



NVIDIA ConnectX-6 Ethernet Adapter Cards User Manual

Table of Contents

- Introduction..... 8**
 - Product Overview 8
 - ConnectX-6 PCIe x16 Card 8
 - ConnectX-6 Socket Direct Cards (2x PCIe x16) 9
- Features and Benefits 11
- Operating Systems/Distributions 13
- Connectivity 13
- Manageability 14
- Interfaces 15**
 - Ethernet QSFP56 Interfaces 15
 - PCI Express Interface 15
 - LED Interface 15
 - SMBus Interface 16
 - Voltage Regulators 16
- Hardware Installation 17**
 - Safety Warnings 17
 - Installation Procedure Overview..... 17
 - System Requirements 18
 - Hardware Requirements 18

Airflow Requirements	19
Software Requirements	20
Safety Precautions	20
Pre-Installation Checklist	21
Bracket Replacement Instructions	21
Installation Instructions	22
Cables and Modules	23
Identifying the Card in Your System	23
ConnectX-6 PCIe x16 Installation Instructions	26
Installing the Card	26
Uninstalling the Card	28
ConnectX-6 Socket Direct (2x PCIe x16) Installation Instructions	30
Installing the Card	31
Uninstalling the Card	41
Driver Installation	44
Linux Driver Installation	44
Prerequisites	44
Installing MLNX_OFED	45
Installation Script	45
Installation Procedure	47
Installation Results	50

Installation Logging	51
Driver Load Upon System Boot	51
mlnxofedinstall Return Codes	52
Installation Logging	53
Uninstalling MLNX_OFED	53
Additional Installation Procedures	54
Installing MLNX_OFED Using YUM	54
Installing MLNX_OFED Using apt-get	60
Performance Tuning	64
Windows Driver Installation	64
Software Requirements	65
Downloading WinOF-2 Driver	65
Installing WinOF-2 Driver	66
Attended Installation	66
Unattended Installation	76
Firmware Upgrade	78
VMware Driver Installation	78
Hardware and Software Requirements	78
Installing NATIVE ESXi Driver for VMware vSphere	79
Removing Earlier NVIDIA Drivers	79
Firmware Programming	80

Updating Adapter Firmware	81
Troubleshooting	83
General Troubleshooting	83
Linux Troubleshooting	84
Windows Troubleshooting.....	85
Specifications.....	86
MCX614106A-CCAT Specifications	86
MCX614105A-VCAT Specifications	88
MCX614106A-VCAT Specifications	90
MCX613106A-VDAT Specifications	92
Board Mechanical Drawing and Dimensions.....	93
Bracket Mechanical Drawing	94
Monitoring.....	97
Thermal Sensors.....	97
Heatsink	97
Finding the MAC on the Adapter Card	98
Document Revision History	99

About This Manual

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

Legacy (EOL) Ordering Part Numbers

NVIDIA SKU	Legacy OPN	Marketing Description
900-9X6B4-0056-DT1	MCX614106A-CCAT	ConnectX®-6 EN adapter card, 100GbE , dual-port QSFP56, Socket Direct 2x PCIe3.0 x16 , tall bracket
900-9X6B4-0018-DT1	MCX614105A-VCAT	ConnectX®-6 EN adapter card, 200GbE , single-port QSFP56, Socket Direct 2x PCIe3.0 x16 , tall bracket
900-9X6B4-0058-DT1	MCX614106A-VCAT	ConnectX®-6 EN adapter card, 200GbE , dual-port QSFP56, Socket Direct 2x PCIe3.0 x16 , tall bracket
900-9X6AF-0058-MT1	MCX613106A-VDAT	ConnectX®-6 EN adapter card, 200GbE , dual-port QSFP56, PCIe4.0 x16 , tall bracket

Intended Audience

This manual is intended for the installer and user of these cards. The manual assumes basic familiarity with 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: enterprisesupport@nvidia.com

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

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

Related Documentation

<i>MLNX_OFED for Linux User Manual and Release Notes</i>	User Manual describing OFED features, performance, band diagnostic, tools content and configuration. See MLNX_OFED for Linux Documentation .
--	--

<i>WinOF-2 for Windows User Manual and Release Notes</i>	User Manual describing WinOF-2 features, performance, Ethernet diagnostic, tools content and configuration. See WinOF-2 for Windows Documentation .
<i>NVIDIA VMware for Ethernet User Manual</i>	User Manual and release notes describing the various components of the NVIDIA ConnectX® NATIVE ESXi stack. See VMware® ESXi Drivers Documentation .
<i>NVIDIA Firmware Utility (mlxup) User Manual and Release Notes</i>	NVIDIA firmware update and query utility used to update the firmware. Refer to Firmware Utility (mlxup) Documentation .
<i>NVIDIA Firmware Tools (MFT) User Manual</i>	User Manual describing the set of MFT firmware management tools for a single node. See MFT User Manual .
<i>IEEE Std 802.3 Specification</i>	IEEE Ethernet Specifications
<i>PCI Express Specifications</i>	Industry Standard PCI Express Base and Card Electromechanical Specifications. Refer to PCI-SIG Specifications .
LinkX Interconnect Solutions	NVIDIA LinkX Ethernet 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 40Gb/s, 56Gb/s, 100Gb/s, and 200Gb/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 Ethernet 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 is provided in [Document Revision History](#).

Introduction

Product Overview

This is the *User Manual* for Ethernet 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 data centers 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 215 million messages per second, enabling the highest performance and most flexible solution for the most demanding High-Performance Computing (HPC), Storage, and data center 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. Please refer to [Feature and Benefits](#) for more details.

 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](#).

The following section provides the ordering part number, port speed, number of ports, and PCI Express speed.

ConnectX-6 PCIe x16 Card

ConnectX-6 with a PCIe x16 slot can support a bandwidth of up to 100GbE in a PCIe Gen 3.0/4.0 slot, or up to 200GbE in a PCIe Gen 4.0 slot.




Part Number	MCX613106A-VDAT
Form Factor/Dimensions	PCIe Half Height, Half Length / 167.65mm x 68.90mm

Part Number	MCX613106A-VDAT
Data Transmission Rate	Ethernet: 10/25/40/50/100/200 Gb/s
Network Connector Type	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 (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.




ConnectX-6 Ethernet Socket Direct Adapter Cards

Model	ConnectX-6 Ethernet Socket Direct Adapter Cards		
Part Number	MCX614105A-VCAT	MCX614106A-VCAT	MCX614106A-CCAT
Form Factor / Dimensions	PCIe Full Height, ¾ Length 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)		
	Cabline CA-II Plus harness Length: 35cm		
Ethernet Data Rate	10/25/40/50/100/200 GbE		10/25/40/50/100 GbE
Network Connector Type	Single-port QSFP56	Dual-port QSFP56	

Model	ConnectX-6 Ethernet Socket Direct Adapter Cards
PCI Express Connectors	Adapter Card Edge Connector: PCIe Gen 3.0/4.0 SERDES @ 16.0GT/s x16
	Auxiliary Connection Card: PCIe Gen 3.0 SERDES @ 16.0GT/s x16
RoHS	RoHS Compliant
IC Part Number	MT28908A0-XCCF-HVM

For more detailed information see [Specifications](#).

Features and Benefits

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

Feature	Description
PCI Express (PCIe)	<p>Uses the following 2x PCIe x16 interfaces:</p> <ul style="list-style-type: none"> • PCIe Gen 3.0/4.0 SERDES@ 8.0GT/s / 16.0GT/s through Edge Connector • PCIe Gen 3.0 SERDES@ 8.0GT/s through Auxiliary Card
Up to 25 Gigabit Ethernet	<p>NVIDIA adapters comply with the following IEEE 802.3 standards:</p> <p>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

Feature	Description
	<ul style="list-style-type: none"> - IEEE 802.1Qau (QCN) - Congestion Notification - IEEE 802.1Qaz (ETS) - IEEE 802.1Qbb (PFC) - IEEE 802.1Qbg - IEEE 1588v2 - Jumbo frame support (9.6KB)
Memory	<ul style="list-style-type: none"> • PCI Express - stores and accesses Ethernet fabric connection information and packet data. • 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, data center operators often create overlay networks that carry traffic from individual virtual machines over logical tunnels in encapsulated formats such as NVGRE and VXLAN. While this solves network scalability issues, it hides the TCP packet from the hardware offloading engines, placing higher loads on the host CPU. ConnectX-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 Ethernet networks. Leveraging data center 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>

Feature	Description
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

- 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 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

Ethernet QSFP56 Interfaces

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 connector on the adapter card.

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

PCI Express Interface

ConnectX®-6 adapter cards support PCI Express Gen 3.0/4.0 (1.1 and 2.0 compatible) through 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 lanes
- Auto-negotiates to x16, x8, x4, x2, or x1
- Support for MSI/MSI-X mechanisms

LED Interface

There are two I/O LEDs per port to indicate speed and link status. LED1 is a bicolor LED (amber and green) and LED2 is a single color LED (green).

Link Indications

State	LED1 - Link speed	LED2 (Green) - Link/Activity
Beacon command for locating the adapter card	1Hz Blinking amber/green according to high/low speeds	OFF

State	LED1 - Link speed	LED2 (Green) - Link/Activity
Error	4Hz Blinking Amber indicates a problem with the link	ON
Physical Activity	The Green LED will blink.	Blinking
Link Up	The Green LED will be solid.	ON

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 or MCTP over PCIe protocols as if it is a standard NVIDIA PCIe stand-up adapter. For configuring the adapter for the specific manageability solution in use by the server, please contact NVIDIA Support.


Voltage Regulators

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

Hardware Installation

Installation and initialization of ConnectX-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](#).

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

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

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 x16 adapter card in the system	ConnectX-6 PCIe x16 Adapter Cards Installation Instructions
	Install the ConnectX-6 2x PCIe x16 Socket Direct adapter card in the system	ConnectX-6 Socket Direct (2x PCIe x16) Installation Instructions

Step	Procedure	Direct Link
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.


ConnectX-6 Configuration	Hardware Requirements
PCIe x16	A system with a 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.
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. 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.

⚠ 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.
- (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
MCX653105A-HDAT MCX653106A-HDAT MCX653105A-ECAT MCX653106A-ECAT MCX653105A-EFAT MCX653106A-EFAT	ConnectX-6 (PCIe x16) Adapter Card
MCX654105A-HCAT MCX654106A-HCAT MCX654106A-ECAT	ConnectX-6 Socket Direct (2x PCIe x16) Adapter Card

Cables and Modules

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”:

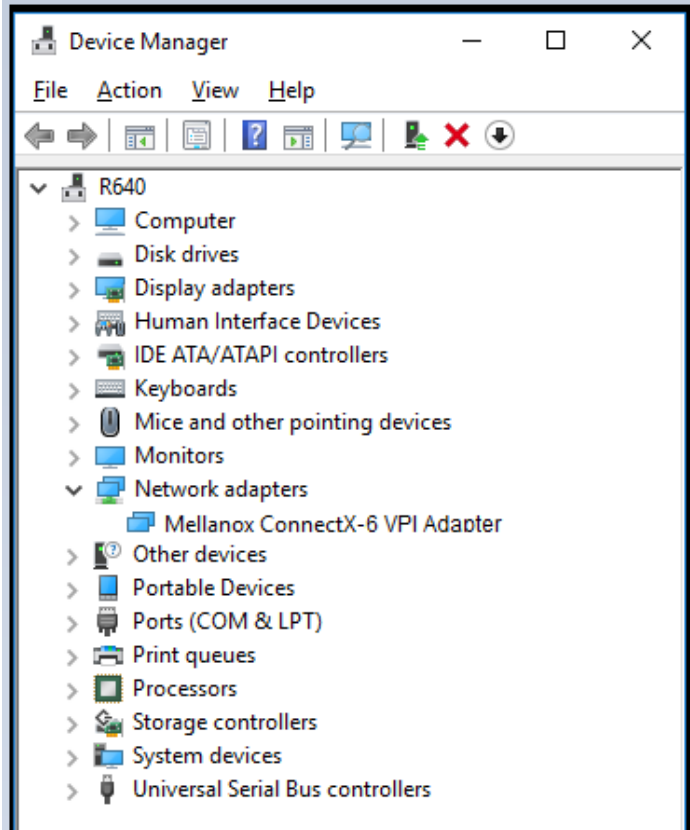
ConnectX-6 Card Configuration	lspci Command Output Example
Single-port Socket Direct Card (2x PCIe x16)	<pre>[root@mftqa-009 ~]# lspci grep mellanox -i a3:00.0 Infiniband controller: Mellanox Technologies MT28908 Family [ConnectX-6] e3:00.0 Infiniband controller: Mellanox Technologies MT28908 Family [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 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 Mellanox Technologies; and DEV is equal to 1018 (for ConnectX-6) - this is a valid Mellanox Technologies PCI Device ID.

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

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

ConnectX-6 PCIe x16 Installation Instructions


Installing the Card

⚠ Applies to MCX613106A-VDAT.


⚠ 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 card out of the PCI Express slot.



ConnectX-6 Socket Direct (2x PCIe x16) Installation Instructions

⚠ Applies to MCX614106A-CCAT, MCX614105A-VCAT and MCX614106A-VCAT.

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.



Driver Installation

Please use the relevant driver installation section.

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

Linux Driver Installation

This section describes how to install and test the MLNX_OFED for Linux package on a single server with a NVIDIA ConnectX-5 adapter card installed.

Prerequisites

Requirements	Description
Platforms	A server platform with a ConnectX-6 EN adapter card installed (firmware: fw-ConnectX6)
Required Disk Space for Installation	1GB
Operating System	Linux operating system. For the list of supported operating system distributions and kernels, please refer to the <i>MLNX_OFED Release Notes</i> file.
Installer Privileges	The installation requires administrator (root) privileges on the target machine.

2. Download the ISO image to your host.

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

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

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

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

Installing MLNX_OFED

Installation Script

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

- Discovers the currently installed kernel
 - Uninstalls any software stacks that are part of the standard operating system distribution or another vendor's commercial stack
 - Installs the MLNX_OFED_LINUX binary RPMs (if they are available for the current kernel)
 - Identifies the currently installed InfiniBand and Ethernet network adapters and automatically upgrades the firmware
- Note:** To perform a firmware upgrade using customized firmware binaries, a path can be provided to the folder that contains the firmware binary files, by running `--fw-image-dir`. Using this option, the firmware version embedded in the MLNX_OFED package will be ignored.

Example:


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

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


Usage


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

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

 Pre-existing configuration files will be saved with the extension “.conf.rpmsave”.

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

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

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

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

```
dracut -f
```

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

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

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

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

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

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

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

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

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

For the list of installation options, run:

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

Installation Procedure


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


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


```
host1# 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-X

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

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

Error message #1:

Device #1:

Device Type: ConnectX-X

Part Number: MCXXXX-XXX

PSID: MT_<version>

PCI Device Name: 0b:00.0

Base MAC: 0000e41d2d5cf810

Versions: Current Available

FW XX.XX.XXXX

Status: No matching image found

Error message #2:

The firmware for this device is not distributed inside NVIDIA driver: 0000:01:00.0 (PSID: IBM2150110033)

To obtain firmware for this device, please contact your HW vendor.

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

Action \ Adapter	Driver Restart	Standard Reboot (Soft Reset)	Cold Reboot (Hard Reset)
Standard ConnectX-4/ConnectX-4 Lx or higher	-	+	-

Adapters with Multi-Host Support	-	-	+
Socket Direct Cards	-	-	+

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

5. (InfiniBand only) Run the `hca_self_test.ofed` utility to verify whether or not the InfiniBand link is up. The utility also checks for and displays additional information such as:
 - HCA firmware version
 - Kernel architecture
 - Driver version
 - Number of active HCA ports along with their states
 - Node GUID

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

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

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

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

Installation Results

Software	<ul style="list-style-type: none"> • Most of <code>MLNX_OFED</code> 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"> • <code>fca</code> and <code>ibutils</code> • <code>iproute2</code> (rdma tool) - installed under <code>/opt/Mellanox/iproute2/sbin/rdma</code> • The kernel modules are installed under <ul style="list-style-type: none"> • <code>/lib/modules/`uname -r`/updates</code> on SLES and Fedora Distributions • <code>/lib/modules/`uname -r`/extra/mlnx-ofa_kernel</code> on RHEL and other RedHat like Distributions
----------	---

	<ul style="list-style-type: none"> • /lib/modules/`uname -r`/updates/dkms/ on Ubuntu
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 OFED ISO image <p>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.</p> • In case your machine has an unsupported network adapter device, no firmware update will occur and the error message below will be printed. "The firmware for this device is not distributed inside NVIDIA driver: 0000:01:00.0 (PSID: IBM2150110033) To obtain firmware for this device, please contact your HW vendor."

Installation Logging

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

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

Example:

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

Driver Load Upon System Boot

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

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

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

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

Software	<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 • iproute2 (rdma tool) - installed under /opt/Mellanox/iproute2/sbin/rdma • The kernel modules are installed under <ul style="list-style-type: none"> • /lib/modules/`uname -r`/updates on SLES and Fedora Distributions
----------	---

	<ul style="list-style-type: none"> • /lib/modules/`uname -r`/extra/mlnx-ofa_kernel on RHEL and other RedHat like Distributions • /lib/modules/`uname -r`/updates/dkms/ on Ubuntu
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 OFED ISO image <p>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.</p> • In case your machine has an unsupported network adapter device, no firmware update will occur and the error message below will be printed. "The firmware for this device is not distributed inside NVIDIA driver: 0000:01:00.0 (PSID: IBM2150110033) To obtain firmware for this device, please contact your HW vendor."

Installation Logging

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

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

Example:

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

Uninstalling MLNX_OFED

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

Additional Installation Procedures

Installing MLNX_OFED Using YUM

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

Setting up MLNX_OFED YUM Repository

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's GPG-KEY:
The key can be downloaded via the following link:
<http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-Mellanox>

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

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

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

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
```

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

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

```
# /mnt/mlnx_add_kernel_support.sh --make-tgz <optional --kmp> -k $(uname -r) -m /mnt/  
Note: This program will create MLNX_OFED_LINUX TGZ for rhel7.6 under /tmp directory.  
Do you want to continue?[y/N]:y  
See log file /tmp/mlnx_iso.4120_logs/mlnx_ofed_iso.4120.log  
  
Checking if all needed packages are installed...  
Building MLNX_OFED_LINUX RPMS . Please wait...  
Creating metadata-rpms for 3.10.0-957.21.3.el7.x86_64 ...  
WARNING: If you are going to configure this package as a repository, then please note  
WARNING: that it contains unsigned rpms, therefore, you need to disable the gpgcheck  
WARNING: by setting 'gpgcheck=0' in the repository conf file.  
Created /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-ext.tgz
```

4. Open the tarball.

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

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

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

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

```

mlnx-ofed-hypervisor-user-only.noarch : MLNX_OFED hypervisor-user-only installer package (User Space packages only)
mlnx-ofed-kernel-only.noarch : MLNX_OFED kernel-only installer package (with KMP support)
mlnx-ofed-vma.noarch : MLNX_OFED vma installer package (with KMP support)
mlnx-ofed-vma-eth.noarch : MLNX_OFED vma-eth installer package (with KMP support)
mlnx-ofed-vma-eth-user-only.noarch : MLNX_OFED vma-eth-user-only installer package (User Space packages only)
mlnx-ofed-vma-user-only.noarch : MLNX_OFED vma-user-only installer package (User Space packages only)
mlnx-ofed-vma-vpi.noarch : MLNX_OFED vma-vpi installer package (with KMP support)
mlnx-ofed-vma-vpi-user-only.noarch : MLNX_OFED vma-vpi-user-only installer package (User Space packages only)

```

where:

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

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

Example:

```

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

```

When using an operating system different than RHEL or SLES, or you have installed a kernel that is not supported by default in MLNX_OFED, you can use the `mlnx_add_kernel_support.sh` script to build MLNX_OFED for your kernel.

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

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

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
sdpnstat.x86_64 0:1.60-26
srptools.x86_64 0:1.0.2-12

Complete!
```

⚠ Installing MLNX_OFED using the “YUM” tool does not automatically update the firmware. To update the firmware to the version included in MLNX_OFED package, run:
`# yum install mlnx-fw-updater`

Installing MLNX_OFED Using apt-get

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

Setting up MLNX_OFED apt-get Repository

1. Log into the installation machine as root.
2. Extract the MLNX_OFED package on a shared location in your network. It can be downloaded from <https://www.nvidia.com/en-us/networking/> → Products → Software → InfiniBand Drivers.
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's Technologies GPG-KEY.

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

5. Verify that the key was successfully imported.

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

6. Update the apt-get cache.

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

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

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

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

Checking if all needed packages are installed...
Building MLNX_OFED_LINUX RPMS . Please wait...
Creating metadata-rpms for 3.10.0-957.21.3.el7.x86_64 ...
WARNING: If you are going to configure this package as a repository, then please note
```

```
WARNING: that it contains unsigned rpms, therefore, you need to disable the gpgcheck
WARNING: by setting 'gpgcheck=0' in the repository conf file.
Created /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-ext.tgz
```

4. Open the tarball.

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

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

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

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:

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

```

mlnx-ofed-all - MLNX_OFED all installer package (with DKMS support)
mlnx-ofed-vma-vpi-user-only - MLNX_OFED vma-vpi-user-only installer package (User Space packages only)
mlnx-ofed-eth-only-user-only - MLNX_OFED eth-only-user-only installer package (User Space packages only)
mlnx-ofed-vma-user-only - MLNX_OFED vma-user-only installer package (User Space packages only)
mlnx-ofed-hpc - MLNX_OFED hpc installer package (with DKMS support)
mlnx-ofed-bluefield-user-only - MLNX_OFED bluefield-user-only installer package (User Space packages only)
mlnx-ofed-dpdk - MLNX_OFED dpdk installer package (with DKMS support)
mlnx-ofed-vma-eth-user-only - MLNX_OFED vma-eth-user-only installer package (User Space packages only)
mlnx-ofed-all-user-only - MLNX_OFED all-user-only installer package (User Space packages only)
mlnx-ofed-vma-eth - MLNX_OFED vma-eth installer package (with DKMS support)
mlnx-ofed-vma - MLNX_OFED vma installer package (with DKMS support)
mlnx-ofed-dpdk-upstream-libs-user-only - MLNX_OFED dpdk-upstream-libs-user-only installer package (User Space packages only)
mlnx-ofed-basic-user-only - MLNX_OFED basic-user-only installer package (User Space packages only)
mlnx-ofed-basic-exact - MLNX_OFED basic installer package (with DKMS support) (exact)
mlnx-ofed-basic - MLNX_OFED basic installer package (with DKMS support)
mlnx-ofed-dpdk-upstream-libs - MLNX_OFED dpdk-upstream-libs installer package (with DKMS support)

```

where:

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


mlnx-ofed-basic-exact	MLNX_OFED mlnx-ofed-basic-exact installer package
dpdk-upstream-libs	MLNX_OFED dpdk-upstream-libs installer package

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

Performance Tuning

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

Windows Driver Installation

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

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

Software Requirements

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

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

Downloading WinOF-2 Driver


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

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

```
echo %PROCESSOR_ARCHITECTURE%
```

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

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

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

Installing WinOF-2 Driver

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

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

- [Attended Installation](#)
An installation procedure that requires frequent user intervention.
- [Unattended Installation](#)
An automated installation procedure that requires no user intervention.

Attended Installation

The following is an example of an installation session.

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


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

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

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

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

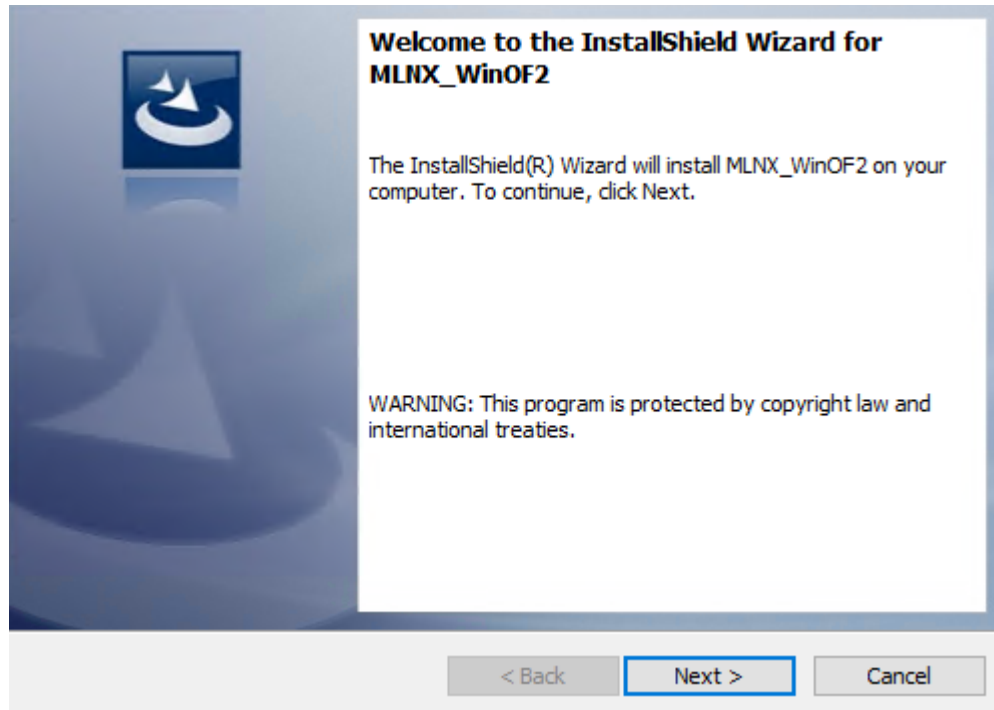
```
MLNX_WinOF2_<revision_version>_All_Arch.exe /v" MT_DISABLE_RSHIM_INSTALL=1 "
```

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

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

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


6. Click Next in the Welcome screen.



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

License Agreement

Please read the following license agreement carefully.



Copyright (c) 2005-2019, Mellanox Technologies All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above


I accept the terms in the license agreement I do not accept the terms in the license agreement

Print


InstallShield

< Back Next > Cancel

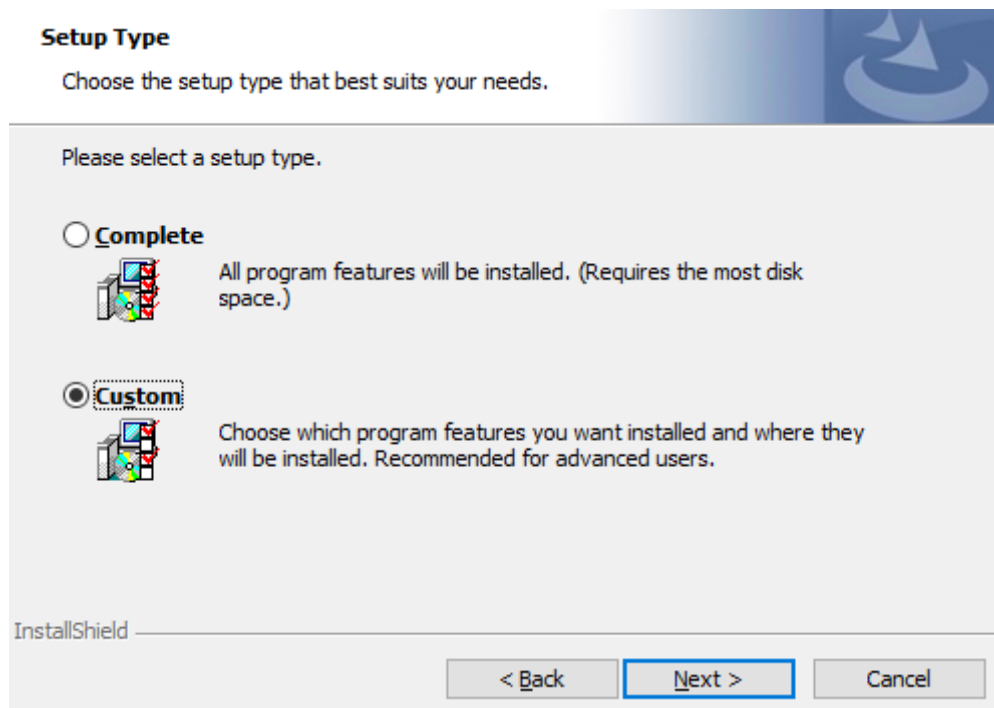
8. Select the target folder for the installation.



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




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





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

- b. Click Next to install the desired tools.



11. Click Install to start the installation.



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


Installing MLNX_WinOF2

The program features you selected are being installed.



Please wait while the InstallShield Wizard installs MLNX_WinOF2. This may take several minutes.

Status:



Firmware upgrade is required. This may take several minutes.

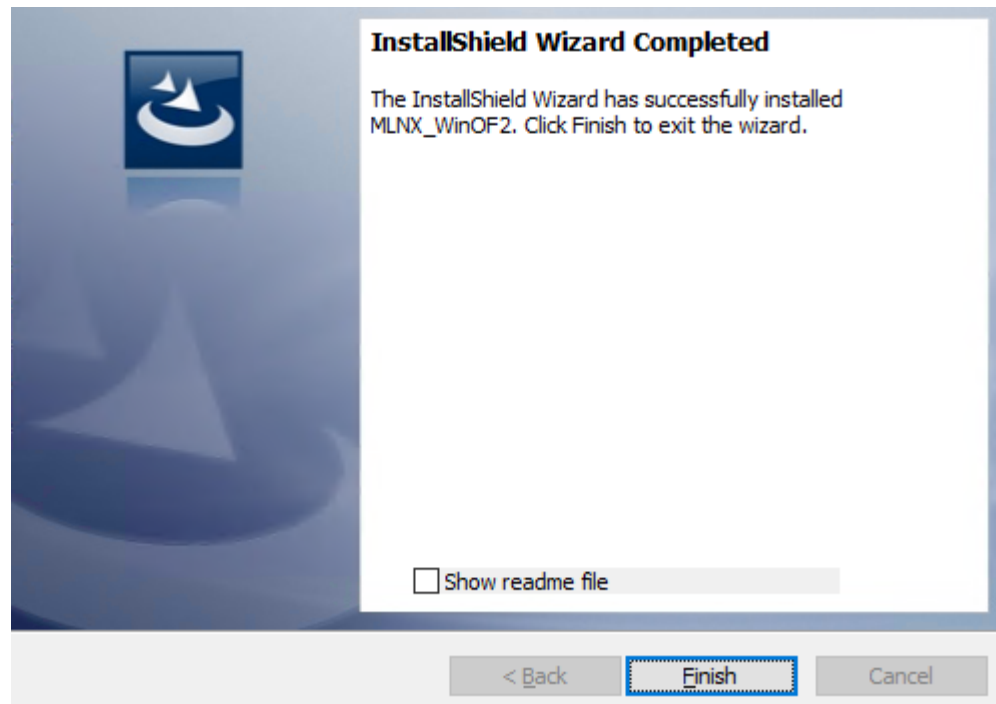
InstallShield

< Back

Next >

Cancel

13. Click Finish to complete the installation.



Unattended Installation

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

The following is an example of an unattended installation session.

1. Open a CMD console-> Click Start-> Task Manager File-> Run new task-> and enter CMD.

2. Install the driver. Run:

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

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

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

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


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

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

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

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

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

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

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

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

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

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

Firmware Upgrade

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

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

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


VMware Driver Installation

This section describes VMware Driver Installation.

Hardware and Software Requirements

Requirement	Description
Platforms	A server platform with an adapter card based on ConnectX®-6 (EN) (firmware: fw-ConnectX6)
Operating System	ESXi 6.5
Installer Privileges	The installation requires administrator privileges on the target machine.

Installing NATIVE ESXi Driver for VMware vSphere

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

To install the driver:

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


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

Example:


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

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

```
esxcli software vib list | grep nmlx  
nmlx5-core      4.16.8.8-10EM.650.0.0.4240417    MEL    PartnerSupported 2017-01-31  
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 NATIVE ESXi Driver for VMware vSphere.](#))
3. Remove each module:

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

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

4. Reboot the server.

Firmware Programming

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

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

Updating Adapter Firmware

Each adapter card is shipped with the latest version of qualified firmware at the time of manufacturing. However, NVIDIA issues firmware updates occasionally that provide new features and bug fixes. To check that your card is programmed with the latest available firmware version, download the mlxup firmware update and query utility. The utility can query for available 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 in [mlxup - Update and Query Utility](#).

Firmware Update Example

```
[server1]# ./mlxup
Querying Mellanox devices firmware ...
Device Type:      ConnectX-6
Part Number:     MCX614105A-VCAT
Description:     ConnectX®-6 EN adapter card, 200GbE, single-port QSFP56, 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:     MCX614105A-VCAT
Description:     ConnectX®-6 EN adapter card, 200GbE, single-port QSFP56, 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-yyyymmdd.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>
Mellanox Firmware Tool (MFT)	<p>Download and install MFT: MFT Documentation Refer to the User Manual for installation instructions. 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 the NVIDIA Update and Query Utility.</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	From the Windows desktop choose the Start menu and run: <code>msinfo32</code> To export system information to a text file, choose the Export option from the File menu. Assign a file name and save.
Mellanox Firmware Tool (MFT)	Download and install MFT: MFT Documentation Refer to the User Manual for installation instructions. Once installed, open a CMD window and run: WinMFT <code>mst start</code> <code>mst status</code> <code>flint -d <mst_device> q</code>
Ports Information	<code>vstat</code>
Firmware Version Upgrade	Download the latest firmware version using the PSID/board ID from here . <code>flint -d <mst_device> -i <firmware_bin_file> b</code>
Collect Log File	<ul style="list-style-type: none">• Event log viewer• MST device logs:<ul style="list-style-type: none">• <code>mst start</code>• <code>mst status</code>• <code>flint -d <mst_device> dc > dump_configuration.log</code>• <code>mstdump <mst_device> dc > mstdump.log</code>

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.

MCX614106A-CCAT Specifications

Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm)		
	Connector: Dual QSFP56 Ethernet (copper and optical)		
Protocol Support	Ethernet: 100GBASE-CR2, 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-KX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR		
	Data Rate	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 and Airflow	Voltage: 12V, 3.3VAUX			
	Power	Cable		
	Typical Power^a	Passive Cables	21.0W	
	Maximum Power	Please refer to ConnectX-6 Ethernet Cards Power Specifications available at NVOnline following login.		
	Voltage: 3.3Aux Maximum current: 100mA			
	Maximum power available through QSFP56 port: 5W			
	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	
Environmental	Temperature	Operational	0°C to 55°C	
		Non-operational	-40°C to 70°C ^b	
	Humidity	Operational	10% to 85% relative humidity	
		Non-operational	10% to 90% relative humidity	
Altitude (Operational)	3050m			
Regulatory	Safety: CB / cTUVus / CE			
	EMC: CE / FCC / VCCI / ICES / RCM / KC			

RoHS: RoHS Compliant

- a. Typical power for ATIS traffic load.
- b. The non-operational storage temperature specifications apply to the product without its package.

MCX614105A-VCAT Specifications

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 Ethernet (copper and optical)	
Protocol Support	Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBASE-CR2, 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-KX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR	
	Data Rate	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^a	Passive Cables 27.1W
	Maximum Power	Please refer to ConnectX-6 Ethernet Cards Power Specifications available at NVOnline following login.
	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		
Airflow	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction	
			Heatsink to Port	Port to Heatsink
		Passive Cables	600 LFM / 55 °C	350 LFM / 35 °C
		NVIDIA Active 4.7W Cables	600 LFM / 55 °C ^b	350 LFM / 35 °C
Environmental	Temperature	Operational	0 °C to 55 °C	
		Non-operational	-40 °C to 70 °C ^c	
	Humidity	Operational	10% to 85% relative humidity	
		Non-operational	10% to 90% relative humidity	
	Altitude (Operational)	3050m		
Regulatory	Safety: CB / cTUVus / CE			
	EMC: CE / FCC / VCCI / ICES / RCM / KC			
	RoHS: RoHS Compliant			

- a. Typical power for ATIS traffic load.
- b. For engineering samples - add 250LFM
- c. The non-operational storage temperature specifications apply to the product without its package.

MCX614106A-VCAT 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 Ethernet (copper and optical)	
Protocol Support	Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBASE-CR2, 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-KX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR	
	Data Rate	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^a	Passive Cables 31.4W
	Maximum Power	Please refer to ConnectX-6 Ethernet Cards Power Specifications available at NVOnline following login.
	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

Airflow	Temperature	Operational	0 °C to 55 °C		
		Non-operational	-40 °C to 70 °C		
	Humidity: 90% relative humidity ^c				
	Altitude (Operational)	3050m			
	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction		
			Heatsink to Port	Port to Heatsink	
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			
Environmental	Temperature	Operational	0 °C to 55 °C		
		Non-operational	-40 °C to 70 °C ^b		
	Humidity	Operational	10% to 85% relative humidity		
		Non-operational	10% to 90% relative humidity		
	Altitude (Operational)	3050m			
Regulatory	Safety: CB / cTUVus / CE				
	EMC: CE / FCC / VCCI / ICES / RCM// KC				
	RoHS: RoHS Compliant				

a. Typical power for ATIS traffic load.

b. The non-operational storage temperature specifications apply to the product without its package.

MCX613106A-VDAT Specifications


Physical	Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm)	
	Connector: Dual QSFP56 Ethernet (copper and optical)	
Protocol Support	Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBASE-CR2, 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-KX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR	
	Data Rate	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^a	Passive Cables 23.6W
	Maximum Power	Please refer to ConnectX-6 Ethernet Cards Power Specifications available at NVOnline following login.
	Voltage: 3.3Aux Maximum current: 100mA	
	Maximum power available through QSFP56 port: 5W	

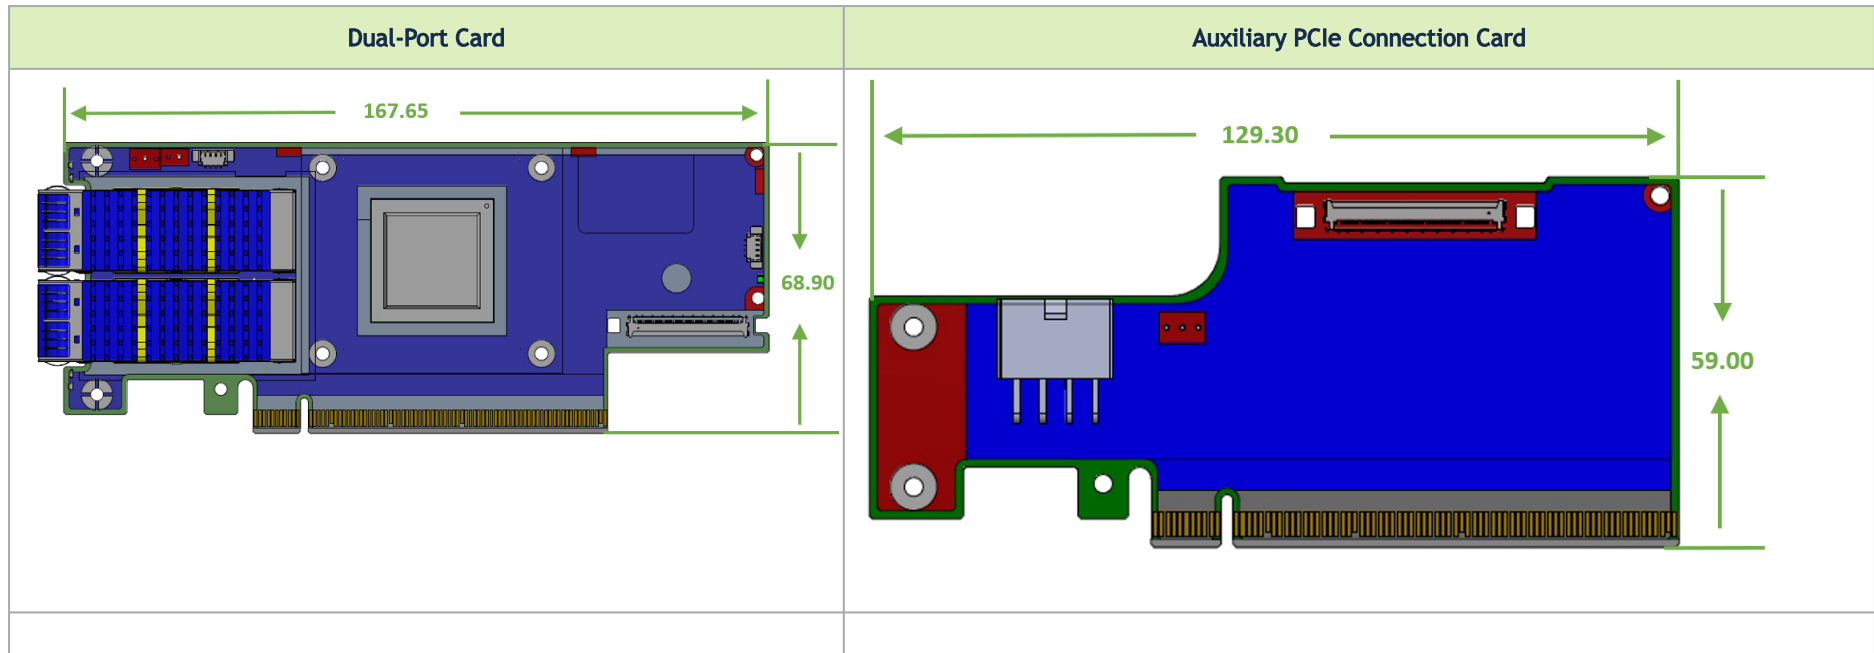
Airflow	Airflow (LFM) / Ambient Temperature	Cable Type	Airflow Direction	
			Heatsink to Port	Port to Heatsink
		Passive Cables	400 LFM / 55° C	300 LFM / 35° C
		NVIDIA Active 4.7W Cables	950 LFM / 55° C 600 LFM / 48° C	300 LFM / 35° C
Environmental	Temperature	Operational	0° C to 55° C	
		Non-operational	-40° C to 70° C ^b	
	Humidity	Operational	10% to 85% relative humidity	
		Non-operational	10% to 90% relative humidity	
	Altitude (Operational)	3050m		
Regulatory	Safety: CB / cTUVus / CE			
	EMC: CE / FCC / VCCI / ICES / RCM / KC			
	RoHS: RoHS Compliant			

a. Typical power for ATIS traffic load.


b. The non-operational storage temperature specifications apply to the product without its package.


Board Mechanical Drawing and Dimensions

 All dimensions are in millimeters. All the mechanical tolerances are +/- 0.1mm.




Bracket Mechanical Drawing

 All dimensions are in millimeters. All the mechanical tolerances are +/- 0.1mm.



Auxiliary PCIe Connection Card Tall Bracket



Monitoring

Thermal Sensors

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

Three thermal threshold definitions impact the overall system operation state:

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

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

Heatsink


The heatsink is attached to the ConnectX-6 IC in order 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](#).

Finding the 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.

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

MCX614106A-CCAT Board Label (Example)



Document Revision History

Date	Description of Changes
Nov. 2023	Added the 100GBASE-CR2 protocol in Specifications tables
May. 2023	Updated Specifications to include non-operational storage temperature specifications
Aug. 2022	<ul style="list-style-type: none">• Updated the "Legacy (EOL) Ordering Part Numbers"• Added a note on the EEPROM memory component under the Features and Benefits table
Jan. 2022	Added table "Legacy (EOL) Ordering Part Numbers"
Jun. 2021	Updated " LED Interfaces "
Mar. 2021	Updated "Troubleshooting"
Jan. 2020	Migrated to on-line format; minor reorganization.
Feb. 2019	Updated power numbers in "Specifications"
Nov. 2018	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

© 2023 NVIDIA Corporation & affiliates. All Rights Reserved.

