d2d5cf810
 Versions: Current Available
 FW XX.XX.XXXX
 Status: No matching image found
 Error message #2:
 The firmware for this device is not distributed inside NVIDIA driver:
 0000:01:00.0 (PSID: IBM2150110033)
 To obtain firmware for this device, please contact your HW vendor.

- d. Case A: If the installation script has performed a firmware update on your network adapter, you need to either restart the driver or reboot your system before the firmware update can take effect. Refer to the table below to find the appropriate action for your specific card.

Action \ Adapter	Driver Restart	Standard Reboot (Soft Reset)	Cold Reboot (Hard Reset)
Standard ConnectX-4/ ConnectX-4 Lx or higher	-	+	-
Adapters with Multi-Host Support	-	-	+
Socket Direct Cards	-	-	+

Case B: If the installations script has not performed a firmware upgrade on your network adapter, restart the driver by running: “/etc/init.d/openibd restart”.

- e. (InfiniBand only) Run the `hca_self_test.ofed` utility to verify whether or not the InfiniBand link is up. The utility also checks for and displays additional information such as:

- HCA firmware version
- Kernel architecture
- Driver version
- Number of active HCA ports along with their states
- Node GUID

For more details on `hca_self_test.ofed`, see the file `docs/readme_and_user_manual/hca_self_test.readme`.

After installation completion, information about the OFED installation, such as prefix, kernel version, and installation parameters can be retrieved by running the command `/etc/infiniband/info`. Most of the OFED components can be configured or reconfigured after the installation, by modifying the relevant configuration files. See the relevant chapters in this manual for details.

The list of the modules that will be loaded automatically upon boot can be found in the `/etc/infiniband/openib.conf` file.

 Installing OFED will replace the RDMA stack and remove existing 3rd party RDMA connectors.

Installation Results

Software	• Most of MLNX_OFED packages are installed under the “/usr” directory except for the following packages which are installed under the “/opt” directory: • fca and ibutils • iproute2 (rdma tool) - installed under /opt/Mellanox/iproute2/sbin/rdma • The kernel modules are installed under • /lib/modules/`uname -r`/updates on SLES and Fedora Distributions • /lib/modules/`uname -r`/extra/mlnx-ofa_kernel on RHEL and other RedHat like Distributions • /lib/modules/`uname -r`/updates/dkms/ on Ubuntu
Firmware	• The firmware of existing network adapter devices will be updated if the following two conditions are fulfilled: • The installation script is run in default mode; that is, without the option ‘--without- fw-update’ • The firmware version of the adapter device is older than the firmware version included with the OFED ISO image Note: If an adapter’s Flash was originally programmed with an Expansion ROM image, the automatic firmware update will also burn an Expansion ROM image. • In case your machine has an unsupported network adapter device, no firmware update will occur and the error message below will be printed. "The firmware for this device is not distributed inside NVIDIA driver: 0000:01:00.0 (PSID: IBM2150110033) To obtain firmware for this device, please contact your HW vendor."

Installation Logging

While installing MLNX_OFED, the install log for each selected package will be saved in a separate log file.

The path to the directory containing the log files will be displayed after running the installation script in the following format:

Example:

```
Logs dir: /tmp/MLNX_OFED_LINUX-4.4-1.0.0.0.IBM2150110033.logs
```

Driver Load Upon System Boot

Upon system boot, the NVIDIA drivers will be loaded automatically.

- To prevent the automatic load of the NVIDIA drivers upon system boot:
- a. Add the following lines to the “/etc/modprobe.d/mlnx.conf” file.

```
blacklist mlx5_core
blacklist mlx5_ib
```

- b. Set “ONBOOT=no” in the “/etc/infiniband/openib.conf” file.
 - c. If the modules exist in the initramfs file, they can automatically be loaded by the kernel. To prevent this behavior, update the initramfs using the operating systems’ standard tools.
- Note: The process of updating the initramfs will add the blacklists from step 1, and will prevent the kernel from loading the modules automatically.

mlnxofedinstall Return Codes

The table below lists the mlnxofedinstall script return codes and their meanings.

Return Code	Meaning
0	The Installation ended successfully
1	The installation failed
2	No firmware was found for the adapter device
22	Invalid parameter
28	Not enough free space
171	Not applicable to this system configuration. This can occur when the required hardware is not present on the system
172	Prerequisites are not met. For example, missing the required software installed or the hardware is not configured correctly
173	Failed to start the mst driver

Software	• Most of MLNX_OFED packages are installed under the “/usr” directory except for the following packages which are installed under the “/opt” directory: • fca and ibutils • iproute2 (rdma tool) - installed under /opt/Mellanox/iproute2/sbin/rdma • The kernel modules are installed under • /lib/modules/`uname -r`/updates on SLES and Fedora Distributions • /lib/modules/`uname -r`/extra/mlnx-ofa_kernel on RHEL and other RedHat like Distributions • /lib/modules/`uname -r`/updates/dkms/ on Ubuntu
----------	---

Firmware	- The firmware of existing network adapter devices will be updated if the following two conditions are fulfilled: - The installation script is run in default mode; that is, without the option '--without- fw-update' - The firmware version of the adapter device is older than the firmware version included with the OFED ISO image Note: If an adapter's Flash was originally programmed with an Expansion ROM image, the automatic firmware update will also burn an Expansion ROM image. - In case your machine has an unsupported network adapter device, no firmware update will occur and the error message below will be printed. "The firmware for this device is not distributed inside NVIDIA driver: 0000:01:00.0 (PSID: IBM2150110033) To obtain firmware for this device, please contact your HW vendor."
----------	---

Installation Logging

While installing MLNX_OFED, the install log for each selected package will be saved in a separate log file.

The path to the directory containing the log files will be displayed after running the installation script in the following format:

Example:

```
Logs dir: /tmp/MLNX_OFED_LINUX-4.4-1.0.0.0.IBMM2150110033.logs
```

Uninstalling MLNX_OFED

Use the script /usr/sbin/ofed_uninstall.sh to uninstall the MLNX_OFED package. The script is part of the ofed-scripts RPM.

Additional Installation Procedures

Installing MLNX_OFED Using YUM

This type of installation is applicable to RedHat/OL and Fedora operating systems.

Setting up MLNX_OFED YUM Repository

- Log into the installation machine as root.
- Mount the ISO image on your machine and copy its content to a shared location in your network.

```
# mount -o ro,loop MLNX_OFED_LINUX-<ver>-<OS label>-<CPU arch>.iso /mnt
```

c. Download and install NVIDIA's GPG-KEY:

The key can be downloaded via the following link:

<http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-Mellanox>

```
# wget http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-Mellanox
--2018-01-25 13:52:30-- http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-Mellanox
Resolving www.mellanox.com... 72.3.194.0
Connecting to www.mellanox.com|72.3.194.0|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 1354 (1.3K) [text/plain]
Saving to: ?RPM-GPG-KEY-Mellanox?

100%[=====] 1,354 --.-K/s in 0s

2018-01-25 13:52:30 (247 MB/s) - ?RPM-GPG-KEY-Mellanox? saved [1354/1354]
```

d. Install the key.

```
# sudo rpm --import RPM-GPG-KEY-Mellanox
warning: rpmts_HdrFromFdno: Header V3 DSA/SHA1 Signature, key ID 6224c050: NOKEY
Retrieving key from file:///repos/MLNX_OFED/<MLNX_OFED file>/RPM-GPG-KEY-Mellanox
Importing GPG key 0x6224C050:
  Userid: "Mellanox Technologies (Mellanox Technologies - Signing Key v2) <support@mellanox.com>"
  From : /repos/MLNX_OFED/<MLNX_OFED file>/RPM-GPG-KEY-Mellanox
Is this ok [y/N]:
```

e. Check that the key was successfully imported.

```
# rpm -q gpg-pubkey --qf '%(NAME)-%(VERSION)-%(RELEASE)\t%(SUMMARY)\n' | grep Mellanox
gpg-pubkey-a9e4b643-520791ba gpg (Mellanox Technologies <support@mellanox.com>)
```

f. Create a yum repository configuration file called "/etc/yum.repos.d/mlnx_ofed.repo" with the following content:

```
[mlnx_ofed]
name=MLNX_OFED Repository
baseurl=file:///<path to extracted MLNX_OFED package>/RPMS
enabled=1
gpgkey=file:///<path to the downloaded key RPM-GPG-KEY-Mellanox>
gpgcheck=1
```

g. Check that the repository was successfully added.

```
# yum repolist
Loaded plugins: product-id, security, subscription-manager
This system is not registered to Red Hat Subscription Management. You can use subscription-manager
to register.
repo id repo name status
mlnx_ofed MLNX_OFED Repository 108
rpmforge RHEL 6Server - RPMforge.net - dag 4,597
repolist: 8,351
```

Setting up MLNX_OFED YUM Repository Using --add-kernel-support

- Log into the installation machine as root.
- Mount the ISO image on your machine and copy its content to a shared location in your network.

```
# mount -o ro,loop MLNX_OFED_LINUX-<ver>-<OS label>-<CPU arch>.iso /mnt
```

c. Build the packages with kernel support and create the tarball.

```
# /mnt/mlnx_add_kernel_support.sh --make-tgz <optional --kmp> -k $(uname -r) -m /mnt/
Note: This program will create MLNX_OFED_LINUX TGZ for rhel7.6 under /tmp directory.
Do you want to continue?[y/N]:y
See log file /tmp/mlnx_iso.4120_logs/mlnx_ofed_iso.4120.log

Checking if all needed packages are installed...
```

```
Building MLNX_OFED_LINUX RPMs . Please wait...
Creating metadata-rpms for 3.10.0-957.21.3.el7.x86_64 ...
WARNING: If you are going to configure this package as a repository, then please note
WARNING: that it contains unsigned rpms, therefore, you need to disable the gpgcheck
WARNING: by setting 'gpgcheck=0' in the repository conf file.
Created /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-ext.tgz
```

d. Open the tarball.

```
# cd /tmp/
# tar -xvf /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-ext.tgz
```

e. Create a YUM repository configuration file called "/etc/yum.repos.d/mlnx_ofed.repo" with the following content:

```
[mlnx_ofed]
name=MLNX_OFED Repository
baseurl=file:///<path to extracted MLNX_OFED package>/RPMS
enabled=1
gpgcheck=0
```

f. Check that the repository was successfully added.

```
# yum repolist
Loaded plugins: product-id, security, subscription-manager
This system is not registered to Red Hat Subscription Management. You can use subscription-manager
to register.
repo id repo name status
mlnx_ofed MLNX_OFED Repository 108
rpmforge RHEL 6Server - RPMforge.net - dag 4,597
repolist: 8,351
```

Installing MLNX_OFED Using the YUM Tool

After setting up the YUM repository for MLNX_OFED package, perform the following:

a. View the available package groups by invoking:

```
# yum search mlnx-ofed-
mlnx-ofed-all.noarch : MLNX_OFED all installer package (with KMP support)
mlnx-ofed-all-user-only.noarch : MLNX_OFED all-user-only installer package (User Space packages
only)
mlnx-ofed-basic.noarch : MLNX_OFED basic installer package (with KMP support)
mlnx-ofed-basic-user-only.noarch : MLNX_OFED basic-user-only installer package (User Space
packages only)
mlnx-ofed-bluefield.noarch : MLNX_OFED bluefield installer package (with KMP support)
mlnx-ofed-bluefield-user-only.noarch : MLNX_OFED bluefield-user-only installer package (User Space
packages only)
mlnx-ofed-dpdk.noarch : MLNX_OFED dpdk installer package (with KMP support)
mlnx-ofed-dpdk-upstream-libs.noarch : MLNX_OFED dpdk-upstream-libs installer package (with KMP
support)
mlnx-ofed-dpdk-upstream-libs-user-only.noarch : MLNX_OFED dpdk-upstream-libs-user-only installer
package (User Space packages only)
mlnx-ofed-dpdk-user-only.noarch : MLNX_OFED dpdk-user-only installer package (User Space packages
only)
mlnx-ofed-eth-only-user-only.noarch : MLNX_OFED eth-only-user-only installer package (User Space
packages only)
mlnx-ofed-guest.noarch : MLNX_OFED guest installer package (with KMP support)
mlnx-ofed-guest-user-only.noarch : MLNX_OFED guest-user-only installer package (User Space
packages only)
mlnx-ofed-hpc.noarch : MLNX_OFED hpc installer package (with KMP support)
mlnx-ofed-hpc-user-only.noarch : MLNX_OFED hpc-user-only installer package (User Space packages
only)
mlnx-ofed-hypervisor.noarch : MLNX_OFED hypervisor installer package (with KMP support)
mlnx-ofed-hypervisor-user-only.noarch : MLNX_OFED hypervisor-user-only installer package (User
Space packages only)
mlnx-ofed-kernel-only.noarch : MLNX_OFED kernel-only installer package (with KMP support)
mlnx-ofed-vma.noarch : MLNX_OFED vma installer package (with KMP support)
mlnx-ofed-vma-eth.noarch : MLNX_OFED vma-eth installer package (with KMP support)
mlnx-ofed-vma-eth-user-only.noarch : MLNX_OFED vma-eth-user-only installer package (User Space
packages only)
mlnx-ofed-vma-user-only.noarch : MLNX_OFED vma-user-only installer package (User Space packages
only)
mlnx-ofed-vma-vpi.noarch : MLNX_OFED vma-vpi installer package (with KMP support)
mlnx-ofed-vma-vpi-user-only.noarch : MLNX_OFED vma-vpi-user-only installer package (User Space
packages only)
```

where:

mlnx-ofed-all	Installs all available packages in MLNX_OFED
mlnx-ofed-basic	Installs basic packages required for running NVIDIA cards
mlnx-ofed-guest	Installs packages required by guest OS
mlnx-ofed-hpc	Installs packages required for HPC
mlnx-ofed-hypervisor	Installs packages required by hypervisor OS
mlnx-ofed-vma	Installs packages required by VMA
mlnx-ofed-vma-eth	Installs packages required by VMA to work over Ethernet
mlnx-ofed-vma-vpi	Installs packages required by VMA to support VPI
bluefield	Installs packages required for BlueField
dpdk	Installs packages required for DPDK
dpdk-upstream-libs	Installs packages required for DPDK using RDMA-Core
kernel-only	Installs packages required for a non-default kernel

Note: MLNX_OFED provides kernel module RPM packages with KMP support for RHEL and SLES. For other operating systems, kernel module RPM packages are provided only for the operating system's default kernel. In this case, the group RPM packages have the supported kernel version in their package's name.

Example:

```
mlnx-ofed-all-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED all installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
mlnx-ofed-basic-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED basic installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
mlnx-ofed-guest-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED guest installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
mlnx-ofed-hpc-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED hpc installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
mlnx-ofed-hypervisor-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED hypervisor installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
mlnx-ofed-vma-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED vma installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
mlnx-ofed-vma-eth-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED vma-eth installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
mlnx-ofed-vma-vpi-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED vma-vpi installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
mlnx-ofed-hypervisor-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED hypervisor installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
mlnx-ofed-vma-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED vma installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
mlnx-ofed-vma-eth-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED vma-eth installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
mlnx-ofed-vma-vpi-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED vma-vpi installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)
```

When using an operating system different than RHEL or SLES, or you have installed a kernel that is not supported by default in MLNX_OFED, you can use the

mlnx_add_kernel_support.sh script to build MLNX_OFED for your kernel.

The script will automatically build the matching group RPM packages for your kernel so that you can still install MLNX_OFED via yum.

Please note that the resulting MLNX_OFED repository will contain unsigned RPMs, therefore, you should set 'gpgcheck=0' in the repository configuration file.

- b. Install the desired group.

```
# yum install mlnx-ofed-all
Loaded plugins: langpacks, product-id, subscription-manager
```

```

Resolving Dependencies
--> Running transaction check
---> Package mlnx-ofed-all.noarch 0:3.1-0.1.2 will be installed
--> Processing Dependency: kmod-iser = 1.0-OFED.3.1.0.1.2.1.g832a737.rhel7u1 for package: mlnx-
ofed-all-3.1-0.1.2.noarch
.....
.....
qperf.x86_64 0:0.4.9-9
rds-devel.x86_64 0:2.0.7-1.12
rds-tools.x86_64 0:2.0.7-1.12
sdpnetstat.x86_64 0:1.60-26
srptools.x86_64 0:1.0.2-12

Complete!

```

⚠ Installing MLNX_OFED using the “YUM” tool does not automatically update the firmware.

To update the firmware to the version included in MLNX_OFED package, run:

```
# yum install mlnx-fw-updater
```

Installing MLNX_OFED Using apt-get

This type of installation is applicable to Debian and Ubuntu operating systems.

Setting up MLNX_OFED apt-get Repository

- Log into the installation machine as root.
- Extract the MLNX_OFED package on a shared location in your network.
It can be downloaded from <https://www.nvidia.com/en-us/networking/> → Products → Software → InfiniBand Drivers.
- Create an apt-get repository configuration file called `"/etc/apt/sources.list.d/mlnx_ofed.list"` with the following content:

```
deb file: /<path to extracted MLNX_OFED package>/DEBS ./
```

- Download and install NVIDIA's Technologies GPG-KEY.

```
# wget -qO - http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-Mellanox | sudo apt-key add -
```

- Verify that the key was successfully imported.

```
# apt-key list
pub 1024D/A9E4B643 2013-08-11
uid Mellanox Technologies <support@mellanox.com>
sub 1024g/09FCC269 2013-08-11
```

- Update the apt-get cache.

```
# sudo apt-get update
```

Setting up MLNX_OFED apt-get Repository Using --add-key-support

- Log into the installation machine as root.
- Mount the ISO image on your machine and copy its content to a shared location in your network.

```
# mount -o ro,loop MLNX_OFED_LINUX-<ver>-<OS label>-<CPU arch>.iso /mnt
```

c. Build the packages with kernel support and create the tarball.

```
# /mnt/mlnx_add_kernel_support.sh --make-tgz <optional --kmp> -k $(uname -r) -m /mnt/
Note: This program will create MLNX_OFED_LINUX TGZ for rhel7.6 under /tmp directory.
Do you want to continue?[y/N]:y
See log file /tmp/mlnx_iso.4120_logs/mlnx_ofed_iso.4120.log

Checking if all needed packages are installed...
Building MLNX_OFED_LINUX RPMs . Please wait...
Creating metadata-rpms for 3.10.0-957.21.3.el7.x86_64 ...
WARNING: If you are going to configure this package as a repository, then please note
WARNING: that it contains unsigned rpms, therefore, you need to disable the gpgcheck
WARNING: by setting 'gpgcheck=0' in the repository conf file.
Created /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-ext.tgz
```

d. Open the tarball.

```
# cd /tmp/
# tar -xvf /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-ext.tgz
```

e. Create an apt-get repository configuration file called "/etc/apt/sources.list.d/mlnx_ofed.list" with the following content:

```
deb [trusted=yes] file:/<path to extracted MLNX_OFED package>/DEBS ./
```

f. Update the apt-get cache.

```
# sudo apt-get update
```

Installing MLNX_OFED Using the apt-get Tool

After setting up the apt-get repository for MLNX_OFED package, perform the following:

a. View the available package groups by invoking:

```
# apt-cache search mlnx-ofed-
apt-cache search mlnx-ofed .....
knem-dkms - DKMS support for mlnx-ofed kernel modules
mlnx-ofed-kernel-dkms - DKMS support for mlnx-ofed kernel modules
mlnx-ofed-kernel-utils - Userspace tools to restart and tune mlnx-ofed kernel modules
mlnx-ofed-vma-vpi - MLNX_OFED vma-vpi installer package (with DKMS support)
mlnx-ofed-kernel-only - MLNX_OFED kernel-only installer package (with DKMS support)
mlnx-ofed-bluefield - MLNX_OFED bluefield installer package (with DKMS support)
mlnx-ofed-hpc-user-only - MLNX_OFED hpc-user-only installer package (User Space packages only)
mlnx-ofed-dpdk-user-only - MLNX_OFED dpdk-user-only installer package (User Space packages only)
mlnx-ofed-all-exact - MLNX_OFED all installer package (with DKMS support) (exact)
mlnx-ofed-all - MLNX_OFED all installer package (with DKMS support)
mlnx-ofed-vma-vpi-user-only - MLNX_OFED vma-vpi-user-only installer package (User Space packages only)
mlnx-ofed-eth-only-user-only - MLNX_OFED eth-only-user-only installer package (User Space packages only)
mlnx-ofed-vma-user-only - MLNX_OFED vma-user-only installer package (User Space packages only)
mlnx-ofed-hpc - MLNX_OFED hpc installer package (with DKMS support)
mlnx-ofed-bluefield-user-only - MLNX_OFED bluefield-user-only installer package (User Space packages only)
mlnx-ofed-dpdk - MLNX_OFED dpdk installer package (with DKMS support)
mlnx-ofed-vma-eth-user-only - MLNX_OFED vma-eth-user-only installer package (User Space packages only)
mlnx-ofed-all-user-only - MLNX_OFED all-user-only installer package (User Space packages only)
mlnx-ofed-vma-eth - MLNX_OFED vma-eth installer package (with DKMS support)
mlnx-ofed-vma - MLNX_OFED vma installer package (with DKMS support)
mlnx-ofed-dpdk-upstream-libs-user-only - MLNX_OFED dpdk-upstream-libs-user-only installer package (User Space packages only)
mlnx-ofed-basic-user-only - MLNX_OFED basic-user-only installer package (User Space packages only)
mlnx-ofed-basic-exact - MLNX_OFED basic installer package (with DKMS support) (exact)
mlnx-ofed-basic - MLNX_OFED basic installer package (with DKMS support)
mlnx-ofed-dpdk-upstream-libs - MLNX_OFED dpdk-upstream-libs installer package (with DKMS support)
```

where:

mlnx-ofed-all	MLNX_OFED all installer package
---------------	---------------------------------

mlnx-ofed-basic	MLNX_OFED basic installer package
mlnx-ofed-vma	MLNX_OFED vma installer package
mlnx-ofed-hpc	MLNX_OFED HPC installer package
mlnx-ofed-vma-eth	MLNX_OFED vma-eth installer package
mlnx-ofed-vma-vpi	MLNX_OFED vma-vpi installer package
knem-dkms	MLNX_OFED DKMS support for mlnx-ofed kernel modules
kernel-dkms	MLNX_OFED kernel-dkms installer package
kernel-only	MLNX_OFED kernel-only installer package
bluefield	MLNX_OFED bluefield installer package
mlnx-ofed-all-exact	MLNX_OFED mlnx-ofed-all-exact installer package
dpdk	MLNX_OFED dpdk installer package
mlnx-ofed-basic-exact	MLNX_OFED mlnx-ofed-basic-exact installer package
dpdk-upstream-libs	MLNX_OFED dpdk-upstream-libs installer package

b. Install the desired group.

```
apt-get install '<group name>'
```

Example:

```
apt-get install mlnx-ofed-all
```



Installing MLNX_OFED using the “apt-get” tool does not automatically update the firmware.

To update the firmware to the version included in MLNX_OFED package, run:

```
# apt-get install mlnx-fw-updater
```

Performance Tuning

Depending on the application of the user's system, it may be necessary to modify the default configuration of network adapters based on the ConnectX® adapters. In case that tuning is required, please refer to the [Performance Tuning Guide for NVIDIA Network Adapters](#).

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.

Software Requirements

Description	Package
Windows Server 2022	MLNX_WinOF2-<version>_All_x64.exe
Windows Server 2019	
Windows Server 2016	
Windows Server 2012 R2	
Windows 11 Client (64 bit only)	
Windows 10 Client (64 bit only)	
Windows 8.1 Client (64 bit only)	

Note: The Operating System listed above must run with administrator privileges.

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”

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.

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 installation 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:
- Performance 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.



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*v* [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 installation 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*v* 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"
```

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.

VMware Driver Installation

This section describes VMware Driver Installation.

Hardware and Software Requirements

Requirement	Description
Platforms	A server platform with an adapter card based on NVIDIA devices: ConnectX®-7 (InfiniBand/Ethernet) (firmware: fw-ConnectX7)
Operating System	For the complete list of VMware supported operating systems, refer to VMware ESXi async Drivers
Installer Privileges	The installation requires administrator privileges on the target machine.

Installing NATIVE ESXi Driver for VMware vSphere

 Please uninstall all previous driver packages prior to installing the new version.

To install the driver:

1. Log into the ESXi server with root permissions.
2. Install the driver.

```
#> esxcli software vib install -d <path>/<bundle_file>
```

Example:

```
#> esxcli software vib install -d /tmp/MLNX-NATIVE-ESX-ConnectX-4-5_4.16.8.8-10EM-650.0.0.4240417.zipesxcli
```

3. Reboot the machine.
4. Verify the driver was installed successfully.

```
esxcli software vib list | grep nmlx  
nmlx5-core      4.16.8.8-10EM.650.0.0.4240417    MEL    PartnerSupported 2017-01-31
```

⚠ After the installation process, all kernel modules are loaded automatically upon boot.

Removing Earlier NVIDIA Drivers

⚠ Please unload the previously installed drivers before removing them.

To remove all the drivers:

1. Log into the ESXi server with root permissions.
2. List all the existing NATIVE ESXi driver modules. (See Step 4 in [Installing NATIVE ESXi Driver for VMware vSphere.](#))
3. Remove each module:

```
#> esxcli software vib remove -n nmlx5-rdma
#> esxcli software vib remove -n nmlx5-core
```

⚠ To remove the modules, you must run the command in the same order as shown in the example above.

4. Reboot the server.

Firmware Programming

1. Download the VMware bootable binary images v4.6.0 from the [Firmware Tools \(MFT\) site](#).
 - a. ESXi 6.5 File: mft-4.6.0.48-10EM-650.0.0.4598673.x86_64.vib
 - b. MD5SUM: 0804cffe30913a7b4017445a0f0adbe1
2. Install the image according to the steps described in the [MFT User Manual](#).

⚠ The following procedure requires custom boot image downloading, mounting and booting from a USB device.

Updating Adapter Firmware

Each adapter card 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 Mellanox adapters and indicate which adapters 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-7
Part Number:      MCX75310AAS-HEAT
Description:      NVIDIA ConnectX-7 adapter card, 200Gb/s NDR200 IB, Single-port OSFP, PCIe 5.0 x16, Secure boot, No
Crypto, Tall 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-7
Part Number:      MCX75310AAS-HEAT
Description:      NVIDIA ConnectX-7 adapter card, 200Gb/s NDR200 IB, Single-port OSFP, PCIe 5.0 x16, Secure boot, No
Crypto, Tall Bracket
PCI Device Name:  0b:00.0
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
```

Setting High-Speed-Port Link Type

 This section applies to ConnectX-7 cards supporting both Ethernet and InfiniBand protocols - see the relevant OPNs in the following table.

The following table lists the ConnectX-7 cards supporting both Ethernet and InfiniBand protocols, the supported speeds and the default networking port link type.

OPN	Data Transmission Rate	Default Protocol and Rate
MCX75310AAS-HEAT	NDR200 / 200GbE	InfiniBand NDR200
MCX75310AAS-NEAT	NDR / 400GbE	InfiniBand NDR
MCX75310AAC-NEAT	NDR / 400GbE	InfiniBand NDR
MCX755106AS-HEAT	NDR200 / 200GbE	Ethernet 200GbE
MCX755106AC-HEAT	NDR200 / 200GbE	Ethernet 200GbE
MCX715105AS-WEAT	NDR / 400GbE	Ethernet 400GbE

To configure the networking high-speed ports mode, you can either use the [mlxconfig](#) or the [UEFI](#) tools.

UEFI can configure the adapter card device before the operating system is up, while mlxconfig configures the card once the operating system is up. According to your preference, use one of the below tools:

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](#).

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](#).

Troubleshooting

General Troubleshooting

Server unable to find the adapter	• Ensure that the adapter is placed correctly • Make sure the adapter slot and the adapter are compatible Install the adapter in a different PCI Express slot • Use the drivers that came with the adapter or download the latest • Make sure your motherboard has the latest BIOS • Try to reboot the server
The adapter no longer works	• Reseat the adapter 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
Adapters stopped working after installing another adapter	• Try removing and re-installing all adapters • 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 adapter port are compatible
Link light is on, but with no communication established	• Check that the latest driver is loaded • Check that both the adapter and its link are set to the same speed and duplex settings

Linux Troubleshooting

Environment Information	<pre>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</pre>
Card Detection	<pre>lspci grep -i Mellanox</pre>

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

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

Specifications

⚠ The ConnectX-7 adapter card is designed and validated for operation in data-center servers and other large environments that guarantee proper power supply and airflow conditions. The adapter card 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 adapter card's functionality and potentially damage it. Failure to meet the environmental requirements listed in this user manual may void the warranty.

⚠ Please make sure to install the ConnectX-7 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.

MCX75310AAC-NEAT / MCX75310AAS-NEAT Specifications

Physical	Adapter Card Size: PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)		
Interfaces	See Supported Interfaces		
	PCI Express Gen 4.0/5.0: SERDES @ 16/32GT/s, x16 lanes (Gen 3.0 compatible)		
	Networking Port: Single OSFP InfiniBand and Ethernet		
Data Rate	InfiniBand (Default Mode)	NDR/NDR200/HDR/HDR100/EDR/FDR/SDR	
	Ethernet	400/200/100/50/40/10/1 Gb/s Ethernet	
Protocol Support	InfiniBand: IBTA v1.5 ^a Auto-Negotiation: 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), 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-4 C2M, 200GBASE-CR4, 100GAUI-2 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	MCX75310AAC-NEAT	Secure Boot Enabled, Crypto Enabled	
	MCX75310AAS-NEAT	Secure Boot Enabled, Crypto Disabled	
Electrical and Thermal Specifications	Voltage: 12V, 3.3VAUX Maximum current: 100mA		
	Maximum power available through OSFP port: 17W (Not thermally supported)		
	Electrical and thermal specifications are provided in <i>NVIDIA ConnectX-7 Electrical and Thermal Specifications</i> document, which is available at NVOnline following login.		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	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-7 adapters 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. Typical power for ATIS traffic load.

c. For both operational and non-operational states.

MCX75310AAS-HEAT Specifications

Physical	Adapter Card Size: PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)		
Interfaces	See Supported Interfaces		
	PCI Express Interface: Gen 4.0/5.0: SERDES @ 16/32GT/s, x16 lanes (Gen 3.0 compatible)		
	Networking Port: Single OSFP InfiniBand and Ethernet		
Data Rate	InfiniBand (Default Mode)	NDR200/HDR/HDR100/EDR/FDR/SDR	
	Ethernet	200/100/50/40/10/1 Gb/s Ethernet	
Protocol Support	InfiniBand: IBTA v1.5 ^a Auto-Negotiation: NDR200 (2 lanes x 100Gb/s per lane) port, HDR (50Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane).		
	Ethernet: 200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 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	MCX75310AAS-HEAT	Secure Boot Enabled, Crypto Disabled	
Electrical and Thermal Specifications	Voltage: 12V, 3.3VAUX Maximum current: 100mA		
	Maximum power available through OSFP port: 17W (Not thermally supported)		
	Electrical and thermal specifications are provided in NVIDIA ConnectX-7 Electrical and Thermal Specifications document, which is available at NVOnline following login.		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	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-7 adapters 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. Typical power for ATIS traffic load.

c. For both operational and non-operational states.

MCX755106AC-HEAT / MCX755106AS-HEAT Specifications

 The Socket-Direct ready cards kit does not include the PCIe passive auxiliary connection card and two Cabline SA-II Plus harnesses. For more information, please refer to [PCIe Auxiliary Card Kit](#).

Physical	Adapter Card Size	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)
	Auxiliary PCIe Connection Card Size	5.09 in. x 2.32 in. (129.30mm x 59.00mm) Two Cabline CA-II Plus harnesses (white and black)
Interface s	See Supported Interfaces	
	PCI Express Interface	Gen 5.0/4.0: SERDES @ 16/32GT/s, x16 lanes (4.0 and 3.0 compatible)
		Optional: Additional PCIe x16 Gen 4.0 @ SERDES 18GT/s through the PCIe auxiliary passive card and Cabline SA-II Plus harnesses
	Networking Ports	Dual QSFP112 InfiniBand and Ethernet
Data Rate	InfiniBand	NDR200/HDR/HDR100/EDR/FDR/SDR
	Ethernet (Default Mode)	200/100/50/25/10 Gb/s
Protocol Support	InfiniBand: IBTA v1.5 ^a Auto-Negotiation: NDR200 (2 lanes x 100Gb/s per lane) port, HDR (50Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane)	
	Ethernet Protocols	200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 C2M, 100GBASE-CR4, 100GBASE-CR, 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	MCX755106AC-HEAT	Secure Boot Enabled, Crypto Enabled

	MCX755106AS -HEAT	Secure Boot Enabled, Crypto Disabled	
Electrical and Thermal Specifications	Voltage: 12V, 3.3VAUX Maximum current: 100mA		
	Maximum power available through QSFP112 port: TBD		
	Electrical and Thermal Specifications: TBD		
Environmental	Temperature	Operational	0 °C to 55 °C
		Non-operational	-40 °C to 70 °C
	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-7 adapters 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. Typical power for ATIS traffic load.

c. For both operational and non-operational states.

MCX715105AS-WEAT Specifications

 The Socket-Direct ready cards kit does not include the PCIe passive auxiliary connection card and two Cabline SA-II Plus harnesses. For more information, please refer to [PCIe Auxiliary Card Kit](#).

Physical	Adapter Card Size	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)
	Auxiliary PCIe Connection Card Size	5.09 in. x 2.32 in. (129.30mm x 59.00mm) Two Cabline CA-II Plus harnesses (white and black)
Interfaces	See Supported Interfaces	
	PCI Express Interface	Gen 5.0/4.0: SERDES @ 16/32GT/s, x16 lanes (4.0 and 3.0 compatible)
		Optional: Additional PCIe x16 Gen 4.0 @ SERDES 18GT/s through the PCIe auxiliary passive card and Cabline SA-II Plus harnesses
	Networking Ports	Single QSFP112 InfiniBand and Ethernet

Data Rate	InfiniBand	NDR/NDR200/HDR/HDR100/EDR/FDR/SDR	
	Ethernet (Default Mode)	400/200/100/50/25/10 Gb/s	
Protocol Support	InfiniBand: IBTA v1.5 ^a Auto-Negotiation: 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), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane)		
	Ethernet Protocols	400GAUI-4 C2M, 400GBASE-CR4, 200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 C2M, 100GBASE-CR4, 100GBASE-CR, 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 Disabled		
Electrical and Thermal Specifications	Voltage: 12V, 3.3VAUX Maximum current: 100mA		
	Maximum power available through QSFP112 port: TBD		
	Power and airflow specifications: TBD		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	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-7 adapters 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. Typical power for ATIS traffic load.

c. For both operational and non-operational states.

MCX75510AAS-HEAT Specifications

 The Socket-Direct ready cards kit does not include the PCIe passive auxiliary connection card and two Cabline SA-II Plus harnesses. For more information, please refer to [PCIe Auxiliary Card Kit](#).

Physical	Adapter Card Size	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	
	Auxiliary PCIe Connection Card Size	5.09 in. x 2.32 in. (129.30mm x 59.00mm) Two Cabline CA-II Plus harnesses (white and black)	
Interfaces	See Supported Interfaces		
	PCI Express Interface	Gen 5.0/4.0: SERDES @ 16/32GT/s, x16 lanes (4.0 and 3.0 compatible)	
		Optional: Additional PCIe x16 Gen 4.0 @ SERDES 16GT/s through the PCIe auxiliary passive card and Cabline SA-II Plus harnesses	
	Networking Ports	Single OSFP InfiniBand	
Data Rate	InfiniBand	NDR200/HDR/HDR100/EDR/FDR/SDR	
Protocol Support	InfiniBand: IBTA v1.5 ^a Auto-Negotiation: NDR200 (2 lanes x 100Gb/s per lane) port, HDR (50Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane)		
	Capabilities	Secure Boot Enabled, Crypto Disabled	
Electrical and Thermal Specifications	Voltage: 12V, 3.3VAUX Maximum current: 100mA		
	Maximum power available through OSFP port: TBD		
	Power and airflow specifications: TBD		
Environmental	Temperature	Operational	0 °C to 55 °C
		Non-operational	-40 °C to 70 °C
	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-7 adapters 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. Typical power for ATIS traffic load.

c. For both operational and non-operational states.

MCX75510AAS-NEAT Specifications

 The Socket-Direct ready cards kit does not include the PCIe passive auxiliary connection card and two Cabline SA-II Plus harnesses. For more information, please refer to [PCIe Auxiliary Card Kit](#).

Physical	Adapter Card Size	PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)	
	Auxiliary PCIe Connection Card Size	5.09 in. x 2.32 in. (129.30mm x 59.00mm) Two Cabline CA-II Plus harnesses (white and black)	
Interfaces	See Supported Interfaces		
	PCI Express Interface	Gen 5.0/4.0: SERDES @ 16/32GT/s, x16 lanes (4.0 and 3.0 compatible) Optional: Additional PCIe x16 Gen 4.0 @ SERDES 16GT/s through the PCIe auxiliary passive card and Cabline SA-II Plus harnesses	
		Networking Ports	Single OSFP InfiniBand
Data Rate	InfiniBand	NDR/NDR200/HDR/HDR100/EDR/FDR/SDR	
Protocol Support	InfiniBand: IBTA v1.5 ^a Auto-Negotiation: 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), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane)		
Capabilities	Secure Boot Enabled, Crypto Disabled		
Electrical and Thermal Specifications	Voltage: 12V, 3.3VAUX Maximum current: 100mA		
	Maximum power available through OSFP port: TBD		
	Electrical and Thermal Specifications: TBD		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	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-7 adapters 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. Typical power for ATIS traffic load.

c. For both operational and non-operational states.

MCX713106AC-CEAT and MCX713106AS-CEAT Specifications

Physical	Adapter Card Size: PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)		
Interfaces	See Supported Interfaces		
	PCI Express Gen 4.0/5.0: SERDES @ 16/32GT/s, x16 lanes (4.0 and 3.0 compatible)		
	Networking Ports: Dual-port QSFP112 Ethernet (copper and optical)		
Capabilities	MCX713106 AC-CEAT	Secure Boot Enabled, Crypto Enabled	
	MCX713106 AS-CEAT	Secure Boot Enabled, Crypto Disabled	
Protocol Support	Data Rate	Ethernet	200/100/50/40/25/10/1GbE
	Ethernet Protocols: 200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 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		
Electrical and Thermal Specifications	Voltage: 12V, 3.3VAUX Maximum current: 100mA		
	Maximum power available through QSFP112 port: 11W per port (Not thermally supported), 2.5W per port (Thermally Supported)		
	Electrical and thermal specifications are provided in <i>NVIDIA ConnectX-7 Electrical and Thermal Specifications</i> document, which is available at NVOnline following login.		
Environmental	Temperature	Operational	0 °C to 55 °C
		Non-operational	-40 °C to 70 °C
	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		

MCX713106AC-VEAT and MCX713106AS-VEAT Specifications

Physical	Adapter Card Size: PCIe Half Height, Half Length 2.71 in. x 6.6 in. (68.90mm x 167.65 mm)		
Interfaces	See Supported Interfaces		
	Gen5.0: SERDES @ 16.0GT/s/32GT/s, x16 lanes (4.0, 3.0, 2.0 and 1.1 compatible)		
	Networking Ports: Dual-port QSFP112 Ethernet (copper and optical)		
Protocol Support	Data Rate	Ethernet	200/100/50/40/25/10/1 GbE
	Ethernet Protocols: 200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 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	MCX713106A C-CEAT	Secure Boot Enabled, Crypto Enabled	
	MCX713106A S-CEAT	Secure Boot Enabled, Crypto Disabled	
Electrical and Thermal Specifications	Voltage: 12V, 3.3VAUX Maximum current: 100mA		
	Maximum power available through QSFP112 port: 11W per port (Not thermally supported), 5.1W per port (Thermally supported)		
	Electrical and thermal specifications are provided in NVIDIA ConnectX-7 Electrical and Thermal Specifications document, which is available at NVOnline following login.		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	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		

MCX713104AC-ADAT and MCX713104AS-ADAT Specifications

⚠ The physical board dimensions are compliant with PCI Express Card Electromechanical Specification Revision 4.0 except for minor differences with the edge finger alignment, bracket mounting scheme, and low-profile bracket opening. These differ slightly from the PCI CEM specification due to the mechanical constraint of the single quad-port SFP56 cage. It is recommended to use the 3D stp file. Please contact your NVIDIA sales representative to get the mechanical simulation.






Physical	PCIe Half Height, Half Length 2.71 in. x 5.64 in. (68.90mm x 143.50 mm)		
Interfaces	See Supported Interfaces		
	PCI Express Gen 4.0: SERDES @ 16GT/s, x16 lanes (4.0 and 3.0 compatible)		
	Networking Port: Quad-port SFP56 Ethernet (copper and optical)		
Protocol Support	Data Rate	Ethernet	50/25GbE
	Ethernet Protocols: 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	MCX713104AC-ADAT: Secure Boot Enabled, Crypto Enabled MCX713104AS-ADAT: Secure Boot Enabled, Crypto Disabled		
Electrical and Thermal Specifications	Voltage: 12V, 3.3VAUX Maximum current: 100mA		
	Maximum power available through Quad-port SFP56: 1.5W per port (Thermally optimised)		
	Power and airflow specifications are provided in <i>NVIDIA ConnectX-7 Electrical and Thermal Specifications</i> document, which is available at NVOnline following login.		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	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		

MCX713114TC-GEAT and MCX713114GC-GEAT Specifications

Physical			
Interfaces	See Supported Interfaces		
	PCI Express Gen 4.0: SERDES @ 16GT/s, x16 lanes (4.0 and 3.0 compatible)		
	Networking Port: Quad-port SFP56 Ethernet (copper and optical)		
Protocol Support	Data Rate	Ethernet	50/25 GbE
	Ethernet Protocols: 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	MCX713114T C-GEAT	Enhanced-SyncE & PTP, PPS In/Out, Secure Boot, Crypto Enabled	
	MCX713114G C-GEAT	Enhanced-SyncE & PTP Grand Master support and GNSS/PPS Out, Secure Boot, Crypto Enabled	
Electrical and Thermal Specifications	Voltage: 12V, 3.3VAUX Maximum current: 100mA		
	Maximum power available through Quad-port SFP56: 1.5W per port, 6W in total		
	Electrical and Thermal Specifications: TBD		
Environmental	Temperature	Operational	0 °C to 55 °C
		Non-operational	-40 °C to 70 °C
	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		

Cards Mechanical Drawings and Dimensions

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








Dual-port x16 QSFP112 Adapter Card	Dual-port x16 QSFP112 Socket-Direct Ready Adapter Card
	
Single-port x16 QSFP112 Adapter Card	HHHL Quad-port SFP56 Adapter Card
	
HHHL Quad-port SFP56 with PPS IN/OUT Adapter Card	
TBD	
Single-port x16 OSFP Adapter Card	Auxiliary PCIe Connection Card
	

Brackets Mechanical Drawings and Dimensions

- ⚠ Applicable to MCX713114TC-GEAT, MCX713114GC-GEAT, MCX713104AC-ADAT and MCX713104AS-ADAT only:
The physical board dimensions are compliant with PCI Express Card Electromechanical Specification Revision 4.0 except for minor differences with the edge finger alignment, bracket mounting scheme, and low-profile bracket opening. These differ slightly from the

PCI CEM specification due to the mechanical constraint of the single quad-port SFP56 cage. It is recommended to use the associated 3D step file. Please contact your NVIDIA sales representative to get the mechanical simulation.

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

Dual-port QSFP112 Tall Bracket	Dual-port QSFP112 Short Bracket
	
Single-port QSFP112 Tall Bracket	Single-port QSFP112 Short Bracket
TBD	TBD
Single-port OSFP Tall Bracket	Single-port OSFP Adapter Card Short Bracket
	
Quad-port SFP56 Tall Bracket	Quad-port SFP56 Short Bracket
	
Quad-port SFP56 with PPS IN/OUT Tall Bracket	Quad-port SFP56 with PPS IN/OUT Short Bracket
Auxiliary PCIe Connection Card Tall Bracket	Auxiliary PCIe Connection Card Short Bracket
	

Monitoring

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:

- **Warning - 105°C:** On managed systems only: When the device crosses the 105°C threshold, a Warning Threshold message is issued by the management SW, indicating to system administration that the card has crossed the warning threshold. Note that this temperature threshold does not require nor lead to any action by hardware (such as adapter card shutdown).
- **Critical - 115°C:** When the device crosses this temperature, the firmware automatically shuts down the device.
- **Emergency - 130°C:** If the firmware fails to shut down the device upon crossing the critical threshold, the device automatically shuts down upon crossing the emergency (130°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.

Adapter Card Heatsink

The heatsink is attached to the ConnectX-7 IC to dissipate the heat from the ConnectX-7 IC. It is attached by spring-loaded screws.

ConnectX-7 IC has a thermal shutdown safety mechanism that automatically shuts down the ConnectX-7 card 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.

Finding the GUID/MAC on the Adapter Card

Each adapter card has a different identifier printed on the label: serial number and the card MAC for the Ethernet protocol and the card GUID for the InfiniBand protocol. InfiniBand/Ethernet cards have both a GUID and a MAC (derived from the GUID). IB only cards have GUID for the InfiniBand protocol.

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

Board Label Example



PCIe Auxiliary Card Kit

 This section applies to the following adapter cards when used as Socket Direct cards in dual-socket servers.

- MCX755106AS-HEAT
- MCX755106AC-HEAT
- MCX75510AAS-NEAT
- MCX75510AAS-HEAT
- MCX715105AS-WEAT

Socket Direct network cards, which cost-effectively integrate a single network adapter silicon on a primary board, and an auxiliary PCIe connection card and Slim Line SAS harnesses connecting the two. Socket Direct enables direct access from each CPU to the network through its dedicated PCIe interface as the card's 32-lane PCIe bus is split into two 16-lane buses, with one bus accessible through a PCIe x16 edge connector and the other bus through an x16 Auxiliary PCIe Connection card. The two cards should be installed into two PCIe x16 slots and connected using two Cabline SA-II Plus harnesses.

The PCIe auxiliary kit can be purchased separately to operate in a dual-socket server. The below table lists the available PCIe auxiliary kit ordering part numbers, depending on the desired length of the Cabline SA-II Plus harnesses.

Ordering Part Number	Passive Auxiliary Connection	Cabline SA-II Plus Harnesses Length
MTMK9100-T15	PCIe Gen 4.0/5.0 x16 connection card	2x 150mm harnesses
MTMK9100-T25	PCIe Gen 4.0 x16 connection card	2x 250mm harnesses
MTMK9100-T35	PCIe Gen 4.0 x16 connection card	2x 350mm harnesses

The two Cabline SA-II Plus harnesses in the PCIe auxiliary kit have different routings. To distinguish between these two harnesses, one black harness is marked with a "WHITE" label while the harness is marked with a "BLACK" label.

The Cabline harness marked with the "WHITE" label should be connected to the connector on the networking card and PCIe Auxiliary card engraved with "White Cable" while the one marked with the "BLACK" label should be connected to the connector on the networking card and the PCIe Auxiliary card engraved with "Black Cable". The Cabline SA-II Plus harness mates with two 60-pin IPEX connectors (P/N 20790-060E-01), on both sides. The black Cabline SA-II Plus harness mates with the connector on the component side (top side) of the network card, while the White Cabline SA-II Plus harnesses mates with the pint side (bottom side) of the main network card. For hardware installation, please refer to [ConnectX-7 Socket Direct \(2x PCIe x16\) Installation Instructions](#).

PCIe Auxiliary Card Package Contents

Category	Qty	Item
Cards	1	MTMK9100-T15: PCIe x16 Gen 5.0/4.0 Auxiliary Connection Card MTMK9100-T25 and MTMK9100-T35: PCIe x16 Gen 4.0 Auxiliary Connection Card
	1	Cabline CA-II Plus harness (white) - Length according to kit OPN (15, 25 or 35cm)
Accessories	1	Cabline CA-II Plus harness (black) - Length according to kit OPN (15, 25 or 35cm)
	2	Retention Clip for Cabeline harness (shipped assembled on the harnesses - optional)
	1	PCIe Auxiliary card short bracket
	1	PCIe Auxiliary card tall bracket (shipped assembled on the Auxiliary card)

Channel Insertion Loss

To be updated in a future version of this document.

Cabline CA-II Plus Harness Pinouts

Component Side

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
1	GND	GND BAR			1
2	PCIE_REFCLK1_P	Micro coax	Primary PCIe clock from the motherboard to the Network Adapter Main card, to be used for the x16 Cabline harness PCIe interface. This clock must meet all the PCIe SIG spec requirements. It should be driven from the motherboard side.	38	2

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
3	PCIE_REFCLK1_N	Micro coax	Primary PCIe clock from the motherboard to the Network Adapter Main card, to be used for the x16 Cabline harness PCIe interface. This clock must meet all the PCIe SIG spec requirements. It should be driven from the motherboard side.	38	3
4	GND	GND BAR			4
5	PCIE_CPU_CX_15N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	5
6	PCIE_CPU_CX_15P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	6
7	GND	GND BAR			7
8	PCIE_CPU_CX_14N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	8
9	PCIE_CPU_CX_14P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	9
10	GND	GND BAR			10
11	PCIE_CPU_CX_13N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	11

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
12	PCIE_CPU_CX_13P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	12
13	GND	GND BAR			13
14	PCIE_CPU_CX_12N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	14
15	PCIE_CPU_CX_12P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	15
16	GND	GND BAR			16
17	PCIE_CPU_CX_11N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	17
18	PCIE_CPU_CX_11P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	18
19	GND	GND BAR			19
20	PCIE_CPU_CX_10N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	20
21	PCIE_CPU_CX_10P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	21

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
22	GND	GND BAR			22
23	PCIE_CPU_CX_9N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	23
24	PCIE_CPU_CX_9P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	24
25	GND	GND BAR			25
26	PCIE_CPU_CX_8N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	26
27	PCIE_CPU_CX_8P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	27
28	GND	GND BAR			28
29	PCIE_CPU_CX_7N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	29
30	PCIE_CPU_CX_7P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	30
31	GND	GND BAR			31

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
32	PCIE_CPU_CX_6N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	32
33	PCIE_CPU_CX_6P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	33
34	GND	GND BAR			34
35	PCIE_CPU_CX_5N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	35
36	PCIE_CPU_CX_5P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	36
37	GND	GND BAR			37
38	PCIE_CPU_CX_4N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	38
39	PCIE_CPU_CX_4P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	39
40	GND	GND BAR			40
41	PCIE_CPU_CX_3N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	41

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
42	PCIE_CPU_CX_3P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	42
43	GND	GND BAR			43
44	PCIE_CPU_CX_2N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	44
45	PCIE_CPU_CX_2P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	45
46	GND	GND BAR			46
47	PCIE_CPU_CX_1N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	47
48	PCIE_CPU_CX_1P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	48
49	GND	GND BAR			49
50	PCIE_CPU_CX_0N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	50
51	PCIE_CPU_CX_0P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	51

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
52	GND	GND BAR			52
53	I2C_SMDAT	Micro coax	PCIe compliant SMBus interface to the network adapter. The Network adapter silicon may be accessed through this SMBus interface after the Power-On sequence of the network adapter is completed. The motherboard side serves as the bus master on the SMBus interface. (Please note on AUX board implementation, an I2C EEPROM at address 0x50 is also present on this interface, to allow the server management entity to read the identity of the auxiliary card. For a direct motherboard interface to the Cabline CA-II Plus cables, mounting such FRU EEPROM on the motherboard should be considered - for motherboard management purposes of the Cabline CA-II Plus interface to the network adapter).	38	53

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
54	I2C_SMCLK	Micro coax	PCIe compliant SMBus interface to the network adapter. The Network adapter silicon may be accessed through this SMBus interface after the Power-On sequence of the network adapter is completed. The motherboard side serves as the bus master on the SMBus interface. (Please note on AUX board implementation, an I2C EEPROM at address 0x50 is also present on this interface, to allow the server management entity to read the identity of the auxiliary card. For a direct motherboard interface to the Cabline CA-II Plus cables, mounting such FRU EEPROM on the motherboard should be considered - for motherboard management purposes of the Cabline CA-II Plus interface to the network adapter).	38	54
55	Reserved_01	Micro coax		38	55

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
56	S_WAKE1_CONN_L	Micro coax	In setups that support operation of the network adapter card during standby operating mode, this optional PCIe compliant Wake signal, may be used for the network adapter silicon to indicate to the motherboard when wake from standby power mode, for example after a magic packet was received. Please consult NVIDIA in case of intention to use this signal, as it has not been populated or used in designs so far.	38	56
57	I2C_AUX_SCL	Micro coax	The ConnectX-7 silicon serves as the I2C bus master on this bus. An I2C EEPROM at I2C address 0x57 needs to be mounted on the motherboard side to report to the Cabline CA-II Plus interface parameters to the main-card network adapter silicon, like Cabline CA-II Plus cables length (contact NVIDIA for the format of this EEPROM). If additional optional I2C slave devices need to be managed by the main-card network adapter silicon, they need to be included on this I2C bus as well.	38	57

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
58	I2C_AUX_SDA	Micro coax	The ConnectX-7 silicon serves as the I2C bus master on this bus. An I2C EEPROM at I2C address 0x57 needs to be mounted on the motherboard side to report to the Cabline CA-II Plus interface parameters to the main-card network adapter silicon, like Cabline CA-II Plus cables length (contact NVIDIA for the format of this EEPROM). If additional optional I2C slave devices need to be managed by the main-card network adapter silicon, they need to be included on this I2C bus as well.	38	58
59	S_PRSNT1_L	Micro coax	Connect this pin to GND No wires are connected to these pins to ensure they do not interfere with the operation of S_PRSNT2_L for the detection when the two Cabline harnesses are installed.	38	59
60		No Wire			60

Print Side

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
1	SER_CLK	Micro coax	38		1
2	SER_CAPTURE	Micro coax	38		2
3	SER_DO	Micro coax	38		3

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
4	Reserved_04	Micro coax	38		4
5	SER_DI	Micro coax	38		5
6	S_BIFFURCATE_AUX_BRD	Micro coax	38	In specific Network Card Assemblies which support bifurcation of the PCIe x16 IPEX interface to two x8 PCIe interfaces, this pin is used for either the motherboard or the network adapter silicon to signal whether the IPEX interface should or should not be bifurcated. This signal has been reserved and hasn't been used in the past, and as a result it's direction has not been fixed yet, to provide for additional options in the future. (pin is marked as Reserved 06 in the generic White IPEX Cable pinout)	6
7	Reserved_07	Micro coax	38	Reserved for future expansion	7
8	Reserved_08	Micro coax	38	Reserved for future expansion	8
9	GND	GND BAR			9
10	PCIE_CX_CPU_0P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	10
11	PCIE_CX_CPU_ON	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	11

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
12	GND	GND BAR			12
13	PCIE_CX_CPU_1P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	13
14	PCIE_CX_CPU_1N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	14
15	GND	GND BAR			15
16	PCIE_CX_CPU_2P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	16
17	PCIE_CX_CPU_2N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	17
18	GND	GND BAR			18
19	PCIE_CX_CPU_3P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	19

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
20	PCIE_CX_CPU_3N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	20
21	GND	GND BAR			21
22	PCIE_CX_CPU_4P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	22
23	PCIE_CX_CPU_4N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	23
24	GND	GND BAR			24
25	PCIE_CX_CPU_5P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	25
26	PCIE_CX_CPU_5N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	26
27	GND	GND BAR			27

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
28	PCIE_CX_CPU_6P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	28
29	PCIE_CX_CPU_6N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	29
30	GND	GND BAR			30
31	PCIE_X_CPU_7P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	31
32	PCIE_CX_CPU_7N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	32
33	GND	GND BAR			33
34	PCIE_CX_CPU_8P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	34

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
35	PCIE_CX_CPU_8N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	35
36	GND	GND BAR			36
37	PCIE_CX_CPU_9P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	37
38	PCIE_CX_CPU_9N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	38
39	GND	GND BAR			39
40	PCIE_CX_CPU_10P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	40
41	PCIE_CX_CPU_10N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	41
42	GND	GND BAR			42

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
43	PCIE_CX_CPU_11P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	43
44	PCIE_CX_CPU_11N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	44
45	GND	GND BAR			45
46	PCIE_CX_CPU_12P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	46
47	PCIE_CX_CPU_12N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	47
48	GND	GND BAR			48
49	PCIE_CX_CPU_13P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	49

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
50	PCIE_CX_CPU_13N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	50
51	GND	GND BAR			51
52	PCIE_CX_CPU_14P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	52
53	PCIE_CX_CPU_14N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	53
54	GND	GND BAR			54
55	PCIE_CX_CPU_15P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	55
56	PCIE_CX_CPU_15N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the ConnectX-7 card.	56
57	GND	GND BAR			57

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
58	S_PERST1_CONN_L	Micro coax	38	PCIe compliant PERST_L (active low PCI Reset) signal for the PCIe Cabline CA-II Plus Connectors. The direction of this PERST_L signal depends on the implementation: When connecting a CPU root complex to the PCIe Cabline CA-II Plus interface, this signal is driven from the motherboard side (from the CPU), to the network adapter. When connecting a GPU or an end point to the PCIe Cabline CA-II Plus interface, this signal is driven from the network adapter side (which operates as a PCIe switch in this case), to the GPU or end-point on the motherboard side.	58
59		No Wire			59

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
60	S_PRSNT2_L	Micro coax	38	Connect to a 4.7K pull-up resistor to 3.3V on the motherboard side, to detect if both the Cabline harnesses are connected or not. This signal is connected to S_PRSNT1_L on the network adapter main card. In the motherboard side, read logic low if both Cabline harnesses are connected. Read logic 1 (3.3V) if one or both the Cabline harnesses are not connected. No wires are connected to these pins to ensure they do not interfere with the operation of S_PRSNT1_L for the detection when the two Cabline harnesses are installed.	60

PCIe Auxiliary Kit Technical Specifications


Technical Specifications

Physical	PCIe Auxiliary Card Size: 5.09 in. x 2.32 in. (129.30mm x 59.00mm) Two Cabline CA-II Plus harnesses (white and black) Length: 15, 25 or 35cm		
PCIe Connectivity	MTMK9100-T15	PCI Express Gen 5.0/4.0: SERDES @ 16/32 GT/s, x16 lanes (Gen 3.0 compatible)	
	MTMK9100-T25 / MTKM9100-T35	PCI Express Gen 4.0: SERDES @ 16GT/s, x16 lanes (Gen 3.0 compatible)	
Power Consumption	Voltage: 12V, 3.3V_AUX Maximum current: 100mA		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	Humidity	Operational	10% to 85% relative humidity

Regulatory		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
	Safety:	CB / cTUVus / CE	
	EMC:	CE / FCC / VCCI / ICES / RCM / KC	
	RoHS:	RoHS Compliant	

PCIe Auxiliary Card Mechanical Drawings and Dimensions


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



Bracket Mechanical Drawings and Dimensions

Auxiliary PCIe Connection Card Tall Bracket	Auxiliary PCIe Connection Card Short Bracket
The drawing shows a side view of the Auxiliary PCIe Connection Card Tall Bracket. It is a grey, rectangular bracket with a series of rectangular slots. The dimensions are indicated by green arrows: a length of 121.05 mm and a height of 21.59 mm.	The drawing shows a side view of the Auxiliary PCIe Connection Card Short Bracket. It is a grey, rectangular bracket with a series of rectangular slots. The dimensions are indicated by green arrows: a length of 80.29 mm and a height of 22.83 mm.

Cabline CA-II Plus Harnesses Mechanical Drawing



Document Revision History

Date	Comments/Changes
Apr. 2023	Updated Setting High-Speed-Port Link Type
Feb. 2023	Updated VMware Driver Installation
Jan. 2023	- Added support for OPN MCX715105AS-WEAT - Updated 400Gb/s Ethernet protocols in Specifications
Dec. 2022	- Modified ordering part numbers and product descriptions in NVIDIA ConnectX-7 Adapter Cards User Manual - Added PCIe Auxiliary Card Kit - Updated Monitoring
Nov. 2022	First release of the consolidated user manual for all ConnectX-7 adapter cards

Notice

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

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

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

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

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

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

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

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

Trademarks

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



Copyright

© 2023 NVIDIA Corporation & affiliates. All Rights Reserved.

