



NVIDIA ConnectX-6 InfiniBand/VPI Adapter Cards User Manual

Table of Contents

Introduction.....	8
Product Overview	8
ConnectX-6 PCIe x8 Card	9
ConnectX-6 PCIe x16 Card	10
ConnectX-6 DE PCIe x16 Card	11
ConnectX-6 for Liquid-Cooled Intel® Server System D50TNP Platforms.....	11
ConnectX-6 Socket Direct™ Cards	12
Package Contents	15
ConnectX-6 PCIe x8/x16 Adapter Cards	15
ConnectX-6 PCIe x16 Adapter Card for liquid-cooled Intel® Server System D50TNP Platforms	15
ConnectX-6 Socket Direct Cards (2x PCIe x16)	15
Features and Benefits	16
Operating Systems/Distributions	19
Connectivity	19
Manageability.....	19
Interfaces	20
InfiniBand Interface	20
Ethernet Interfaces	20
PCI Express Interface	20
LED Interface	20
Heatsink Interface	22
SMBus Interface	22
Voltage Regulators	22
Thermal Sensors.....	23
Hardware Installation	24
Safety Warnings	24
Installation Procedure Overview.....	25
System Requirements	25
Hardware Requirements	25
Airflow Requirements	26
Software Requirements	27

Safety Precautions	27
Pre-Installation Checklist	27
Bracket Replacement Instructions	27
Installation Instructions	28
Cables and Modules	29
Identifying the Card in Your System	29
PCIe x8/16 Cards Installation Instructions	31
Installing the Card	31
Uninstalling the Card	33
Socket Direct (2x PCIe x16) Cards Installation Instructions	34
Installing the Card	35
Uninstalling the Card	41
Cards for Intel Liquid-Cooled Platforms Installation Instructions	42
Installing the Card	43
Driver Installation	47
Linux Driver Installation	47
Prerequisites	47
Downloading NVIDIA OFED	47
Installing NVIDIA OFED	49
Installing MLNX_OFED Using YUM	54
Installing MLNX_OFED Using apt-get Tool	56
Updating Firmware After Installation	58
UEFI Secure Boot	59
Performance Tuning	61
Windows Driver Installation	61
Software Requirements	61
Downloading NVIDIA WinOF-2 Driver	61
Uninstalling NVIDIA WinOF-2 Driver	69
Extracting Files Without Running Installation	69
Firmware Upgrade	72
VMware Driver Installation	72
Hardware and Software Requirements	72
Installing NVIDIA NATIVE ESXi Driver for VMware vSphere	73
Removing Earlier NVIDIA Drivers	73

Firmware Programming	74
Updating Adapter Firmware	75
Troubleshooting	76
General Troubleshooting	76
Linux Troubleshooting	77
Windows Troubleshooting	77
Specifications	79
MCX651105A-EDAT Specifications	79
MCX653105A-HDAT Specifications	80
MCX653106A-HDAT Specifications	82
MCX653105A-HDAL Specifications	83
MCX653106A-HDAL Specifications	85
MCX653105A-ECAT Specifications	86
MCX653106A-ECAT Specifications	88
MCX654105A-HCAT Specifications	89
MCX654106A-HCAT Specifications	91
MCX654106A-ECAT Specifications	93
MCX653105A-EFAT Specifications	95
MCX653106A-EFAT Specifications	97
MCX683105AN-HDAT Specifications	98
Adapter Card and Bracket Mechanical Drawings and Dimensions	100
Adapter Cards	100
Brackets Dimensions	101
PCI Express Pinouts Description for Single-Slot Socket Direct Card	102
Finding the GUID/MAC on the Adapter Card	103
Document Revision History	104

About This Manual

This User Manual describes NVIDIA® ConnectX®-6 InfiniBand/VPI adapter cards. It provides details as to the interfaces of the board, specifications, required software and firmware for operating the board, and relevant documentation.

Ordering Part Numbers

The table below provides the ordering part numbers (OPN) for the available ConnectX-6 InfiniBand/VPI adapter cards.

OPN	Marketing Description
MCX65410 6A-ECAT	ConnectX®-6 VPI adapter card, 100Gb/s (HDR100, EDR InfiniBand and 100GbE), dual-port QSFP56, Socket Direct 2x PCIe 3.0/4.0 x16 , tall bracket
MCX65310 5A-EFAT	ConnectX®-6 VPI adapter card, 100Gb/s (HDR100, EDR IB and 100GbE), single-port QSFP56, PCIe3.0/4.0 Socket Direct 2x8 in a row , tall bracket
MCX65310 6A-EFAT	ConnectX®-6 VPI adapter card, 100Gb/s (HDR100, EDR IBand100GbE), dual-port QSFP56, PCIe3.0/4.0 Socket Direct 2x8 in a row , tall bracket
MCX65110 5A-EDAT	ConnectX®-6 VPI adapter card, 100Gb/s (HDR100, EDR IB and 100GbE), single-port QSFP56, PCIe4.0 x8 , tall bracket
MCX65310 5A-ECAT	ConnectX®-6 VPI adapter card, 100Gb/s (HDR100, EDR IB and 100GbE), single-port QSFP56, PCIe3.0/4.0 x16 , tall bracket
MCX65310 6A-ECAT	ConnectX®-6 VPI adapter card, 100Gb/s (HDR100, EDR IB and 100GbE), dual-port QSFP56, PCIe3.0/4.0 x16 , tall bracket
MCX65310 5A-HDAT	ConnectX®-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE , single-port QSFP56, PCIe3.0/4.0 x16 , tall bracket
MCX65310 6A-HDAT	ConnectX®-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE , dual-port QSFP56, PCIe3.0/4.0 x16 , tall bracket
MCX65410 5A-HCAT	ConnectX®-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE , single-port QSFP56, Socket Direct 2x PCIe3.0/4.0x16 , tall bracket
MCX65410 6A-HCAT	ConnectX®-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE , dual-port QSFP56, Socket Direct 2x PCIe3.0/4.0x16 , tall bracket
MCX65310 5A-HDAL	ConnectX®-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE , single-port QSFP56, PCIe3.0/4.0 x16, cold plate for liquid-cooled Intel® Server System D50TNP platforms , tall bracket, ROHS R6
MCX65310 6A-HDAL	ConnectX®-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE , dual-port QSFP56, PCIe3.0/4.0 x16, cold plate for liquid-cooled Intel® Server System D50TNP platforms , tall bracket, ROHS R6
MCX68310 5AN-HDAT	ConnectX®-6 DE adapter card, HDR IB (200Gb/s) , single-port QSFP, PCIe4.0 x16 , no crypto, tall bracket

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.

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: Network-support@nvidia.com

Customers who purchased NVIDIA M-1 Global Support Services, please see your contract for details regarding Technical Support.

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

Related Documentation

MLNX_OFED for Linux User Manual and Release Notes	User Manual describing OFED features, performance, band diagnostic, tools content, and configuration. See MLNX_OFED for Linux Documentation .
WinOF-2 for Windows User Manual and Release Notes	User Manual describing WinOF-2 features, performance, Ethernet diagnostic, tools content, and configuration. See WinOF-2 for Windows Documentation .
NVIDIA VMware for Ethernet User Manual	User Manual and release notes describing the various components of the NVIDIA ConnectX® NATIVE ESXi stack. See VMware® ESXi Drivers Documentation .
NVIDIA Firmware Utility (mlxup) User Manual and Release Notes	NVIDIA firmware update and query utility used to update the firmware. Refer to Firmware Utility (mlxup) Documentation .
NVIDIA Firmware Tools (MFT) User Manual	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 1.3	InfiniBand Specifications
IEEE Std 802.3 Specification	IEEE Ethernet Specifications
PCI Express Specifications	Industry Standard PCI Express Base and Card Electromechanical Specifications. Refer to PCI-SIG Specifications .
LinkX Interconnect Solutions	LinkX InfiniBand 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 broadest portfolio of QDR/FDR10 (40Gb/s), FDR (56Gb/s), EDR/HDR100 (100Gb/s), HDR (200Gb/s) and NDR (400Gb/s) cables, 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 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 .

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 the size in MegaBits. In this document, PCIe is used to mean PCI Express.

Revision History

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

Introduction

Product Overview

This is the user guide for VPI adapter cards based on the ConnectX-6 integrated circuit device. ConnectX-6 connectivity provides the highest performing low latency and most flexible interconnect solution for PCI Express Gen 3.0/4.0 servers used in enterprise datacenters and high-performance computing environments.

ConnectX-6 Virtual Protocol Interconnect® adapter cards provide up to two ports of 200Gb/s for InfiniBand and Ethernet connectivity, sub-600ns latency and 200 million messages per second, enabling the highest performance and most flexible solution for the most demanding High-Performance Computing (HPC), storage, and datacenter applications.

ConnectX-6 is a groundbreaking addition to the NVIDIA ConnectX series of industry-leading adapter cards. In addition to all the existing innovative features of past ConnectX versions, ConnectX-6 offers a number of enhancements that further improve the performance and scalability of datacenter applications. In addition, specific PCIe stand-up cards are available with a cold plate for insertion into liquid-cooled Intel® Server System D50TNP platforms.

ConnectX-6 is available in two form factors: low-profile stand-up PCIe and Open Compute Project (OCP) Spec 3.0 cards with QSFP connectors. Single-port, HDR, stand-up PCIe adapters are available based on either ConnectX-6 or ConnectX-6 DE (ConnectX-6 Dx enhanced for HPC applications).



Make sure to use a PCIe slot that is capable of supplying the required power and airflow to the ConnectX-6 as stated in [Specifications](#).

Configuration	OPN	Marketing Description
ConnectX-6 PCIe x8 Card	MCX651105 A-EDAT	ConnectX-6 VPI adapter card, 100Gb/s (HDR100, EDR IB and 100GbE), single-port QSFP56, PCIe4.0 x8, tall bracket
ConnectX-6 PCIe x16 Card	MCX653105 A-HDAT	ConnectX-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE, single-port QSFP56, PCIe4.0 x16, tall bracket
	MCX653106 A-HDAT	ConnectX-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE, dual-port QSFP56, PCIe3.0/4.0 x16, tall bracket
	MCX653105 A-ECAT	ConnectX-6 VPI adapter card, 100Gb/s (HDR100, EDR IB and 100GbE), single-port QSFP56, PCIe3.0/4.0 x16, tall bracket
	MCX653106 A-ECAT	ConnectX-6 VPI adapter card, 100Gb/s (HDR100, EDR IB and 100GbE), dual-port QSFP56, PCIe3.0/4.0 x16, tall bracket
ConnectX-6 DE PCIe x16 Card	MCX683105 AN-HDAT	ConnectX-6 DE InfiniBand adapter card, HDR, single-port QSFP, PCIe 3.0/4.0 x16, No Crypto, Tall Bracket
ConnectX-6 PCIe x16 Cards for liquid-cooled Intel® Server System D50TNP platforms	MCX653105 A-HDAL	ConnectX-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE, single-port QSFP56, PCIe4.0 x16, cold plate for liquid-cooled Intel® Server System D50TNP platforms, tall bracket, ROHS R6
	MCX653106 A-HDAL	ConnectX-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE, dual-port QSFP56, PCIe4.0 x16, cold plate for liquid-cooled Intel® Server System D50TNP platforms, tall bracket, ROHS R6

Configuration	OPN	Marketing Description
ConnectX-6 Dual-slot Socket Direct Cards (2x PCIe x16)	MCX654105 A-HCAT	ConnectX-6 VPI adapter card kit, HDR IB (200Gb/s) and 200GbE, single-port QSFP56, Socket Direct 2x PCIe3.0 x16, tall brackets
	MCX654106 A-HCAT	ConnectX-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE, dual-port QSFP56, Socket Direct 2x PCIe3.0/4.0x16, tall bracket
	MCX654106 A-ECAT	ConnectX-6 VPI adapter card, 100Gb/s (HDR100, EDR InfiniBand and 100GbE), dual-port QSFP56, Socket Direct 2x PCIe3.0/4.0 x16, tall bracket
ConnectX-6 Single-slot Socket Direct Cards (2x PCIe x8 in a row)	MCX653105 A-EFAT	ConnectX-6 VPI adapter card, 100Gb/s (HDR100, EDR IB and 100GbE), single-port QSFP56, PCIe3.0/4.0 Socket Direct 2x8 in a row, tall bracket
	MCX653106 A-EFAT	ConnectX-6 VPI adapter card, 100Gb/s (HDR100, EDR IBand100GbE), dual-port QSFP56, PCIe3.0/4.0 Socket Direct 2x8 in a row, tall bracket

ConnectX-6 PCIe x8 Card

ConnectX-6 with a single PCIe x8 slot can support a bandwidth of up to 100Gb/s in a PCIe Gen 4.0 slot.



Part Number	MCX651105A-EDAT
Form Factor/Dimensions	PCIe Half Height, Half Length / 167.65mm x 68.90mm
Data Transmission Rate	Ethernet: 10/25/40/50/100 Gb/s InfiniBand: SDR, DDR, QDR, FDR, EDR, HDR100
Network Connector Type	Single-port QSFP56

Part Number	MCX651105A-EDAT
PCIe x8 through Edge Connector	PCIe Gen 3.0 / 4.0 SERDES @ 8.0GT/s / 16.0GT/s
RoHS	RoHS Compliant
Adapter IC Part Number	MT28908A0-XCCF-HVM

ConnectX-6 PCIe x16 Card

ConnectX-6 with a single PCIe x16 slot can support a bandwidth of up to 100Gb/s in a PCIe Gen 3.0 slot, or up to 200Gb/s in a PCIe Gen 4.0 slot. This form-factor is available also for Intel® Server System D50TNP Platforms where an Intel liquid-cooled cold plate is used for adapter cooling mechanism.



Part Number	MCX653105A-ECAT	MCX653106A-ECAT	MCX653105A-HDAT	MCX653106A-HDAT
Form Factor/Dimensions	PCIe Half Height, Half Length / 167.65mm x 68.90mm			
Data Transmission Rate	Ethernet: 10/25/40/50/100 Gb/s InfiniBand: SDR, DDR, QDR, FDR, EDR, HDR100		Ethernet: 10/25/40/50/100/200 Gb/s InfiniBand: SDR, DDR, QDR, FDR, EDR, HDR100, HDR	
Network Connector Type	Single-port QSFP56	Dual-port QSFP56	Single-port QSFP56	Dual-port QSFP56
PCIe x16 through Edge Connector	PCIe Gen 3.0 / 4.0 SERDES @ 8.0GT/s / 16.0GT/s			

Part Number	MCX653105A-ECAT	MCX653106A-ECAT	MCX653105A-HDAT	MCX653106A-HDAT
RoHS	RoHS Compliant			
Adapter IC Part Number	MT28908A0-XCCF-HVM			


ConnectX-6 DE PCIe x16 Card

ConnectX-6 DE (ConnectX-6 Dx enhanced for HPC applications) with a single PCIe x16 slot can support a bandwidth of up to 100Gb/s in a PCIe Gen 3.0 slot, or up to 200Gb/s in a PCIe Gen 4.0 slot.

Part Number	MCX683105AN-HDAT
Form Factor/Dimensions	PCIe Half Height, Half Length / 167.65mm x 68.90mm
Data Transmission Rate	InfiniBand: SDR, DDR, QDR, FDR, EDR, HDR100, HDR
Network Connector Type	Single-port QSFP56
PCIe x16 through Edge Connector	PCIe Gen 3.0 / 4.0 SERDES @ 8.0GT/s / 16.0GT/s
RoHS	RoHS Compliant
Adapter IC Part Number	MT28924A0-NCCF-VE

ConnectX-6 for Liquid-Cooled Intel® Server System D50TNP Platforms

The below cards are available with a cold plate for insertion into liquid-cooled Intel® Server System D50TNP platforms.



Part Number	MCX653105A-HDAL	MCX653106A-HDAL
Form Factor/Dimensions	PCIe Half Height, Half Length / 167.65mm x 68.90mm	
Data Transmission Rate	Ethernet: 10/25/40/50/100/200 Gb/s InfiniBand: SDR, DDR, QDR, FDR, EDR, HDR100, HDR	
Network Connector Type	Single-port QSFP56	Dual-port QSFP56
PCIe x16 through Edge Connector	PCIe Gen 3.0 / 4.0 SERDES @ 8.0GT/s / 16.0GT/s	
RoHS	RoHS Compliant	
Adapter IC Part Number	MT28908A0-XCCF-HVM	

ConnectX-6 Socket Direct™ 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.


Please note that ConnectX-6 Socket Direct cards do not support Multi-Host functionality (i.e. connectivity to two independent CPUs). For ConnectX-6 Socket Direct card with Multi-Host functionality, please contact NVIDIA.

ConnectX-6 Socket Direct cards are available in two configurations: Dual-slot Configuration (2x PCIe x16) and Single-slot Configuration (2x PCIe x8).

ConnectX-6 Dual-slot Socket Direct Cards (2x PCIe x16)

In order to obtain 200Gb/s speed, NVIDIA offers ConnectX-6 Socket Direct that enable 200Gb/s connectivity also for servers with PCIe Gen 3.0 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, as shown in the below figure.




Part Number	MCX654105A-HCAT	MCX654106A-HCAT	MCX654106A-ECAT
Form Factor/Dimensions	Adapter Card: PCIe Half Height, Half Length / 167.65mm x 68.90mm Auxiliary PCIe Connection Card: 5.09 in. x 2.32 in. (129.30mm x 59.00mm) Two 35cm Cabline CA-II Plus harnesses		
Data Transmission Rate	Ethernet: 10/25/40/50/100/200 Gb/s InfiniBand: SDR, DDR, QDR, FDR, EDR, HDR100, HDR		Ethernet: 10/25/40/50/100 Gb/s InfiniBand: SDR, DDR, QDR, FDR, EDR, HDR100
Network Connector Type	Single-port QSFP56	Dual-port QSFP56	
PCIe x16 through Edge Connector	PCIe Gen 3.0 / 4.0SERDES@ 8.0GT/s / 16.0GT/s		
PCIe x16 through Auxiliary Card	PCIe Gen 3.0SERDES@ 8.0GT/s		
RoHS	RoHS Compliant		
Adapter IC Part Number	MT28908A0-XCCF-HVM		

ConnectX-6 Single-slot Socket Direct Cards (2x PCIe x8 in a row)

The PCIe x16 interface comprises two PCIe x8 in a row, such that each of the PCIe x8 lanes can be connected to a dedicated CPU in a dual-socket server. In such a configuration, Socket Direct brings lower latency and lower CPU utilization as the direct connection from each CPU to the network means the interconnect can bypass a QPI (UPI) and the other CPU, optimizing performance and improving latency. CPU utilization is improved as each CPU handles only its own traffic and not traffic from the other CPU.

A system with a custom PCI Express x16 slot that includes special signals is required for installing the card. Please refer to [PCI Express Pinouts Description for Single-Slot Socket Direct Card](#) for pinout definitions.



Part Number	MCX653105A-EFAT	MCX653106A-EFAT
Form Factor/Dimensions	PCIe Half Height, Half Length / 167.65mm x 68.90mm	
Data Transmission Rate	Ethernet: 10/25/40/50/100 Gb/s InfiniBand: SDR, DDR, QDR, FDR, EDR, HDR100	
Network Connector Type	Single-port QSFP56	Dual-port QSFP56
PCIe x16 through Edge Connector	PCIe Gen 3.0 / 4.0 SERDES @ 8.0GT/s / 16.0GT/s Socket Direct 2x8 in a row	
RoHS	RoHS Compliant	
Adapter IC Part Number	MT28908A0-XCCF-HVM	

Package Contents

ConnectX-6 PCIe x8/x16 Adapter Cards



Applies to MCX651105A-EDAT, MCX653105A-ECAT, MCX653106A-ECAT, MCX653105A-HDAT, MCX653106A-HDAT, MCX653105A-EFAT, MCX653106A-EFAT, and MCX683105AN-HDAT.

Category	Qty	Item
Cards	1	ConnectX-6 adapter card
Accessories	1	Adapter card short bracket
	1	Adapter card tall bracket (shipped assembled on the card)

ConnectX-6 PCIe x16 Adapter Card for liquid-cooled Intel® Server System D50TNP Platforms



Applies to MCX653105A-HDAL and MCX653106A-HDAL.

Category	Qty	Item
Cards	1	ConnectX-6 adapter card
Accessories	1	Adapter card short bracket
	1	Adapter card tall bracket (shipped assembled on the card)
	1	Accessory Kit with two 2 TIMs (MEB000386)

ConnectX-6 Socket Direct Cards (2x PCIe x16)



Applies to MCX654105A-HCAT, MCX654106A-HCAT and MCX654106A-ECAT.

Category	Qty.	Item
Cards	1	ConnectX-6 adapter card
	1	PCIe Auxiliary Card
Harnesses	1	35cm Cabline CA-II Plus harness (white)
	1	35cm Cabline CA-II Plus harness (black)
Accessories	2	Retention Clip for Cabline harness (optional accessory)
	1	Adapter card short bracket
	1	Adapter card tall bracket (shipped assembled on the card)
	1	PCIe Auxiliary card short bracket
	1	PCIe Auxiliary card tall bracket (shipped assembled on the card)

Features and Benefits



Make sure to use a PCIe slot that is capable of supplying the required power and airflow to the ConnectX-6 cards as stated in [Specifications](#).

PCI Express (PCIe)	<p>Uses the following PCIe interfaces:</p> <ul style="list-style-type: none"> • PCIe x8/x16 configurations: PCIe Gen 3.0 (8GT/s) and Gen 4.0 (16GT/s) through an x8/x16 edge connector. • 2x PCIe x16 configurations: PCIe Gen 3.0/4.0 SERDES @ 8.0/16.0 GT/s through Edge Connector PCIe Gen 3.0 SERDES @ 8.0GT/s through PCIe Auxiliary Connection Card
200Gb/s Virtual Protocol Interconnect (VPI) Adapter	<p>ConnectX-6 offers the highest throughput VPI adapter, supporting HDR 200b/s InfiniBand and 200Gb/s Ethernet and enabling any standard networking, clustering, or storage to operate seamlessly over any converged network leveraging a consolidated software stack.</p>

InfiniBand Architecture Specification v1.3 compliant	<p>ConnectX-6 delivers low latency, high bandwidth, and computing efficiency for performance-driven server and storage clustering applications. ConnectX-6 is InfiniBand Architecture Specification v1.3 compliant.</p>
Up to 200 Gigabit Ethernet	<p>NVIDIA adapters comply with the following IEEE 802.3 standards: 200GbE / 100GbE / 50GbE / 40GbE / 25GbE / 10GbE / 1GbE</p> <ul style="list-style-type: none"> - IEEE 802.3bj, 802.3bm 100 Gigabit Ethernet - IEEE 802.3by, Ethernet Consortium25, 50 Gigabit Ethernet, supporting all FEC modes - IEEE 802.3ba 40 Gigabit Ethernet - IEEE 802.3by 25 Gigabit Ethernet - IEEE 802.3ae 10 Gigabit Ethernet - IEEE 802.3ap based auto-negotiation and KR startup - IEEE 802.3ad, 802.1AX Link Aggregation - IEEE 802.1Q, 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 - Jumbo frame support (9.6KB)
InfiniBand HDR100	<p>A standard InfiniBand data rate, where each lane of a 2X port runs a bit rate of 53.125Gb/s with a 64b/66b encoding, resulting in an effective bandwidth of 100Gb/s.</p>
InfiniBand HDR	<p>A standard InfiniBand data rate, where each lane of a 4X port runs a bit rate of 53.125Gb/s with a 64b/66b encoding, resulting in an effective bandwidth of 200Gb/s.</p>
Memory Components	<ul style="list-style-type: none"> • SPI Quad - includes 256Mbit SPI Quad Flash device (MX25L25645GXDI-08G device by Macronix) • 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	<p>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-6 effectively addresses this by providing advanced NVGRE and VXLAN hardware offloading engines that encapsulate and de-capsulate the overlay protocol.</p>

RDMA and RDMA over Converged Ethernet (RoCE)	<p>ConnectX-6, 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-6 advanced congestion control hardware mechanisms, RoCE provides efficient low-latency RDMA services over Layer 2 and Layer 3 networks.</p>
NVIDIA PeerDirect™	<p>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-6 advanced acceleration technology enables higher cluster efficiency and scalability to tens of thousands of nodes.</p>
CPU Offload	<p>Adapter functionality enables reduced CPU overhead leaving more CPU resources available for computation tasks.</p> <p>Open vSwitch (OVS) offload using ASAP²(™)</p> <ul style="list-style-type: none"> • Flexible match-action flow tables • Tunneling encapsulation/decapsulation
Quality of Service (QoS)	<p>Support for port-based Quality of Service enabling various application requirements for latency and SLA.</p>
Hardware-based I/O Virtualization	<p>ConnectX-6 provides dedicated adapter resources and guaranteed isolation and protection for virtual machines within the server.</p>
Storage Acceleration	<p>A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols can leverage:</p> <ul style="list-style-type: none"> • RDMA for high-performance storage access • NVMe over Fabric offloads for target machine • Erasure Coding • T10-DIF Signature Handover
SR-IOV	<p>ConnectX-6 SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server.</p>
High-Performance Accelerations	<ul style="list-style-type: none"> • Tag Matching and Rendezvous Offloads • Adaptive Routing on Reliable Transport • Burst Buffer Offloads for Background Checkpointing

Operating Systems/Distributions



ConnectX-6 Socket Direct cards 2x PCIe x16 (OPNs: MCX654105A-HCAT, MCX654106A-HCAT and MCX654106A-ECAT) are not supported in Windows and WinOF-2.

- OpenFabrics Enterprise Distribution (OFED)
- RHEL/CentOS
- Windows
- FreeBSD
- VMware
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinOF-2)

Connectivity

- Interoperable with 1/10/25/40/50/100/200 Gb/s InfiniBand/VPI and Ethernet switches
- Passive copper cable with ESD protection
- Powered connectors for optical and active cable support

Manageability

ConnectX-6 technology maintains support for manageability through a BMC. ConnectX-6 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.

Interfaces

InfiniBand Interface

The network ports of the ConnectX®-6 adapter cards are compliant with the *InfiniBand Architecture Specification, Release 1.3*. InfiniBand traffic is transmitted through the cards' QSFP56 connectors.

Ethernet Interfaces



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

The network ports of the ConnectX-6 adapter card are compliant with the IEEE 802.3 Ethernet standards listed in [Features and Benefits](#). Ethernet traffic is transmitted through the QSFP56/QSFP connectors on the adapter card.

PCI Express Interface

ConnectX®-6 adapter cards support PCI Express Gen 3.0/4.0 (1.1 and 2.0 compatible) through x8/x16 edge connectors. The device can be either a master initiating the PCI Express bus operations or a subordinate responding to PCI bus operations.

The following lists PCIe interface features:

- PCIe Gen 3.0 and 4.0 compliant, 2.0 and 1.1 compatible
- 2.5, 5.0, 8.0, or 16.0 GT/s link rate x16/x32
- Auto-negotiates to x32, x16, x8, x4, x2, or x1
- Support for MSI/MSI-X mechanisms

LED Interface




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

There are two I/O LEDs per port:

- LED 1 and 2: Bi-color I/O LED which indicates link status. LED behavior is described below for Ethernet and InfiniBand port configurations.

- LED 3 and 4: Reserved for future use.



LED1 and LED2 Link Status Indications - Ethernet Protocol:

LED Color and State	Description									
Off	A link has not been established									
Beacon command for locating the adapter card	1Hz blinking Yellow									
Error	<p>4Hz blinking Yellow Indicates an error with the link. The error can be one of the following:</p> <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
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								
Solid green	Indicates a valid link with no active traffic									
Blinking green	Indicates a valid link with active traffic									

LED1 and LED2 Link Status Indications - InfiniBand Protocol:

LED Color and State	Description
Off	A link has not been established
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
Solid amber	Indicates an active link		
Solid green	Indicates a valid (data activity) link with no active traffic		
Blinking green	Indicates a valid link with active traffic		

Heatsink Interface

The heatsink is attached to the ConnectX-6 IC to dissipate the heat from the ConnectX-6 IC. It is attached either by using four spring-loaded push pins that insert into four mounting holes or by screws.

ConnectX-6 IC has a thermal shutdown safety mechanism that automatically shuts down the ConnectX-6 card in cases of high-temperature events, improper thermal coupling or heatsink removal.

For the required airflow (LFM) per OPN, please refer to [Specifications](#).

For MCX653105A-HDAL and MCX653106A-HDAL cards, the heatsink is compatible with a cold plate for liquid-cooled Intel® Server System D50TNP platforms only.

SMBus Interface

ConnectX-6 technology maintains support for manageability through a BMC. ConnectX-6 PCIe stand-up adapter can be connected to a BMC using MCTP over SMBus protocol 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.

Thermal Sensors

Unable to render include or excerpt-include. Could not retrieve page.

Hardware Installation

Installation and initialization of ConnectX-6 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](#) document available on [nvidia.com](https://www.nvidia.com).

Please observe all safety warnings to avoid injury and prevent damage to system components.

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-6 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-6 PCIe x8/x16 adapter card in the system	ConnectX-6 PCIe x8/x16 Adapter Cards Installation Instructions
	Install the ConnectX-6 2x PCIe x16 Socket Direct adapter card in the system	Socket Direct (2x PCIe x16) Cards Installation Instructions
	Install the ConnectX-6 card for Intel Liquid-cooled platforms	Cards for Intel Liquid-Cooled Platforms Installation Instructions
8	Connect cables or modules to the card	Cables and Modules
9	Identify ConnectX-6 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 and 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-6 cards in a PCIe slot that is capable of supplying the required power as stated in [Specifications](#).

ConnectX-6 Configuration	Hardware Requirements
PCIe x8/x16	A system with a PCI Express x8/x16 slot is required for installing the card.
Cards for liquid-cooled Intel® Server System D50TNP platforms	Intel® Server System D50TNP Platform with an available PCI Express x16 slot is required for installing the card.
Socket Direct 2x PCIe x8 in a row (single slot)	A system with a custom PCI Express x16 slot (four special pins) is required for installing the card. Please refer to PCI Express Pinouts Description for Single-Slot Socket Direct Card for pinout definitions.
Socket Direct 2x PCIe x16 (dual slots)	A system with two PCIe x16 slots is required for installing the cards.

Airflow Requirements

ConnectX-6 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.

Airflow from the heatsink to the network ports	Airflow from the network ports to the heatsink
	



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

Software Requirements

- See [Operating Systems/Distributions](#) section under the Introduction section.
- Software Stacks - NVIDIA OpenFabric software package MLNX_OFED for Linux, 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-6 Card; Unpack and remove the ConnectX-6 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-6 card, make sure that the system is disconnected from power.
- (Optional) Check the mounting bracket on the ConnectX-6 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-6 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.



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

To replace the bracket you will need the following parts:

- The new 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-6 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-6 configuration you have purchased.

OPNs	Installation Instructions
MCX651105A-EDAT MCX653105A-HDAT MCX653106A-HDAT MCX653105A-ECAT MCX653106A-ECAT MCX653105A-EFAT MCX653106A-EFAT MCX683105AN-HDAT	PCIe x8/16 Cards Installation Instructions
MCX654105A-HCAT MCX654106A-HCAT MCX654106A-ECAT	Socket Direct (2x PCIe x16) Cards Installation Instructions
MCX653105A-HDAL MCX653106A-HDAL	Cards for Intel Liquid-Cooled Platforms Installation Instructions


Cables and Modules

To obtain the list of supported NVIDIA cables for your adapter, please refer to the Cables Reference Table at <http://www.nvidia.com/products/interconnect/cables-configurator.php>.

Cable Installation

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 as not to impede the air exhaust flow through the ventilation holes. Use cable lengths which 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. 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”:

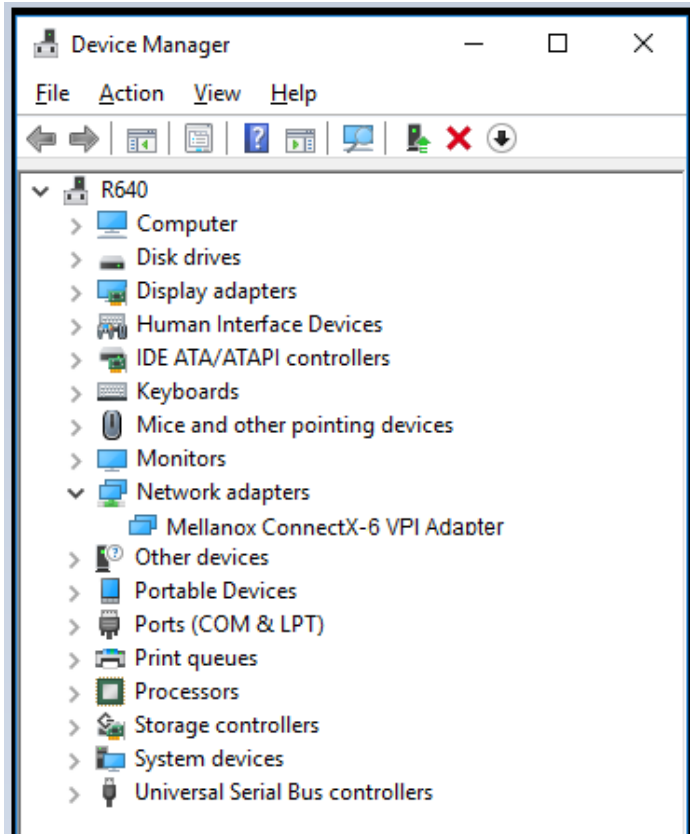
ConnectX-6 Card Configuration	lspci Command Output Example
--------------------------------------	-------------------------------------

Single-port Socket Direct Card (2x PCIe x16)	<pre>[root@mftqa-009 ~]# lspci grep mellanox -i a3:00.0 Infiniband controller: Mellanox Technologies MT28908 Family [ConnectX-6] e3:00.0 Infiniband controller: Mellanox Technologies MT28908 Family [ConnectX-6]</pre>
Dual-port Socket Direct Card (2x PCIe x16)	<pre>[root@mftqa-009 ~]# lspci grep mellanox -i 05:00.0 Infiniband controller: Mellanox Technologies MT28908A0 Family [ConnectX-6] 05:00.1 Infiniband controller: Mellanox Technologies MT28908A0 Family [ConnectX-6] 82:00.0 Infiniband controller: Mellanox Technologies MT28908A0 Family [ConnectX-6] 82:00.1 Infiniband controller: Mellanox Technologies MT28908A0 Family [ConnectX-6]</pre> <p>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 (hexa-decimal), 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-6 Socket Direct adapter are viewed as four net devices by the system.</p>
Single-port PCIe x8/x16 Card	<pre>[root@mftqa-009 ~]# lspci grep mellanox -i 3:00.0 Infiniband controller: Mellanox Technologies MT28908 Family [ConnectX-6]</pre>
Dual-port PCIe x16 Card	<pre>[root@mftqa-009 ~]# lspci grep mellanox -i 86:00.0 Network controller: Mellanox Technologies MT28908A0 Family [ConnectX-6] 86:00.1 Network controller: Mellanox Technologies MT28908A0 Family [ConnectX-6]</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 NVIDIA ConnectX-6 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 2012/R2/2016) 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_1003": VEN is equal to 0x15B3 - this is the Vendor ID of NVIDIA; and DEV is equal to 1018 (for ConnectX-6) - this is a valid NVIDIA PCI Device ID.



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



The list of NVIDIA PCI Device IDs can be found in the PCI ID repository at <http://pci-ids.ucw.cz/read/PC/15b3>.

PCIe x8/16 Cards Installation Instructions

Installing the Card




Applies to OPNs MCX651105A-EDAT, MCX654105A-HCAT, MCX654106A-HCAT, MCX683105AN-HDAT, MCX653106A-ECAT and MCX653105A-ECAT.




Please make sure to install the ConnectX-6 cards in a PCIe slot that is capable of supplying the required power and airflow as stated in [Specifications](#).

➤ 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.

Step 1: 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-6 and Auxiliary Connections cards out of the PCI Express slot.



Socket Direct (2x PCIe x16) Cards Installation Instructions

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 supplied 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 "WHITE" label should be connected to the connector on the ConnectX-6 and PCIe card engraved with "White Cable" while the one marked with "BLACK" label should be connected to the connector on the ConnectX-6 and PCIe card engraved with "Black Cable".



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

Installing the Card



Applies to MCX654105A-HCAT, MCX654106A-HCAT and MCX654106A-ECAT.



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-6 cards in a PCIe slot that is capable of supplying the required power and airflow as stated in [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 making sure the clip opening is facing the plugs.



Step 2: Plug the Cabline CA-II Plus harnesses on the ConnectX-6 adapter card while paying attention to the color-coding. As indicated on both sides of the card; plug the black harness to the component side and the white harness to 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 on the PCIe Auxiliary Card. As indicated on both sides of the Auxiliary connection card; plug the black harness to the component side and the white harness to the print side.




Step 7: Verify the plugs are locked.




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



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




Verify that both latches are firmly snapped



➤ Connect the ConnectX-6 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-6 adapter and PCIe Auxiliary Connection Cards to the chassis.

Step 1: Secure the brackets 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-6 and Auxiliary Connections cards out of the PCI Express slot.



Cards for Intel Liquid-Cooled Platforms Installation Instructions

The below instructions apply to ConnectX-6 cards designed for Intel liquid-cooled platforms with ASIC interposer cooling mechanism. OPNs: MCX653105A-HDAL and MCX653106A-HDAL.



The below figures are for illustration purposes only.



The below instructions should be used in conjunction with the server's documentation.

Installing the Card



Please make sure the system is capable of supplying the required power as stated in [Specifications](#).




Pay extra attention to the black bumpers located on the print side of the card. Failure to do so may harm the bumpers.

➤ Apply the supplied thermal pad (one of the two) on top of the ASIC interposer or onto the coldplate.



The thermal pads are shipped with two protective liners covering the pad on both sides. It is highly important to peel the liners as instructed below prior to applying them to the card.

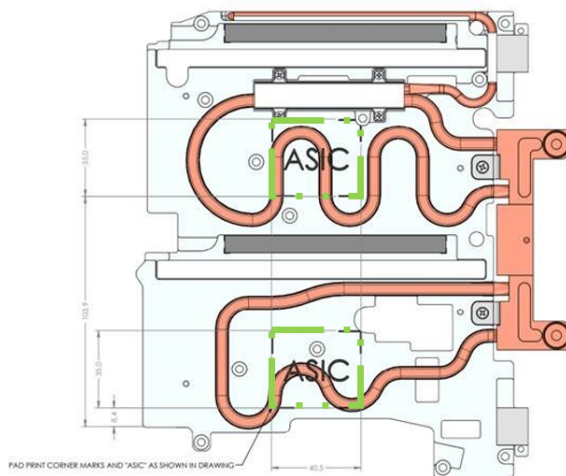
1. Gently peel the liner from the thermal pad's tacky side.
2. Carefully apply the thermal pad on the cool block (ASIC interposer) while ensuring it thoroughly covers it. Extra care should be taken not to damage the pad.
The thermal pad should be applied on the cool block from its tacky (wet) side. The pad should be applied with its non-tacky side facing up.




OR

Carefully apply the thermal pad on the coldplate while ensuring it thoroughly covers it. The below figure indicates the position of the thermal pad. Extra care should be taken not to damage the pad.

The thermal pad should be applied on the coldplate from its tacky (wet) side. The pad should be applied with its non-tacky side facing up.




2. Ensure the thermal pad is in place and intact.
3. Once the thermal pad is applied to the ASIC interposer, the non-tacky side should be visible on the card's faceplate.




4. Gently peel the liner of the pad's non-tacky side visible on the card's faceplate. Failure to do so may degrade the thermal performance of the product.

➤ Install the adapter into the riser and attach the card to the PCIe x16 slot.

1. Disengage the adapter riser from the blade. Please refer to the blade documentation for instructions.




2. Applying even pressure at both corners of the card, insert the adapter card into the adapter riser until firmly seated. Care must be taken to not harm the black bumpers located on the print side of the card.




- Vertically insert the riser that populates the adapter card into the server blade.

1. Applying even pressure on the riser, gently insert the riser into the server.



2. Secure the riser with the supplied screws. Please refer to the server blade documentation for more information.



Driver Installation

Please use the relevant driver installation section.



ConnectX-6 Socket Direct cards 2x PCIe x16 (OPNs: MCX654106A-HCAT and MCX654106A-ECAT) are not supported in Windows and WinOF-2.

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

Linux Driver Installation

This section describes how to install and test the NVIDIA OFED for Linux package on a single server with a NVIDIA ConnectX-6 adapter card installed.

Prerequisites

Requirements	Description
Platforms	A server platform with a ConnectX-6 InfiniBand/VPI adapter card installed.
Required Disk Space for Installation	1GB
Device ID	For the latest list of device IDs, please visit the NVIDIA website at http://www.nvidia.com/page/firmware_HCA_FW_identification .
Operating System	Linux operating system. For the list of supported operating system distributions and kernels, please refer to the <i>NVIDIA OFED Release Notes</i> file.
Installer Privileges	The installation requires administrator (root) privileges on the target machine.

Downloading NVIDIA OFED

1. Verify that the system has a NVIDIA network adapter installed by running `lspci` command. The below table provides output examples per ConnectX-6 card configuration.

ConnectX-6 Card Configuration	
Single-port Socket Direct Card (2x PCIe x16)	<pre>[root@mftqa-009 ~]# lspci grep mellanox -i a3:00.0 Infiniband controller: Mellanox Technologies MT28908 Family [ConnectX-6] e3:00.0 Infiniband controller: Mellanox Technologies MT28908 Family [ConnectX-6]</pre>
Dual-port Socket Direct Card (2x PCIe x16)	<pre>[root@mftqa-009 ~]# lspci grep mellanox -i 05:00.0 Infiniband controller: Mellanox Technologies MT28908A0 Family [ConnectX-6] 05:00.1 Infiniband controller: Mellanox Technologies MT28908A0 Family [ConnectX-6] 82:00.0 Infiniband controller: Mellanox Technologies MT28908A0 Family [ConnectX-6] 82:00.1 Infiniband controller: Mellanox Technologies MT28908A0 Family [ConnectX-6]</pre> <p>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-6 Socket Direct adapter are viewed as four net devices by the system.</p>
Single-port PCIe x16 Card	<pre>[root@mftqa-009 ~]# lspci grep mellanox -ia 3:00.0 Infiniband controller: Mellanox Technologies MT28908 Family [ConnectX-6]</pre>
Dual-port PCIe x16 Card	<pre>[root@mftqa-009 ~]# lspci grep mellanox -ia 86:00.0 Network controller: Mellanox Technologies MT28908A0 Family [ConnectX-6] 86:00.1 Network controller: Mellanox Technologies MT28908A0 Family [ConnectX-6]</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 <http://www.nvidia.com> > Products > Software > Ethernet Drivers > Linux SW/Drivers > Download..
 - a. Scroll down to the Download wizard, and click the Download tab.
 - b. Choose your relevant package depending on your host operating system.
 - c. Click the desired ISO/tgz package.
 - d. To obtain the download link, accept the End User License Agreement (EULA).
3. Use the md5sum 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.

```
md5sum MLNX_OFED_LINUX-<ver>-<OS label>.iso
```

Installing NVIDIA 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: The firmware will not be updated if you run the install script with the '`--without-fw-update`' option.

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

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

Usage

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



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


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

- If you need to install NVIDIA 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.

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 6.3 under the /tmp directory.

```
# ./MLNX_OFED_LINUX-x.x-x-rhel6.3-x86_64/mlnx_add_kernel_support.sh -m /tmp/MLNX_OFED_LINUX-x.x-x-rhel6.3-x86_64/ --make-tgz
Note: This program will create MLNX_OFED_LINUX TGZ for rhel6.3 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-rhel6.3-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 user space 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, NVIDIA OFED for Linux has already installed the OpenSM Subnet Manager on your machine.

For the list of installation options, run:

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

 The DKMS (on Debian based OS) and the weak-modules (RedHat OS) mechanisms rebuild the initrd/initramfs for the respective kernel in order to add the MLNX_OFED drivers. When installing MLNX_OFED without DKMS support on Debian based OS, or without KMP support on RedHat or any other distribution, the initramfs will not be changed. Therefore, the inbox drivers may be loaded on boot. In this case, openibd service script will automatically unload them and load the new drivers that come with MLNX_OFED.

Installation Procedure

1. Login to the installation machine as root.
2. Mount the ISO image on your machine.

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

3. 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-6
Part Number:     MCX654106A-HCAT
Description:     ConnectX@-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE, dual-port
QSFP56, Socket Direct 2x PCIe3.0 x16, tall bracket
PSID:           MT_2190110032
PCI Device Name: 0b:00.0
Base MAC:       0000e41d2d5cf810
Versions:      Current      Available
FW            16.22.0228    16.22.0228
Status:        Up to date

```

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

Error message 1:

```

Device #1:
-----
Device Type:      ConnectX-6
Part Number:     MCX654106A-HCAT
Description:     ConnectX@-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE, dual-port
QSFP56, Socket Direct 2x PCIe3.0 x16, tall bracket
PSID:           MT_2190110032
PCI Device Name: 0b:00.0
Base MAC:       0000e41d2d5cf810
Versions:      Current      Available
FW            16.22.0228    N/A
Status:        No matching image found

```

Error message 2:

```
The firmware for this device is not distributed inside Mellanox driver: 0000:01:00.0 (PSID:
IBM2150110033)
To obtain firmware for this device, please contact your HW vendor.
```

4. If the installation script has performed a firmware update on your network adapter, complete the step relevant to your adapter card type to load the firmware:
ConnectX-6 Socket Direct - perform a cold reboot (power cycle)
Otherwise, restart the driver by running: `/etc/init.d/openibd restart`

After installation completion, information about the NVIDIA OFED installation, such as prefix, kernel version, and installation parameters can be retrieved by running the command `/etc/infiniband/info`.

Most of the NVIDIA 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.

Installation Results

Software	<ul style="list-style-type: none">• Most of MLNX_OFED packages are installed under the “/usr” directory except for the following packages which are installed under the “/opt” directory:<ul style="list-style-type: none">• fca and ibutils• The kernel modules are installed under<ul style="list-style-type: none">• /lib/modules/`uname -r`/updates on SLES and Fedora Distributions• /lib/modules/`uname -r`/extra/mlnx-ofa_kernel on RHEL and other Red Hat like Distributions
Firmware	<ul style="list-style-type: none">• The firmware of existing network adapter devices will be updated if the following two conditions are fulfilled:<ul style="list-style-type: none">• 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 NVIDIA 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 that your machine has an unsupported network adapter device, no firmware update will occur and the error message below will be printed. <pre>The firmware for this device is not distributed inside Mellanox driver: 0000:01:00.0 (PSID: IBM2150110033) To obtain firmware for this device, please contact your HW vendor.</pre>

Installation Logs

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: `"Logs dir: /tmp/MLNX_OFED_LINUX-<version>.<PD>.logs"`.

Example:

```
Logs dir: /tmp/MLNX_OFED_LINUX-4.4-1.0.0.0.63414.logs
```

openibd Script

As of MLNX_OFED v2.2-1.0.0 the openibd script supports pre/post start/stop scripts:
This can be controlled by setting the variables below in the /etc/infiniband/openibd.conf file.

```
OPENIBD_PRE_START  
OPENIBD_POST_START  
OPENIBD_PRE_STOP  
OPENIBD_POST_STOP
```

Example:

```
OPENIBD_POST_START=/sbin/openibd_post_start.sh
```



An example of OPENIBD_POST_START script for activating all interfaces is provided in the MLNX_OFED package under the docs/scripts/openibd-post-start-configure-interfaces/ folder.

Driver Load Upon System Boot

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

To prevent automatic load of the NVIDIA drivers upon system boot:

1. Add the following lines to the "/etc/modprobe.d/mlnx.conf" file.

```
blacklist mlx4_core  
blacklist mlx4_en  
blacklist mlx5_core  
blacklist mlx5_ib
```

2. Set "ONBOOT=no" in the "/etc/infiniband/openib.conf" file.
3. 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

Return Code	Meaning
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

Uninstalling MLNX_OFED

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

Installing MLNX_OFED Using YUM

This type of installation is applicable to RedHat/OL, Fedora, XenServer Operating Systems.

Setting up MLNX_OFED YUM Repository

1. Log into the installation machine as root.
2. 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
```

3. Download and install NVIDIA GPG-KEY:
The key can be downloaded via the following link: <http://www.nvidia.com/downloads/ofed/RPM-GPG-KEY-Mellanox>

```
# wget http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-Mellanox
--2014-04-20 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
2014-04-20 13:52:30 (247 MB/s) - ?RPM-GPG-KEY-Mellanox? saved [1354/1354]
```

4. 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]:
```

5. 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>)
```

6. 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
```

7. 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:

1. 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-basic.noarch : MLNX_OFED basic installer package (with KMP support)
mlnx-ofed-guest.noarch : MLNX_OFED guest installer package (with KMP support)
mlnx-ofed-hpc.noarch : MLNX_OFED hpc installer package (with KMP support)
mlnx-ofed-hypervisor.noarch : MLNX_OFED hypervisor 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-vpi.noarch : MLNX_OFED vma-vpi installer package (with KMP support)
```

Where:

mlnx-ofed-all	Installs all available packages in MLNX_OFED.
mlnx-ofed-basic	Installs basic packages required for running Mellanox 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.

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 systems' 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)

```

If you have 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.

2. 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!

```

Uninstalling MLNX_OFED Using the YUM Tool

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

Installing MLNX_OFED Using apt-get Tool

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

Setting up MLNX_OFED apt-get Repository

1. Log into the installation machine as root.
2. Extract the MLNX_OFED package on a shared location in your network.
You can download it from <http://www.nvidia.com> > Products > Software > Ethernet Drivers.
3. 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 ./
```

4. Download and install NVIDIA GPG-KEY.

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

5. Check 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
```

6. 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:

1. View the available package groups by invoking:

```
<pre># apt-cache search mlnx-ofed-
mlnx-ofed-vma-eth - MLNX_OFED vma-eth installer package (with DKMS support)
mlnx-ofed-hpc - MLNX_OFED hpc installer package (with DKMS support)
mlnx-ofed-vma-vpi - MLNX_OFED vma-vpi installer package (with DKMS support)
mlnx-ofed-basic - MLNX_OFED basic installer package (with DKMS support)
mlnx-ofed-vma - MLNX_OFED vma installer package (with DKMS support)
mlnx-ofed-all - MLNX_OFED all 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.
```

2. 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`
Or, update the firmware to the latest version available on the NVIDIA website as described in [Updating Adapter Firmware](#).

Uninstalling MLNX_OFED Using the apt-get Tool

Use the script `/usr/sbin/ofed_uninstall.sh` to uninstall the NVIDIA OFED package. The script is part of the `ofed-scripts` package.

Updating Firmware After Installation

The firmware can be updated either manually or automatically (upon system boot), as described in the sections below.

Updating the Device Online

To update the device online on the machine from the NVIDIA site, use the following command line:

```
mlxfwmanager --online -u -d <device>
```

Example:

```
mlxfwmanager --online -u -d 0000:09:00.0
Querying Mellanox devices firmware ...
Device #1:
-----
Device Type:          ConnectX-5
Part Number:
Description:
PSID:                 MT_1020120019
PCI Device Name:     0000:09:00.0
Port1 GUID:          0002c9000100d051
Port2 MAC:           0002c9000002
Versions:             Current      Available
                    FW 2.32.5000   2.33.5000
Status:              Update required
-----
Found 1 device(s) requiring firmware update. Please use -u flag to perform the update.
```

Updating the Device Manually

To update the device manually, please refer to the [OEM Firmware Download page](http://www.nvidia.com/page/firmware_table_dell?mtag=oem_firmware_download) at http://www.nvidia.com/page/firmware_table_dell?mtag=oem_firmware_download.

In case that you ran the mlnxofedinstall script with the '--without-fw-update' option or you are using an OEM card and now you wish to (manually) update firmware on your adapter card(s), you need to perform the steps below. The following steps are also appropriate in case that you wish to burn newer firmware that you have downloaded from the NVIDIA web site (<http://www.nvidia.com> > Support > Firmware Download).

1. Get the device's PSID.

```
mlxfwmanager_pci | grep PSID
PSID: MT_1210110019
```

2. Download the firmware BIN file from the NVIDIA website or the OEM website.
3. Burn the firmware.

```
mlxfwmanager_pci -i <fw_file.bin>
```

4. Reboot your machine after the firmware burning is completed.

Updating the Device Firmware Automatically upon System Boot

As of MLNX_OFED v3.1-x.x.x, firmware can be automatically updated upon system boot. The firmware update package (mlnx-fw-updater) is installed in the “/opt/mellanox/mlnx-fw-updater” folder, and openibd service script can invoke the firmware update process if requested on boot.

If the firmware is updated, the following message is printed to the system’s standard logging file:

```
fw_updater: Firmware was updated. Please reboot your system for the changes to take effect.
```

Otherwise, the following message is printed:

```
fw_updater: Didn't detect new devices with old firmware.
```

Please note, this feature is disabled by default. To enable the automatic firmware update upon system boot, set the following parameter to “yes” “RUN_FW_UPDATER_ONBOOT=yes” in the openibd service configuration file “/etc/infiniband/openib.conf”.

You can opt to exclude a list of devices from the automatic firmware update procedure. To do so, edit the configurations file “/opt/mellanox/mlnx-fw-updater/mlnx-fw-updater.conf” and provide a comma separated list of PCI devices to exclude from the firmware update.

Example:

```
MLNX_EXCLUDE_DEVICES="00:05:0,00:07:0"
```

UEFI Secure Boot

All kernel modules included in MLNX_OFED for RHEL7 and SLES12 are signed with x.509 key to support loading the modules when Secure Boot is enabled.

Enrolling NVIDIA's x.509 Public Key on Your Systems

In order to support loading MLNX_OFED drivers when an OS supporting Secure Boot boots on a UEFI-based system with Secure Boot enabled, the NVIDIA x.509 public key should be added to the UEFI Secure Boot key database and loaded onto the system key ring by the kernel.

Follow these steps below to add the Mellanox's x.509 public key to your system:



Prior to adding the Mellanox's x.509 public key to your system, please make sure that (1) The 'mokutil' package is installed on your system, and (2) The system is booted in UEFI mode.

1. Download the x.509 public key.

```
# wget http://www.mellanox.com/downloads/ofed/mlnx_signing_key_pub.der
```

2. Add the public key to the MOK list using the mokutil utility.

```
# mokutil --import mlnx_signing_key_pub.der
```

3. Reboot the system.

The pending MOK key enrollment request will be noticed by shim.efi and it will launch MokManager.efi to allow you to complete the enrollment from the UEFI console. You will need to enter the password you previously associated with this request and confirm the enrollment. Once done, the public key is added to the MOK list, which is persistent. Once a key is in the MOK list, it will be automatically propagated to the system key ring and subsequent will be booted when the UEFI Secure Boot is enabled.



To see what keys have been added to the system key ring on the current boot, install the 'keyutils' package and run: `#keyctl list %:..system_keyring#`

Removing Signature from kernel Modules

The signature can be removed from a signed kernel module using the 'strip' utility which is provided by the 'binutils' package. The strip utility will change the given file without saving a backup. The operation can be undo only by resigning the kernel module. Hence, we recommend backing up a copy prior to removing the signature.

To remove the signature from the MLNX_OFED kernel modules:

1. Remove the signature.

```
# rpm -qa | grep -E "kernel-ib|mlnx-ofa_kernel|iser|srp|knem|mlnx-rds|mlnx-nfsrdma|mlnx-nvme|mlnx-rdma-rxe" | xargs rpm -ql | grep "\.ko$" | xargs strip -g
```

After the signature has been removed, a message as the below will no longer be presented upon module loading:

```
"Request for unknown module key 'Mellanox Technologies signing key: 61feb074fc7292f958419386ffdd9d5ca999e403' err -11"
```

However, please note that a similar message as the following will still be presented:

```
"my_module: module verification failed: signature and/or required key missing - tainting kernel"
```

This message is only presented once, upon first module boot that either has no signature or whose key is not in the kernel key ring. Therefore, this message may go unnoticed. Once the system is rebooted after unloading and reloading a kernel module, the message will appear. (Note that this message cannot be eliminated.)

2. Update the initramfs on RHEL systems with the stripped modules.

```
mkinitrd /boot/initramfs-$(uname -r).img $(uname -r) --force
```

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](https://community.nvidia.com/docs/DOC-2489) at <https://community.nvidia.com/docs/DOC-2489>.

Windows Driver Installation

For Windows, download and install the latest NVIDIA WinOF-2 for Windows software package available via the NVIDIA web site at: <http://www.nvidia.com> > Products > Software > Ethernet Drivers > Download. Follow the installation instructions included in the download package (also available from the download page).



Windows driver is currently not supported in the following ConnectX-6 OPNs:

- MCX654106A-HCAT
- MCX654106A-ECAT



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 2012 R2	MLNX_WinOF2-2_10_All_x64.exe
Windows Server 2012	
Windows Server 2016	
Windows Server 2019	
Windows 8.1 Client (64 bit only)	
Windows 10 Client (64 bit only)	

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


Downloading NVIDIA 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 NVIDIA WinOF-2 web page at:
<http://www.nvidia.com> > Products > InfiniBand/VPI Drivers => Windows SW/Drivers.
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 NVIDIA 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:

- [.Windows Driver Installation v3.0](#) - An installation procedure that requires frequent user intervention.
- [Unattended Installation](#) - An automated installation procedure that requires no user intervention.

 Both Attended and Unattended installations require administrator privileges.

Attended Installation

The following is an example of an installation session.

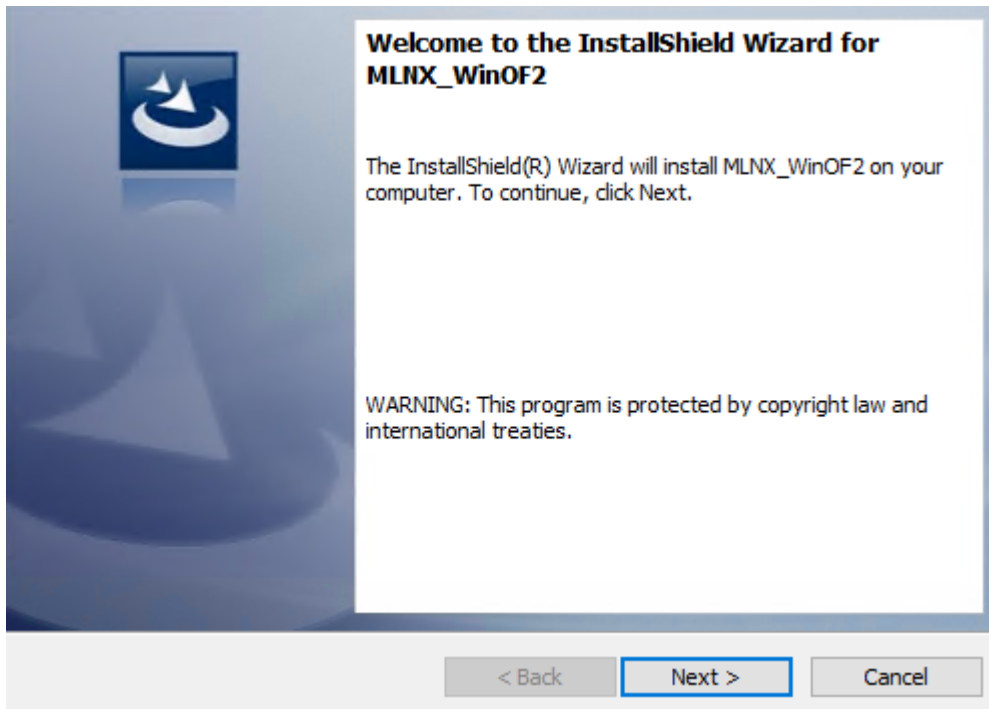
1. Double click the .exe and follow the GUI instructions to install MLNX_WinOF2.
2. [Optional] Manually configure your setup to contain the logs option (replace “LogFile” with the relevant directory):

```
MLNX_WinOF2-[Driver/Version]_<revision_version>_All_Arch.exe /v"/1*vx [LogFile]"
```

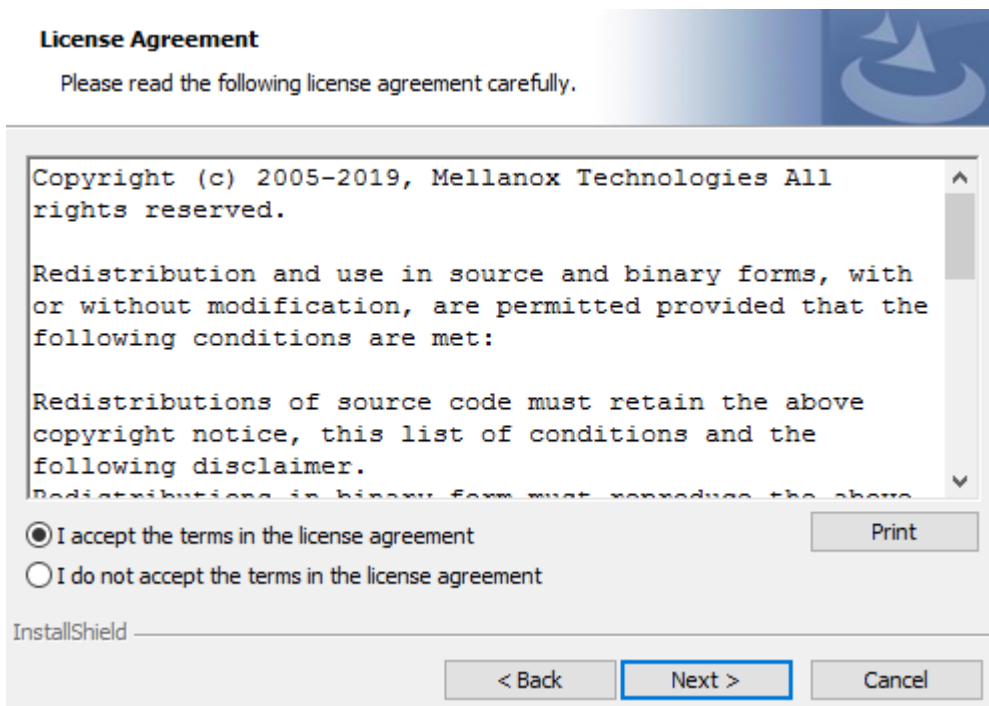
3. [Optional] If you do not want to upgrade your firmware version. (Note: MT_SKIPFWUPGRD default value is False.)

```
MLNX_WinOF2-2_10_50000_All_x64.exe /v"/1*vx MyLog.txt=1"
```


4. Click Next in the Welcome screen.




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



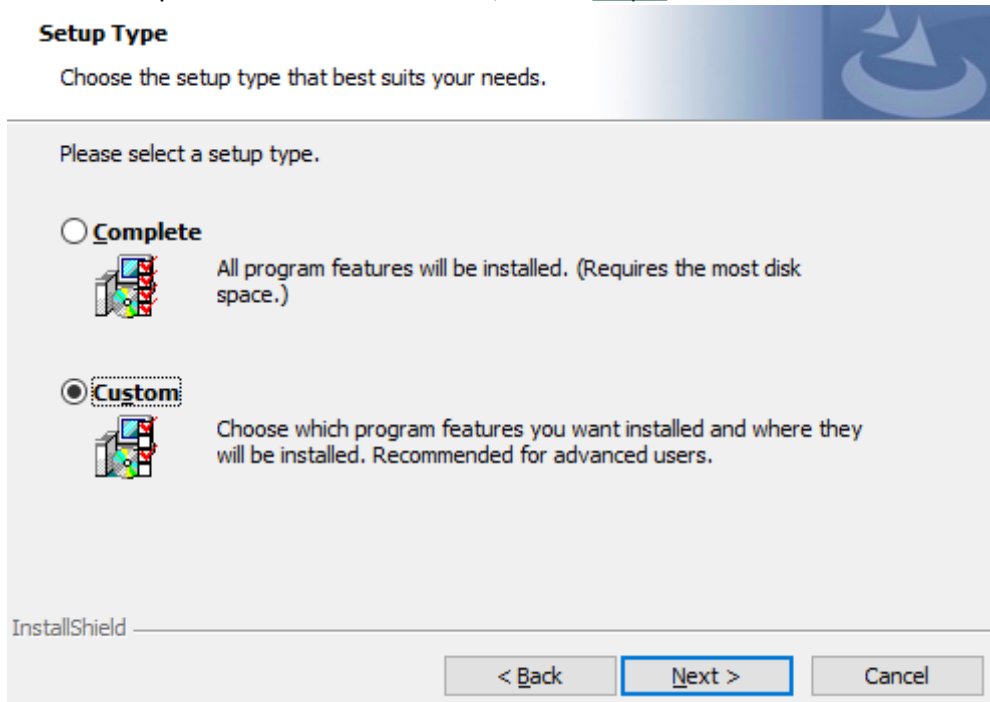
6. Select the target folder for the installation.



7. 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.

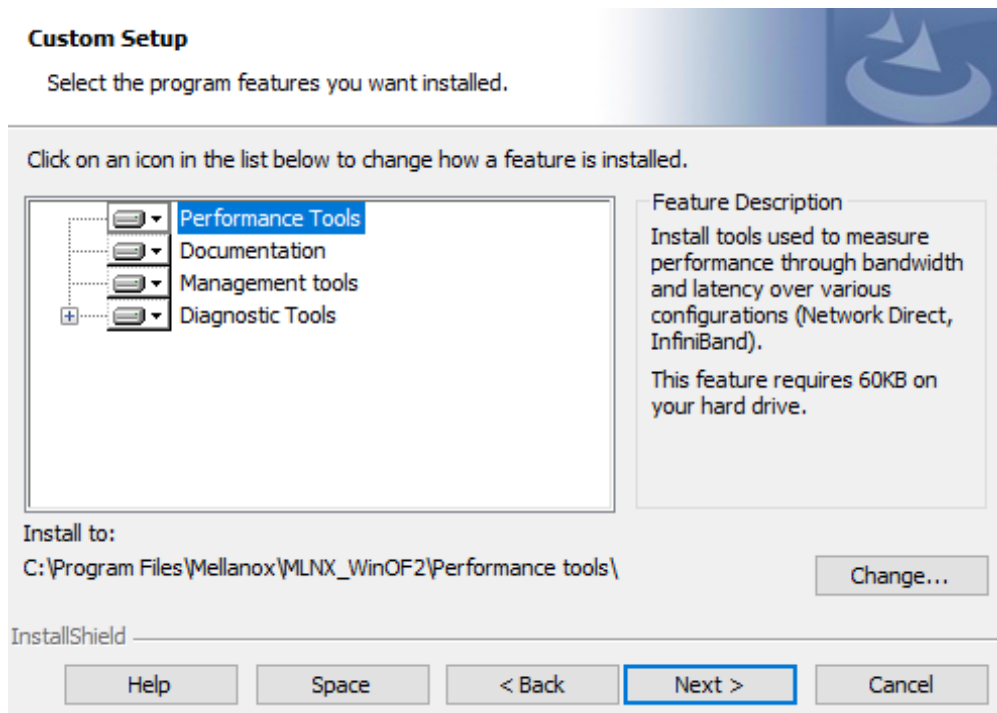


8. Select a Complete or Custom installation, follow [Step a](#) onward.

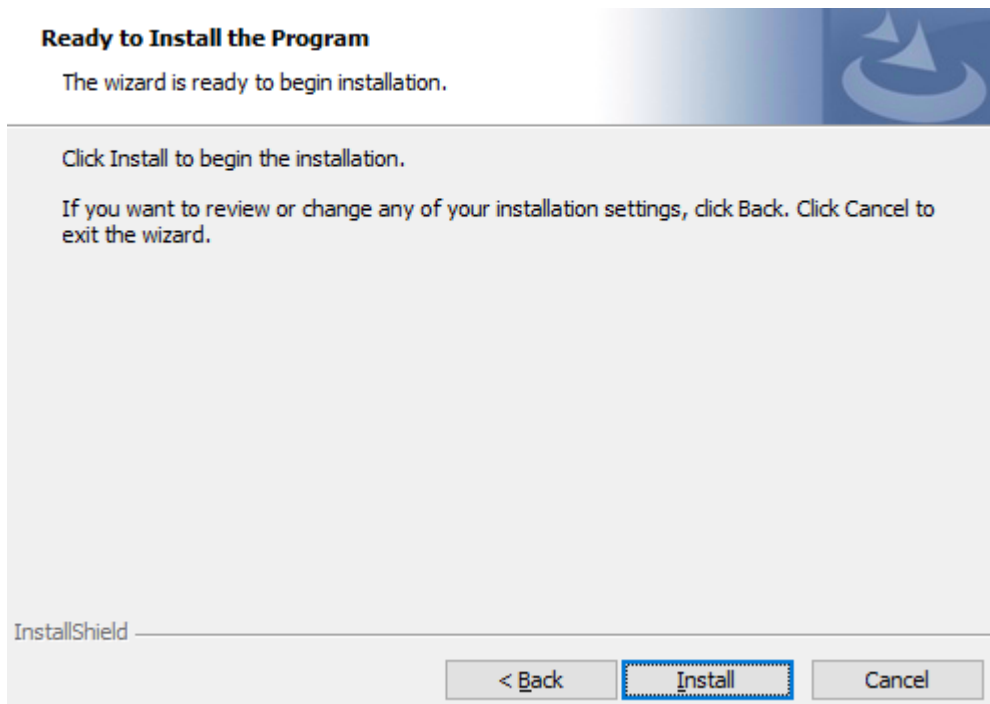



- a. Select the desired feature to install:
- Performances tools - install the performance tools that are used to measure performance in user environment
 - Documentation - contains the User Manual and Release Notes
 - Management tools - installation tools used for management, such as mlxstat
 - Diagnostic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


- b. Click Next to install the desired tools.



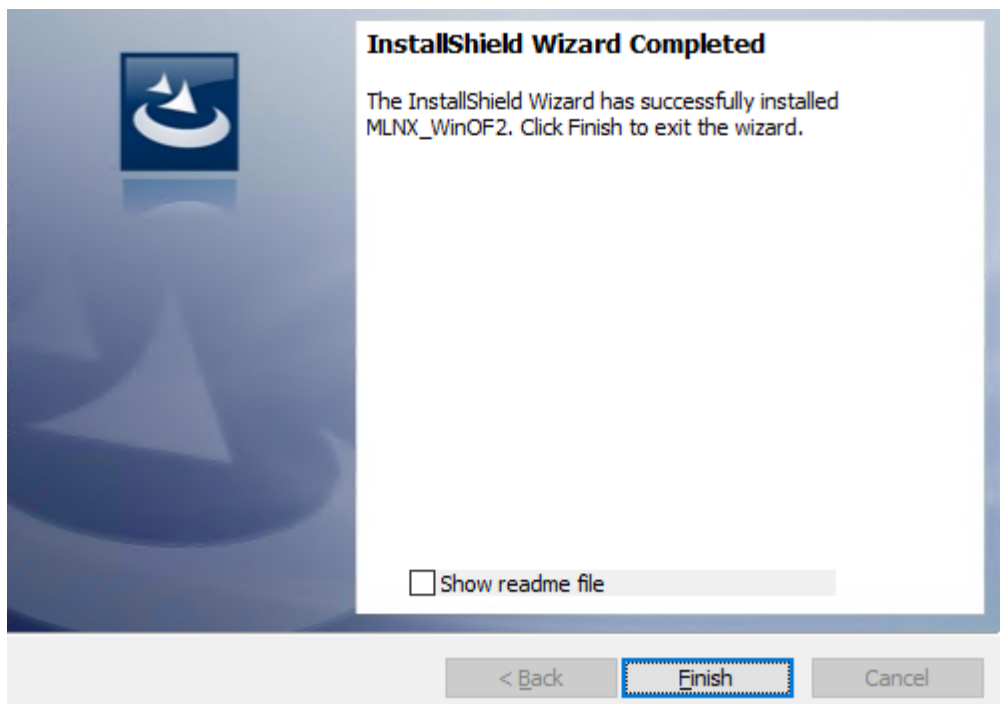
9. Click Install to start the installation.



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



11. 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:

```
_All_Arch.exe /S /v/qn /v"/1*vx [Log- File]" " v:shapes="_x0000_s1026">
```

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`).

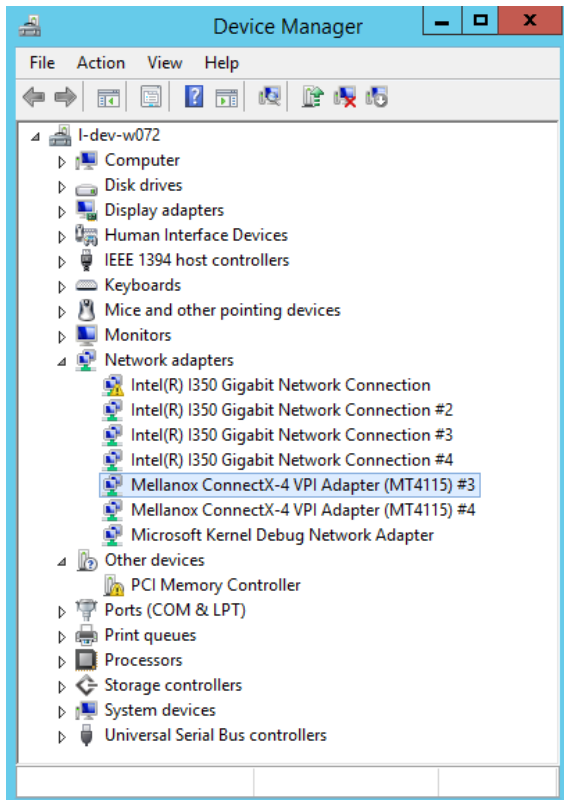
```
MLNX_WinOF2-[Driver/Version]_<revision_version>_All_Arch.exe /vMT_SKIPFWUPGRD=1
```

Installation Results

Upon installation completion, you can verify the successful addition of the network card(s) through the Device Manager. The inf files can be located at:

```
%ProgramFiles%\Mellanox\MLNX_WinOF2\Drivers\
```

To see the NVIDIA network adapters, display the Device Manager and pull down the “Network adapters” menu.



Uninstalling NVIDIA WinOF-2 Driver

Attended Uninstallation

To uninstall MLNX_WinOF2 on a single node:

1. Click Start > Control Panel > Programs and Features > MLNX_WinOF2 > Uninstall.
(NOTE: This requires elevated administrator privileges)

Unattended Uninstallation

To uninstall MLNX_WinOF2 in unattended mode:

1. Open a CMD console. (Click Task Manager > File > Run new task, and enter CMD.)
2. To uninstall the driver, run:

```
MLNX_WinOF2-2_0_A11_x64.exe /S /x /v"/qn"
```

Extracting Files Without Running Installation

➤ To extract the files without running installation, perform the following steps:

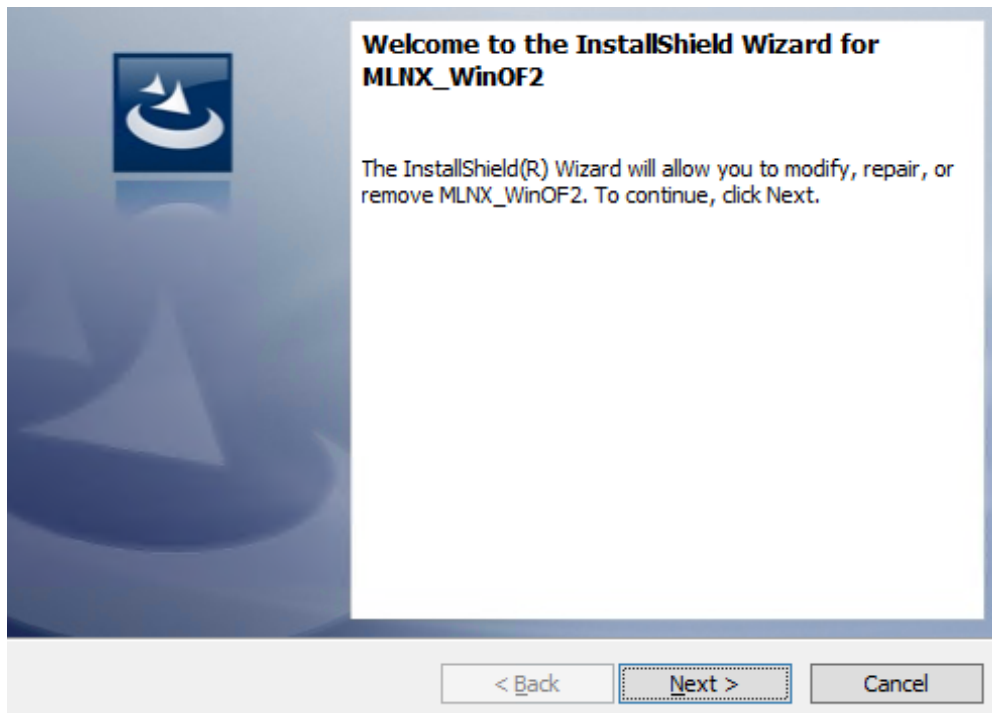
1. Open a CMD console-> Click Start-> Task Manager-> File-> Run new task-> and enter CMD.
2. Extract the driver and the tools:

```
MLNX_WinOF2-2_0_<revision_version>_All_x64 /a
```


To extract only the driver file

```
MLNX_WinOF2-2_0_<revision_version>_All_x64 /a /vMT_DRIVERS_ONLY=1
```

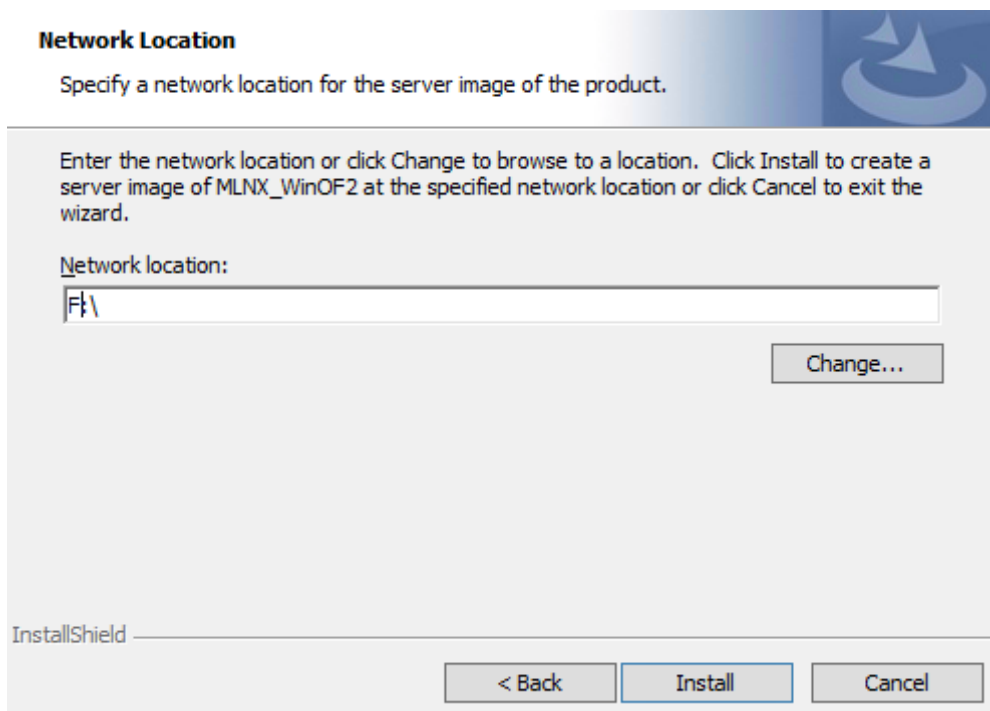
3. Click Next to create a server image.




4. Click Change and specify the location in which the files are extracted to.



5. Click Install to extract this folder, or click Change to install to a different folder.



6. To complete the extraction, click Finish.



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 WinOF-2 package installation.

For information on how to upgrade firmware manually, please refer to the MFT User Manual at www.nvidia.com >Products > Ethernet Drivers > Firmware 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 one of the following NVIDIA devices: <ul style="list-style-type: none">• ConnectX®-6 (InfiniBand/VPI/EN) (firmware: fw-ConnectX6)
Device ID	For the latest list of device IDs, please visit the NVIDIA website.
Operating System	ESXi 6.5

Requirement	Description
Installer Privileges	The installation requires administrator privileges on the target machine.

Installing NVIDIA NATIVE ESXi Driver for VMware vSphere



Please uninstall all previous NVIDIA driver packages prior to installing the new version. See [Removing Earlier NVIDIA Drivers](#) for further information.

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
nmlx5-rdma     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 NVIDIA 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 [NVIDIA 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 NVIDIA 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 from <http://www.nvidia.com> > Products > Software > Firmware Tools.

Firmware Update Example

```
[server1]# ./mlxup
Querying Mellanox devices firmware ...
Device Type:      ConnectX-6
Part Number:      MCX654106A-HCAT
Description:      ConnectX@-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE, dual-port QSFP56, Socket Direct 2x
PCIe3.0 x16, tall bracket
PSID:             MT_2190110032
PCI Device Name:  0000:06:00.0
Base GUID:        e41d2d0300fd8b8a
Versions:         Current      Available
                  FW 16.23.1020  16.24.1000

Status:           Update required

Device Type:      ConnectX-6
Part Number:      MCX654106A-HCAT
Description:      ConnectX@-6 VPI adapter card, HDR IB (200Gb/s) and 200GbE, dual-port QSFP56, Socket Direct 2x
PCIe3.0 x16, tall bracket
PSID:             MT_2170110021
PCI Device Name:  0000:07:00.0
Base MAC:         0000e41d2da206d4
Versions:         Current      Available
                  FW 16.24.1000  16.24.1000

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-yyyyymmdd.log
```

Troubleshooting

General Troubleshooting

Server unable to find the adapter	<ul style="list-style-type: none">• 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	<ul style="list-style-type: none">• 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	<ul style="list-style-type: none">• 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	<ul style="list-style-type: none">• 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	<ul style="list-style-type: none">• Check that the latest driver is loaded• Check that both the adapter and its link are set to the same speed and duplex settings
Event message received of insufficient power	<ul style="list-style-type: none">• When [adapter's current power consumption] > [PCIe slot advertised power limit] - a warning message appears in the server's system even logs (Eg. dmesg: "Detected insufficient power on the PCIe slot")• It's recommended to use a PCIe slot that can supply enough power.• If a message of the following format appears - "mlx5_core 0003:01:00.0: port_module:254:(pid 0): Port module event[error]: module 0, Cable error, One or more network ports have been powered down due to insufficient/unadvertised power on the PCIe slot" please upgrade your Adapter's firmware.• If the message remains - please consider switching from Active Optical Cable (AOC) or transceiver to Direct Attached Copper (DAC) connectivity.

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>
NVIDIA Firmware Tool (MFT)	<p>Download and install MFT: http://www.nvidia.com/content/pages.php?pg=management_tools&menu_section=34 Refer to the User Manual for installation instructions.</p> <p>Once installed, run:</p> <pre>mst start mst status flint -d <mst_device> q</pre>
Ports Information	<pre>ibstat ibv_devinfo</pre>
Firmware Version Upgrade	<p>To download the latest firmware version refer to http://www.nvidia.com/supportdownloader</p>
Collect Log File	<pre>cat /var/log/messages dmesg >> system.log journalctl (Applicable on new operating systems) cat /var/log/syslog</pre>

Windows Troubleshooting

Environment Information	<p>From the Windows desktop choose the Start menu and run: msinfo32</p> <p>To export system information to a text file, choose the Export option from the File menu. Assign a file name and save.</p>
NVIDIA Firmware Tool (MFT)	<p>Download and install MFT: http://www.nvidia.com/content/pages.php?pg=management_tools&menu_section=34 Refer to the User Manual for installation instructions.</p> <p>Once installed, open a CMD window and run:</p> <pre>WinMFT mst start mst status flint -d <mst_device> q</pre>
Ports Information	<pre>vstat</pre>
Firmware Version Upgrade	<p>Download the latest firmware version using the PSID/board ID: http://www.nvidia.com/supportdownloader/ flint -d <mst_device> -i <firmware_bin_file> b</p>

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

MCX651105A-EDAT Specifications



Please make sure to install the ConnectX-6 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.

Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm)		
	Connector: Single QSFP56 InfiniBand and Ethernet (copper and optical)		
Protocol Support	InfiniBand: IBTA v1.4 ^a		
	Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), HDR (50Gb/s per lane) port		
	Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR		
	Data Rate	InfiniBand	SDR/DDR/QDR/FDR/EDR/HDR100
		Ethernet	1/10/25/40/50/100 Gb/s
	PCI Express Gen3/4:	SERDES @ 8.0GT/s/16GT/s, x8 lanes (2.0 and 1.1 compatible)	
Adapter Card Power	Voltage: 12V, 3.3VAUX		
	Power	Cable	
	Typical Power^b	Passive Cables	10.1W
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVONline login credentials)	
	Voltage: 3.3Aux Maximum current: 100mA		
	Maximum power available through QSFP56 port: 5W		
Environmental	Temperature	Operational	0° C to 55° C
		Non-operational	-40° C to 70° C

	Humidity: 90% relative humidity ^ε		
	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction
			Heatsink to Port
		Passive Cables	TBD
NVIDIA Active W Cables	TBD	TBD	
Regulatory	Safety: CB / cTUVus / CE		
	EMC: CE / FCC / VCCI / ICES / RCM / KC		
	RoHS: RoHS Compliant		

MCX653105A-HDAT Specifications



Please make sure to install the ConnectX-6 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.

Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm)
	Connector: Single QSFP56 InfiniBand and Ethernet (copper and optical)
Protocol Support	InfiniBand: IBTA v1.4 ^a
	Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), HDR (50Gb/s per lane) port
	Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR

	Data Rate	InfiniBand	SDR/DDR/QDR/FDR/EDR/HDR100/HDR		
		Ethernet	1/10/25/40/50/100/200 Gb/s		
	PCI Express Gen3/4:	SERDES @ 8.0GT/s/16GT/s, x16 lanes (2.0 and 1.1 compatible)			
Adapter Card Power	Voltage: 12V, 3.3VAUX				
	Power	Cable			
	Typical Power^b	Passive Cables	19.3W		
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVONline login credentials)			
	Voltage: 3.3Aux Maximum current: 100mA				
	Maximum power available through QSFP56 port: 5W				
Environmental	Temperature	Operational	0°C to 55°C		
		Non-operational	-40°C to 70°C		
	Humidity: 90% relative humidity ^c				
	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction		
			Heatsink to Port	Port to Heatsink	
	Passive Cables	350 LFM / 55°C	250 LFM / 35°C		
	NVIDIA Active 4.7W Cables	500 LFM / 55°C ^d	250 LFM / 35°C		
Regulator y	Safety: CB / cTUVus / CE				
	EMC: CE / FCC / VCCI / ICES / RCM / KC				
	RoHS: RoHS Compliant				

- Notes: a. The ConnectX-6 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.
- d. For engineering samples - add 250LFM

MCX653106A-HDAT Specifications



Please make sure to install the ConnectX-6 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.

Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm)		
	Connector: Dual QSFP56 InfiniBand and Ethernet (copper and optical)		
Protocol Support	InfiniBand: IBTA v1.4 ^a Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), HDR (50Gb/s per lane) port		
	Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR		
	Data Rate	InfiniBand	SDR/DDR/QDR/FDR/EDR/HDR100/HDR
		Ethernet	1/10/25/40/50/100/200 Gb/s
PCI Express Gen3/4: SERDES @ 8.0GT/s/16GT/s, x16 lanes (2.0 and 1.1 compatible)			
Adapter Card Power	Voltage: 12V, 3.3VAUX		
	Power	Cable	
	Typical Power^b	Passive Cables	23.6W
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVONline login credentials)	
	Voltage: 3.3Aux		
	Maximum current: 100mA		
Maximum power available through QSFP56 port: 5W			

Environmental	Temperature	Operational	0° C to 55° C		
		Non-operational	-40° C to 70° C		
	Humidity: 90% relative humidity ^c				
	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction		
			Heatsink to Port	Port to Heatsink	
Passive Cables		400 LFM / 55° C	300 LFM / 35° C		
NVIDIA Active 4.7W Cables	950 LFM / 55° C 600 LFM / 48° Cd	300 LFM / 35° C			
Regulatory	Safety: CB / cTUVus / CE				
	EMC: CE / FCC / VCCI / ICES / RCM / KC				
	RoHS: RoHS Compliant				

Notes: a. The ConnectX-6 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.

MCX653105A-HDAL Specifications



Please make sure to install the ConnectX-6 card in an liquid-cooled Intel® Server System D50TNP platform.


Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm)
	Connector: Single QSFP56 InfiniBand and Ethernet (copper and optical)

Protocol Support	InfiniBand: IBTA v1.4 ^a Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), HDR (50Gb/s per lane) port			
	Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR			
	Data Rate	InfiniBand	SDR/DDR/QDR/FDR/EDR/HDR100/HDR	
		Ethernet	1/10/25/40/50/100/200 Gb/s	
	PCI Express Gen3/4: SERDES @ 8.0GT/s/16GT/s, x16 lanes (2.0 and 1.1 compatible)			
Adapter Card Power	Voltage: 12V, 3.3VAUX			
	Power	Cable		
	Typical Power ^b	Passive Cables	18.5W	
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVOnline login credentials)		
	Voltage: 3.3Aux Maximum current: 100mA			
	Maximum power available through QSFP56 port: 5W			
Environmental	Temperature	Operational	0°C to 55°C	
		Non-operational	-40°C to 70°C	
	Humidity: 90% relative humidity ^c			
	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction	
			Heatsink to Port	Port to Heatsink
Passive Cables		TBD	TBD	
NVIDIA Active 4.7W Cables	TBD	TBD		
Regulatory	Safety: CB / cTUVus / CE			
	EMC: CE / FCC / VCCI / ICES / RCM / KC			

	RoHS: RoHS Compliant
--	-----------------------------

Notes: a. The ConnectX-6 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.

MCX653106A-HDAL Specifications

 Please make sure to install the ConnectX-6 card in an liquid-cooled Intel® Server System D50TNP platform.

Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm)		
	Connector: Dual QSFP56 InfiniBand and Ethernet (copper and optical)		
Protocol Support	InfiniBand: IBTA v1.4 ^a Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), HDR (50Gb/s per lane) port		
	Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR		
	Data Rate	InfiniBand	SDR/DDR/QDR/FDR/EDR/HDR100/HDR
		Ethernet	1/10/25/40/50/100/200 Gb/s
	PCI Express Gen3/4: SERDES @ 8.0GT/s/16GT/s, x16 lanes (2.0 and 1.1 compatible)		
Adapter Card Power	Voltage: 12V, 3.3VAUX		
	Power	Cable	
	Typical Power^b	Passive Cables	20.85W
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVOnline login credentials)	
	Voltage: 3.3Aux Maximum current: 100mA		
	Maximum power available through QSFP56 port: 5W		

Environmental	Temperature	Operational	0°C to 55°C	
		Non-operational	-40°C to 70°C	
	Humidity: 90% relative humidity ^c			
	TB Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction	
			Heatsink to Port	Port to Heatsink
Passive Cables		TBD	TBD	
	NVIDIA Active 4.7W Cables	TBD	TBD	
Regulatory	Safety: CB / cTUVus / CE			
	EMC: CE / FCC / VCCI / ICES / RCM / KC			
	RoHS: RoHS Compliant			

- Notes: a. The ConnectX-6 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.

MCX653105A-ECAT Specifications



Please make sure to install the ConnectX-6 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.

Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm)
	Connector: Single QSFP56 InfiniBand and Ethernet (copper and optical)
Protocol Support	InfiniBand: IBTA v1.4 ^a Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane)
	Ethernet: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR

	Data Rate	InfiniBand	SDR/DDR/QDR/FDR/EDR/HDR100		
		Ethernet	1/10/25/40/50/100 Gb/s		
PCIe Gen3/4: SERDES @ 8.0GT/s/16GT/s, x16 lanes (2.0 and 1.1 compatible)					
Adapter Card Power	Voltage: 12V, 3.3VAUX				
	Power	Cable			
	Typical Power^b	Passive Cables	15.6W		
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVOnline login credentials)			
	Voltage: 3.3Aux Maximum current: 100mA				
	Maximum power available through QSFP56 port: 5W				
Environmental	Temperature	Operational	0°C to 55°C		
		Non-operational	-40°C to 70°C		
	Humidity: 90% relative humidity^c				
	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction		
			Heatsink to Port		Port to Heatsink
Passive Cables			300 LFM / 55°C	200 LFM / 35°C	
	NVIDIA Active 2.7W Cables	300 LFM / 55°C	200 LFM / 35°C		
Regulatory	Safety: CB / cTUVus / CE				
	EMC: CE / FCC / VCCI / ICES / RCM / KC				
	RoHS: RoHS Compliant				

- Notes: a. The ConnectX-6 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.

MCX653106A-ECAT Specifications



Please make sure to install the ConnectX-6 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.



For power specifications when using a single-port configuration, please refer to [MCX653105A-ECAT Specifications](#)

Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm)		
	Connector: Dual QSFP56 InfiniBand and Ethernet (copper and optical)		
Protocol Support	InfiniBand: IBTA v1.4 ^a		
	Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane) port		
	Ethernet: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR		
	Data Rate	InfiniBand	SDR/DDR/QDR/FDR/EDR
		Ethernet	1/10/25/40/50/100 Gb/s
Gen3/4: SERDES @ 8.0GT/s/16GT/s, x16 lanes (2.0 and 1.1 compatible)			
Adapter Card Power	Voltage: 12V, 3.3VAUX		
	Power	Cable	
	Typical Power^b	Passive Cables	21.0W
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVOnline login credentials)	
	Voltage: 3.3Aux Maximum current: 100mA		

	Maximum power available through QSFP56 port: 5W		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	Humidity: 90% relative humidity[€]		
	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction
Heatsink to Port			Port to Heatsink
Passive Cables		350 LFM / 55°C	250 LFM / 35°C
NVIDIA Active 2.7W Cables	550 LFM / 55°C	250 LFM / 35°C	
Regulatory	Safety: CB / cTUVus / CE		
	EMC: CE / FCC / VCCI / ICES / RCM / KC		
	RoHS: RoHS Compliant		

- Notes: a. The ConnectX-6 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.

MCX654105A-HCAT Specifications



Please make sure to install the ConnectX-6 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.

Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm) 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) Length: 35cm
-----------------	--

	Connector: Single QSFP56 InfiniBand and Ethernet (copper and optical)					
Protocol Support	InfiniBand: IBTA v1.4 ^a					
	Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), HDR (50Gb/s per lane) port					
	Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR					
	Data Rate	<table border="1"> <tr> <td>InfiniBand</td> <td>SDR/DDR/QDR/FDR/EDR/HDR100/HDR</td> </tr> <tr> <td>Ethernet</td> <td>1/10/25/40/50/100/200 Gb/s</td> </tr> </table>	InfiniBand	SDR/DDR/QDR/FDR/EDR/HDR100/HDR	Ethernet	1/10/25/40/50/100/200 Gb/s
	InfiniBand	SDR/DDR/QDR/FDR/EDR/HDR100/HDR				
Ethernet	1/10/25/40/50/100/200 Gb/s					
Gen3: SERDES @ 8.0GT/s/, x16 lanes (2.0 and 1.1 compatible)						
Adapter Card Power	Voltage: 12V, 3.3VAUX					
	Power	Cable				
	Typical Power^b	Passive Cables	27.1W			
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVONline login credentials)				
	Voltage: 3.3Aux Maximum current: 100mA					
	Maximum power available through QSFP56 port: 5W					
Active Auxiliary PCIe Connection Card Power	Typical Power	3.0W				
	Maximum Power	4.0W				
Environmental	Temperature	Operational	0°C to 55°C			
		Non-operational	-40°C to 70°C			
	Humidity: 90% relative humidity ^c					

	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction	
			Heatsink to Port	Port to Heatsink
			Passive Cables	600 LFM / 55° C
NVIDIA Active 4.7W Cables	600 LFM / 55° C ^d	350 LFM / 35° C		

Regulatory	Safety: CB / cTUVus / CE
	EMC: CE / FCC / VCCI / ICES / RCM / KC
	RoHS: RoHS Compliant

Notes: a. The ConnectX-6 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.

d. For engineering samples - add 250LFM

MCX654106A-HCAT Specifications



Please make sure to install the ConnectX-6 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.



For power specifications when using a single-port configuration, please refer to [MCX654105A-HCAT Specifications](#)

Physical	Low Profile Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm) 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) Length: 35cm
	Connector: Dual QSFP56 InfiniBand and Ethernet (copper and optical)

Protocol Support	InfiniBand: IBTA v1.4 ^a		
	Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), HDR (50Gb/s per lane) port		
	Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBAConnectX-6SE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR		
	Data Rate	SDR/DDR/QDR/FDR/EDR/HDR100/HDR	
		1/10/25/40/50/100/200 Gb/s	
Gen3: SERDES @ 8.0GT/s, x16 lanes (2.0 and 1.1 compatible)			
Adapter Card Power	Voltage: 12V, 3.3VAUX		
	Power	Cable	
	Typical Power^b	Passive Cables	31.4W
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVONline login credentials)	
	Voltage: 3.3Aux Maximum current: 100mA		
	Maximum power available through QSFP56 port: 5W		
Active Auxiliary PCIe Connection Card Power	Typical Power	3.0W	
	Maximum Power	4.0W	
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	Humidity: 90% relative humidity _c		

	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction	
			Heatsink to Port	Port to Heat sink
		Passive Cables	700 LFM / 55°C	400 LFM / 35°C
NVIDIA Active 4.7W Cables	1050 LFM / 55°C 600 LFM / 48°C	400 LFM / 35°C		
Regulatory	Safety: CB / cTUVus / CE			
	EMC: CE / FCC / VCCI / ICES / RCM// KC			
	RoHS: RoHS Compliant			

Notes: a. The ConnectX-6 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.

MCX654106A-ECAT Specifications



Please make sure to install the ConnectX-6 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.

Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm) 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) Length: 35cm
	Connector: Dual QSFP56 InfiniBand and Ethernet (copper and optical)
Protocol Support	InfiniBand: IBTA v1.4 ^a Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane)

	Ethernet: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR		
	Data Rate	InfiniBand	SDR/DDR/QDR/FDR/HDR100/EDR
		Ethernet	1/10/25/40/50/100 Gb/s
	Gen3: SERDES @ 8.0GT/s, x16 lanes (2.0 and 1.1 compatible)		
Adapter Card Power	Voltage: 12V, 3.3VAUX		
	Power	Cable	
	Typical Power^b	Passive Cables	27.1W
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVONline login credentials)	
	Voltage: 3.3Aux Maximum current: 100mA		
	Maximum power available through QSFP56 port: 5W		
Active Auxiliary PCIe Connection Card Power	Typical Power	3.0W	
	Maximum Power	4.0W	
Environmental	Temperature	Operational	0 °C to 55 °C
		Non-operational	-40 °C to 70 °C
	Humidity: 90% relative humidity ^c		
	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction
Heatsink to Port			Port to Heat sink
	Passive Cables	600 LFM / 55 °C	400 LFM / 35 °C

		NVIDIA Active 2.7W Cables	700 LFM / 55° C	400 LFM / 35° C
Regulatory	Safety: CB / cTUVus / CE			
	EMC: CE / FCC / VCCI / ICES / RCM / KC			
	RoHS: RoHS Compliant			

- Notes: a. The ConnectX-6 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.

MCX653105A-EFAT Specifications



Please make sure to install the ConnectX-6 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.

Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm)		
	Connector: Single QSFP56 InfiniBand and Ethernet (copper and optical)		
Protocol Support	InfiniBand: IBTA v1.4 ^a Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane) port		
	Ethernet: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR		
	Data Rate	InfiniBand	SDR/DDR/QDR/FDR/EDR/HDR100
		Ethernet	1/10/25/40/50/100 Gb/s
	Gen3/4: SERDES @ 8.0GT/s/16GT/s, x16 lanes, Socket Direct (2.0 and 1.1 compatible)		
Adapter Card Power	Voltage: 12V, 3.3VAUX		
	Power	Cable	

	Typical Power^b	Passive Cables	19.4W		
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVONline login credentials)			
	Voltage: 3.3Aux Maximum current: 100mA				
	Maximum power available through QSFP56 port: 5W				
Environmental	Temperature	Operational	0°C to 55°C		
		Non-operational	-40°C to 70°C		
	Humidity: 90% relative humidity ^c				
	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction		
			Heatsink to Port	Port to Heatsink	
Passive Cables			300 / 55°C	200 / 35°C	
NVIDIA Active 2.75W Cables	300 / 55°C	200 / 35°C			
Regulatory	Safety: CB / cTUVus / CE				
	EMC: CE / FCC / VCCI / ICES / RCM / KC				
	RoHS: RoHS Compliant				

Notes: a. The ConnectX-6 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.

MCX653106A-EFAT Specifications



Please make sure to install the ConnectX-6 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.



For power specifications when using a single-port configuration, please refer to [MCX653105A-EFAT Specifications](#).

Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm)		
	Connector: Dual QSFP56 InfiniBand and Ethernet (copper and optical)		
Protocol Support	InfiniBand: IBTA v1.4 ^a Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane) port		
	Ethernet: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR		
	Data Rate	InfiniBand	SDR/DDR/QDR/FDR/EDR, HDR100
		Ethernet	1/10/25/40/50/100 Gb/s
	Gen3/4: SERDES @ 8.0GT/s/16GT/s, x16 lanes, Socket Direct (2.0 and 1.1 compatible)		
Adapter Card Power	Voltage: 12V, 3.3VAUX		
	Power	Cable	
	Typical Power^b	Passive Cables	21.6W
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVOnline login credentials)	
	Voltage: 3.3Aux Maximum current: 100mA		
	Maximum power available through QSFP56 port: 5W		

Environmental	Temperature	Operational	0 °C to 55 °C		
		Non-operational	-40 °C to 70 °C		
	Humidity: 90% relative humidity ^c				
	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction		
			Heatsink to Port	Port to Heat sink	
Passive Cables		350 LFM / 55 °C	250 LFM / 35 °C		
NVIDIA Active 2.75W Cables	550 LFM / 55 °C	250 LFM / 35 °C			
Regulatory	Safety: CB / cTUVus / CE				
	EMC: CE / FCC / VCCI / ICES / RCM / KC				
	RoHS: RoHS Compliant				

Notes: a. The ConnectX-6 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.

MCX683105AN-HDAT Specifications



Please make sure to install the ConnectX-6 card in a PCIe slot that is capable of supplying the required power and airflow as stated in the below table.


Physical	Adapter Card Size: 5.59 in. x 2.71 in. (142mm x 68.90mm)
	Connector: Single QSFP InfiniBand (copper and optical)

Protocol Support	InfiniBand: IBTA v1.4^a HDR (4 lanes x 50Gb/s per lane) , HDR100 (2 lane x 50Gb/s per lane), EDR (4 lanes x 25Gb/s per lane), FDR (4 lanes x 14.0625Gb/s), 1X/2X/4X SDR (4 lanes x 2.5Gb/s per lane).			
	Data Rate	InfiniBand	HDR/HDR100/EDR/FDR/SDR	
	Gen3/4: SERDES @ 8.0GT/s/16GT/s, x16 lanes			
Adapter Card Power	Voltage: 12V, 3.3VAUX			
	Power	Cable		
	Typical Power	Passive Cables	19.58W	
	Maximum Power	Please refer to ConnectX-6 VPI Power Specifications (requires NVONline login credentials)		
	Voltage: 3.3Aux Maximum current: 100mA			
	Maximum power available through QSFP port: 5W			
Environmental	Temperature	Operational	0° C to 55° C	
		Non-operational	-40° C to 70° C	
	Humidity: 90% relative humidity			
	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction	
			Heatsink to Port	Port to Heat sink
			Passive Cables	TBD
NVIDIA Active 2.75W Cables	TBD	TBD		
Regulatory	Safety: CB / cTUVus / CE			
	EMC: CE / FCC / VCCI / ICES / RCM / KC			
	RoHS: RoHS Compliant			





Notes: a. The ConnectX-6 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.



Adapter Card and Bracket Mechanical Drawings and Dimensions

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

Adapter Cards

ConnectX-6 PCIe x16 Adapter Card	ConnectX-6 PCIe x8 Adapter Card
	
MCX683105AN-HDAT	Auxiliary PCIe Connection Card
	

Brackets Dimensions

Tall Bracket	Short Bracket
 <p>Technical drawing of a Tall Bracket. The drawing shows a rectangular component with a central display area and two rows of buttons on either side. The overall length is dimensioned as 120.00mm. The height of the main body is dimensioned as 21.6mm. The part number MEC00000X-YY-MM is printed at the bottom center.</p>	 <p>Technical drawing of a Short Bracket. The drawing shows a similar component to the Tall Bracket but shorter. The overall length is dimensioned as 79.3. The height of the main body is dimensioned as 18.4. The part number MEC00000X-YY-MM is printed at the bottom center.</p>

PCI Express Pinouts Description for Single-Slot Socket Direct Card



This section applies to ConnectX-6 single-slot cards (MCX653105A-EFAT and MCX653106A-EFAT).


ConnectX-6 single-slot Socket Direct cards offer 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. The PCIe x16 interface is split into two PCIe x8 in a row, such that each of the PCIe x8 lanes can be connected to a dedicated CPU in a dual-socket server. In such a configuration, Socket Direct brings lower latency and lower CPU utilization as the direct connection from each CPU to the network means the Interconnect can bypass a QPI (UPI) and the other CPU, optimizing performance and improving latency. CPU utilization is improved as each CPU handles only its own traffic and not traffic from the other CPU.

In order to allow this capability, a system with a special PCI Express x16 slot is required. Table 31 provides the pin definitions of the required four special PCIe pins.

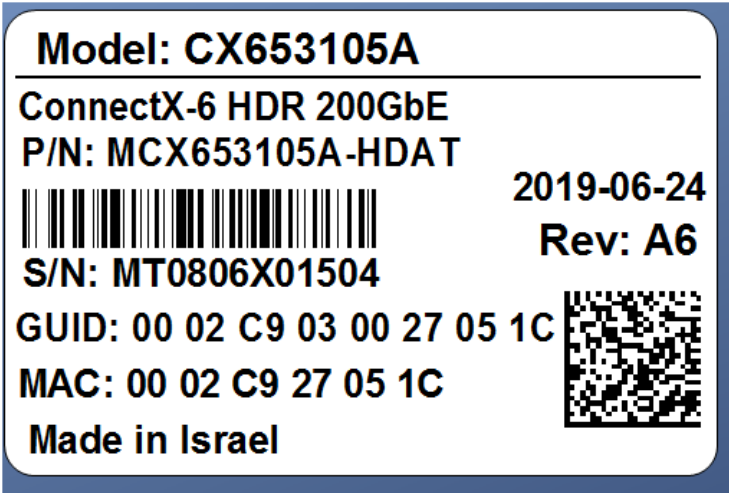
PCIe pin #	Server Connection for 2x PCIe x8 mode	Server Connection for 1x PCIe x16 mode	Comments
B82	1Kohm pull up to 3.3VAUX	Either leave unconnected or pulled to GND	Configure the card to work at 2x PCIe x8 or 1x PCIe x16 modes
A32	P signal of differential PCIe clock (100MHz nominally) of the CPU which connects to PCIe lanes 15-8 of the PCIe connector	Not Connected	PCIe clock for ConnectX-6 PCIe bus lanes [15:8]
A33	N signal of differential PCIe clock (100MHz nominally) of the CPU which connects to PCIe lanes 15-8 of the PCIe connector	Not Connected	
A50	PERST signal from the CPU which connects to PCIe lanes 15-8 of the PCIe connector	Not Connected	PERST (PCIe Reset) for ConnectX-6 PCIe bus lanes [15:8]

Finding the GUID/MAC on the Adapter Card

Each NVIDIA 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. VPI cards have both a GUID and a MAC (derived from the GUID).

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

MCX653105A-HDAT Board Label (Example)



MCX683105AN-HDAT Board Label (Example)



Document Revision History

Date	Revision	Comments/Changes
Aug. 2022	3.3	<ul style="list-style-type: none"> Added mechanical drawing and board label example for MCX683105AN-HDAT. Added a note about FRU EEPROM memory component under the Features and Benefits table.
Jul. 2022	3.2	Updated the "SMBus Interface" under Interfaces .
Jan. 2022	3.1	Added MCX683105AN-HDAT.
Mar. 2021	3.0	Updated "Troubleshooting".
Dec. 2020	2.9	Updated installation instructions.
Dec. 2020	2.8	Added MCX653105A-HDAL and MCX653106A-HDAL support across the document.
Mar. 2020	2.7	Added MCX651105A-EDAT support across the document.
Sep. 2019	2.6	Added a note to the hardware installation instructions.
Aug. 2019	2.5	Updated "Package Contents" and "Hardware Installation"
Aug. 2019	2.4	Updated "PCI Express Pinouts Description".
Aug. 2019	2.3	Updated "Hardware Installation".
Jul. 2019	2.2	Updated "Linux Driver" and "Identifying the card in the system" to include lspci command output examples.
Jun. 2019	2.1	<ul style="list-style-type: none"> Added MCX653105A-HDAT and MCX654105A-HCAT to the UM. Updated "LED Interfaces".
Jun. 2019	2.0	<ul style="list-style-type: none"> Added a note to "Windows Driver Installation". Added short and tall brackets dimensions.
May. 2019	1.9	<ul style="list-style-type: none"> Added mechanical drawings to "Specifications". Updated PCB mechanical tolerance in "Specifications".
May. 2019	1.8	<ul style="list-style-type: none"> Updated "LEDs Interface" specifications. Updated PCB mechanical tolerance in "Specifications".
Apr. 2019	1.7	<ul style="list-style-type: none"> Migrated to on-line format; minor reorganization. Added a note to "Introduction"
Feb. 2019	1.6	Updated "Specifications"
Feb. 2019	1.5	Updated "Specifications"
Jan. 2019	1.4	<ul style="list-style-type: none"> Updated "Airflow Specifications" Added a note to "Installation Instructions"
Dec. 2018	1.3	Updated "Airflow Specifications"

Date	Revision	Comments/Changes
Dec. 2018	1.2	<ul style="list-style-type: none">• Updated “Hardware Requirements”• Updated “Product Overview”
Nov. 2018	1.1	Updated “Hardware Requirements”
Oct. 2018	1.0	First release

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

© 2022 NVIDIA Corporation & affiliates. All Rights Reserved.

