NVIDIA ConnectX-4 Lx Ethernet Adapter Cards User Manual

Table of Contents

ⓘ   You can download a PDF version of the document here.

ⓘ   The ConnectX-4 Lx product line has moved to end-of-life.

About This Manual

This User Manual describes NVIDIA® ConnectX®-4 Lx 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.

EOL'd (End of Life) Ordering Part Numbers

The table below provides the ordering part numbers (OPN) for ConnectX-4 Lx Ethernet adapter cards.

NVIDIA SKU	Legacy OPN	Marketing Description
900-9X4B0-0012-0T1	MCX4111A-XCAT	ConnectX®-4 Lx EN network interface card, **10GbE single**-port **SFP28**, PCIe 3.0 x8, tall bracket
900-9X4B0-0052-0T1	MCX4121A-XCAT	ConnectX®-4 Lx EN network interface card, **10GbE dual**-port **SFP28**, PCIe3.0 x8, tall bracket
900-9X4B0-0052-ST0	MCX4121A-XCHT	ConnectX®-4 Lx EN network interface card, with **Host Management**, **10GbE dual**-port **SFP28**, PCIe3.0 x8, **UEFI Enabled**, tall bracket
900-9X4B0-0013-0Q0	MCX4111A-ACAT	ConnectX®-4 Lx EN network interface card, **25GbE single**-port **SFP28**, PCIe3.0 x8, tall bracket
900-9X4B0-0053-0T1	MCX4121A-ACAT	ConnectX®-4 Lx EN network interface card, **25GbE dual**-port **SFP28**, PCIe3.0 x8, tall bracket
900-9X4B0-0013-ST0	MCX4111A-ACUT	ConnectX®-4 Lx EN network interface card, **25GbE single**-port **SFP28**, PCIe3.0 x8, **UEFI Enabled**, tall bracket
900-9X4B0-0053-ST0	MCX4121A-ACUT	ConnectX®-4 Lx EN network interface card, **25GbE dual**-port **SFP28**, PCIe3.0 x8, **UEFI Enabled**, tall bracket
900-9X4B0-0053-ST2	MCX4121A-ACHT	ConnectX®-4 Lx EN network interface card, with **Host Management**, **25GbE dual**-port **SFP28**, PCIe3.0 x8, **UEFI Enabled**, tall bracket
900-9X4B0-0053-ST3	MCX4121A-ACST	ConnectX®-4 Lx EN network interface card, **Secure Firmware**, **25GbE dual**-port **SFP28**, PCIe3.0 x8, tall bracket
900-9X429-0014-0T0	MCX4131A-BCAT	ConnectX®-4 Lx EN network interface card, **40GbE single**-port **QSFP28**, PCIe3.0 x8, tall bracket
900-9X429-0015-0T0	MCX4131A-GCAT	ConnectX®-4 Lx EN network interface card, **50GbE single**-port **QSFP28**, PCIe3.0 x8, 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

NVIDIA OFED for Linux User Manual and Release Notes	User Manual describing OFED features, performance, band diagnostic, tools content and configuration. See NVIDIA OFED for Linux Documentation.
WinOF-2 for Windows User Manual and Release Notes	User Manual describing WinOF-2 features, performance, Ethernet diagnostic, tools content and configuration. See WinOF-2 for Windows Documentation.
NVIDIA VMware for Ethernet User Manual and Release Notes	User Manual describing the various components of the NVIDIA ConnectX® NATIVE ESXi stack. See http://www.nvidia.com Products > Software > Ethernet Drivers > VMware Driver
NVIDIA Firmware Utility (mlxup) User Manual and Release Notes	NVIDIA firmware update and query utility used to update the firmware. See http://www.nvidia.com > Products > Software > Firmware Tools >mlxup Firmware Utility
NVIDIA Firmware Tools (MFT) User Manual	User Manual describing the set of MFT firmware management tools for a single node. See MFT User Manual.
IEEE Std 802.3 Specification	IEEE Ethernet specification at http://standards.ieee.org/
PCI Express Specifications	Industry Standard PCI Express Base and Card Electromechanical Specifications at https://pcisig.com/specifications
NVIDIA LinkX Interconnect Solutions	NVIDIA LinkX InfiniBand cables and transceivers are designed to maximize the performance of High-Performance Computing networks, requiring high-bandwidth, low-latency connections between compute nodes and switch nodes. NVIDIA offers one of industry's broadest portfolio of QDR/FDR10 (40Gb/s), FDR (56Gb/s), EDR/HDR100 (100Gb/s) and HDR (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 IBTA standards, NVIDIA tests every product in an end-to-end environment ensuring a Bit Error Rate of less than 1E-15. Read more at https://www.nvidia.com/products/interconnect/infiniband-overview.php

Document Conventions

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

Introduction

Product Overview

This is the User Guide for Ethernet adapter cards based on the ConnectX®-4 Lx integrated circuit device. These adapters connectivity provide the highest performing low latency and most flexible interconnect solution for PCI Express Gen 3.0 servers used in Enterprise Data Centers and High-Performance Computing environments.

The following provides the ordering part number, port speed, number of ports, and PCI Express speed. Each adapter is shipped with two bracket heights - short and tall (assembled).

ConnectX-4 Lx 10Gb/s Ethernet Adapter Cards

Model	ConnectX-4 Lx 10Gb/s Ethernet Adapter Cards		
Part Number	MCX4111A-XCAT	MCX4121A-XCAT	MCX4121A-XCHT
Dimensions	2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – Low profile		
PCI Express Connectors	PCIe Gen 3.0 SERDES @ 8.0GT/s x8 lanes		
Ethernet Data Rate	1/10 Gb/s	1/10 Gb/s	1/10 Gb/s
Network Connector Type	Single-port SFP28	Dual-port SFP28	Dual-port SFP28
UEFI	-	-	✓
Host Management	-	-	✓
RoHS	RoHS Compliant		
IC Part Number	MT27712A0-FDCF-XE		
Device ID	4117 for Physical Function (PF) 4118 for Virtual Function (VF)		

ConnectX-4 Lx 25Gb/s Ethernet Adapter Cards

Model	ConnectX-4 Lx 25Gb/s Ethernet Adapter Cards					
Part Number	MCX4111A-ACAT	MCX4111A-ACUT	MCX4121A-ACAT	MCX4121A-ACUT	MCX4121A-ACST	MCX4121A-ACHT
Network Connector Type	1Port SFP28	1Port SFP28	2Ports SFP28	2Ports SFP28	2Ports SFP28	2Ports SFP28
Ethernet Data Rate	1/10/25 Gb/s					
PCI Express Connectors	PCIe Gen 3.0 SERDES @ 8.0GT/s x8 Lanes					
Dimensions	2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile					
UEFI	-	✓	-	✓	-	✓
Secure Firmware	-	-	-	-	✓	-
Host Management	-	-	-	-	-	✓
RoHS	RoHS Compliant					
IC Part Number	MT27712A0-FDCF-AE					
Device ID	4117 for Physical Function (PF) 4118 for Virtual Function (VF)					

ConnectX-4 Lx 40/50 Gb/s Ethernet Adapter Cards

Model	ConnectX-4 Lx 40/50 Gb/s Ethernet Adapter Cards	
Part Number	MCX4131A-BCAT	MCX4131A-GCAT
Ethernet Data Rate	1/10/25/40 Gb/s	1/10/25/40/50 Gb/s
Network Connector Type	Single-port QSFP28	
PCI Express Connectors	PCIe Gen 3.0 x8 SerDes @ 8.0GT/s x8 Lanes	
Dimensions	2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile	
RoHS	RoHS Compliant	
IC Part Number	MT27711A0-FDCF-BE	MT27711A0-FDCF-GE
Device ID	4117 for Physical Function (PF) 4118 for Virtual Function (VF)	

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)	Uses PCIe Gen 3.0 (8GT/s) through an x8 edge connector. Gen 1.1 and 2.0 compatible.
Up to 50 Gigabit Ethernet	NVIDIA adapters comply with the following IEEE 802.3 standards: • 50GbE / 40GbE / 25GbE / 10GbE / 1GbE • 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 • Proprietary Ethernet protocols (20/40GBASE-R2, 50GBASE-R4) • IEEE 802.3ad, 802.1AX Link Aggregation • IEEE 802.1Q, 802.1P VLAN tags and priority • IEEE 802.1Qau (QCN) • Congestion Notification • IEEE 802.1Qaz (ETS) • IEEE 802.1Qbb (PFC) • IEEE 802.1Qbg • IEEE 1588v2 • Jumbo frame support (9.6KB)
Memory	• PCI Express - stores and accesses Ethernet fabric connection information and packet data. • SPI - includes one 16MB SPI Flash device (W25Q128FVSIG device by WINBOND-NUVOTON).

Feature	Description
Overlay Networks	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-4 Lx effectively addresses this by providing advanced NVGRE and VXLAN hardware offloading engines that encapsulate and de-capsulate the overlay protocol.
RDMA and RDMA over Converged Ethernet (RoCE)	ConnectX-4 Lx, utilizing IBTA RDMA (Remote Data Memory Access) and RoCE (RDMA over Converged Ethernet) technology, delivers low-latency and high-performance over Band and Ethernet networks. Leveraging data center bridging (DCB) capabilities, as well as ConnectX-4 Lx, advanced congestion control hardware mechanisms, RoCE provides efficient low-latency RDMA services over Layer 2 and Layer 3 networks.
NVIDIA PeerDirect™	PeerDirect™ communication provides high-efficiency RDMA access by eliminating unnecessary internal data copies between components on the PCIe bus (for example, from GPU to CPU), and therefore significantly reduces application run time. ConnectX-4 Lx advanced acceleration technology enables higher cluster efficiency and scalability to tens of thousands of nodes.
CPU Offload	Adapter functionality enabling reduced CPU overhead allowing more available CPU for computation tasks.
Quality of Service (QoS)	Support for port-based Quality of Service enabling various application requirements for latency and SLA.
Hardware-based I/O Virtualization	ConnectX-4 Lx provides dedicated adapter resources and guaranteed isolation and protection for virtual machines within the server.
Storage Acceleration	A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols can leverage RDMA for high-performance storage access. • NVMe over Fabric offloads for the target machine
SR-IOV	ConnectX-4 Lx SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server.
NC-SI	The adapter supports a Network Controller Sideband Interface (NC-SI), MCTP over SMBus and MCTP over PCIe - Baseboard Management Controller interface.
High-Performance Accelerations	• Tag Matching and Rendezvous Offloads • Adaptive Routing on Reliable Transport • Burst Buffer Offloads for Background Checkpointing
UEFI	UEFI is a standard firmware interface designed to replace BIOS. NVIDIA UEFI Network driver allows boot over the network via PXE (Preboot eXecution Environment). This network driver allows remote boot over Ethernet, or Boot over iSCSI (Bo-iSCSI) in UEFI mode, and also supports the Secure Boot standard. The UEFI Network driver allows IT managers the flexibility to deploy servers with a single adapter card into Ethernet networks while also enabling booting from LAN or remote storage targets. In addition to boot capabilities, NVIDIA UEFI Network driver provides firmware management and diagnostic protocols compliant with the UEFI specification. For further information, refer to the NVIDIA PreBoot Drivers User Manual. Supported in MCX4111A-ACUT, MCX4121A-XCHT, MCX4121A-ACUT and MCX4121A-ACHT.

Feature	Description
Secure Firmware	A "Secure firmware update" is the ability of a device to verify digital signatures of new firmware binaries, in order to assure that only officially approved versions can be installed from the host, the network or a Board Management Controller (BMC). The firmware of devices with "secure firmware update" functionality (secure FW), restricts access to specific commands and registers that can be used to modify the firmware binary image on the flash, as well as commands that can jeopardize security in general. For further information, refer to the MFT User Manual. Supported in MCX4121A-ACST.

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 Gb/s Ethernet switches
- Passive copper cable with ESD protection
- Powered connectors for optical and active cable support

Interfaces

Ethernet SFP28 and QSFP28 Interfaces

The network ports of the ConnectX®-4 Lx adapter card are compliant with the IEEE 802.3 Ethernet standards listed in Features and Benefits. Ethernet traffic is transmitted through the SFP28/QSFP28 connectors 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

The table below describes the supported PCIe interface in ConnectX-4 Lx adapter cards.

Supported PCIe Interface	Features
PCIe Gen 3.0 (1.1 and 2.0 compatible) through x8 edge connectors	Link Rates: 2.5. 5.0 or 8.0GT/s. Auto Negotiation to: x8, x4, x2 or x1. Support for MSI/MSI-X mechanisms.

LED Interface

There is one bi-color I/O LED per port to indicate link status located on the adapter card.

Link Indications (Ethernet Mode)

LED Color and State	Description
Off	A link has not been established
Blinking Amber[a]	4 Hz blinking amber indicates a problem with the link
Solid Green	Indicates a valid link with no active traffic
Blinking Green	Indicates a valid logical link with active traffic

Note a. 1Hz Blinking amber occurs due to running a beacon command for locating the adapter card.

Heatsink Interface

A heatsink is attached to the ConnectX-4 Lx IC in order to dissipate the heat from the ConnectX-4 Lx IC. It is attached either by using four spring-loaded push pins that insert into four mounting holes, or by screws.

ConnectX-4 Lx IC has a thermal shutdown safety mechanism which automatically shuts down the ConnectX-4 Lx card in cases of high-temperature event, improper thermal coupling or heatsink removal.

SMBus Interface

ConnectX-4 Lx technology maintains support for manageability through a BMC. ConnectX-4 Lx 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-4 Lx adapter cards require attention to the mechanical attributes, power specifications, and precautions for electronic equipment.

Safety Warnings

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

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

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

Installation Procedure Overview

The installation procedure of ConnectX-4 Lx adapter cards involves the following steps:

Step	Procedure	Direct Link
1	Check the system's hardware and software requirements.	Refer to System Requirements
2	Pay attention to the airflow consideration within the host system.	Refer to Airflow Requirements
3	Unpack the product.	Refer to Unpacking the Product
4	(Optional) Replace the full-height mounting bracket with the supplied short bracket.	Refer to Bracket Replacement Instructions
5	Install the ConnectX-4 Lx adapter card in the system.	Refer to Installation Instructions
6	Connect cables or modules to the card.	Refer to Cables and Modules
7	Identify the ConnectX-4 Lx adapter card in the system.	Refer to Identifying the Card in Your System

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 operation environment should meet severity level G1 as per ISA 71.04 for gaseous contamination and ISO 14644-1 class 8 for cleanliness level.

A system with a PCI Express x8 edge connector is required for installing the card.

> ⬥ 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.

Airflow Requirements

ConnectX-4 Lx 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 chapter for airflow numbers for each specific card model.

> ⬥ 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.

Unpacking the Product

> ⬥ 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.

1. Unpack the adapter package and place them on an antistatic surface, and verify you have received the following items:
   a. ConnectX-4 Lx Adapter Card, with a tall bracket assembled on the card.
   b. Short bracket (accessory)
2. Check the parts for visible damage that may have occurred during shipping.

> ⬥ 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.

EMI Fingers

Plastic Zip Lock

3. Shut down your system if active:
   Turn off the power to the system, and disconnect the power cord. Refer to the system documentation for instructions. Before you install the ConnectX-4 Lx card, make sure that the system is disconnected from power.

Bracket Replacement Instructions

The card is 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-4 Lx 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.

> ⚠ Please note that the following figures are for illustration purposes only.

1. Before installing the card, make sure that the system is off and the power cord is not connected to the server. Please follow proper electrical grounding procedures.
2. Open the system case.
3. Place the adapter in an available PCI Express slot.
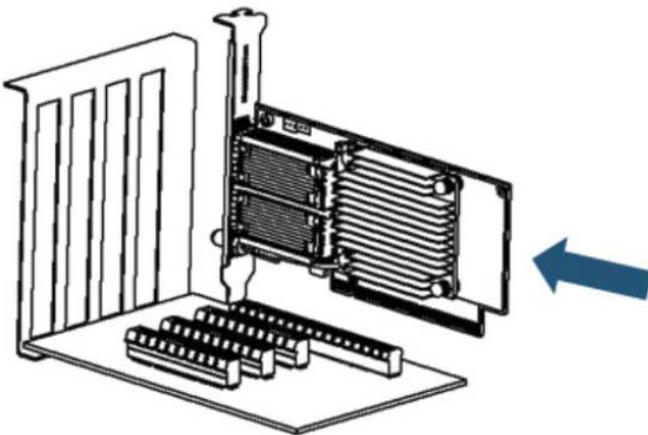
4. Applying even pressure at both corners of the card, insert the adapter card into the PCI Express slot until firmly seated.
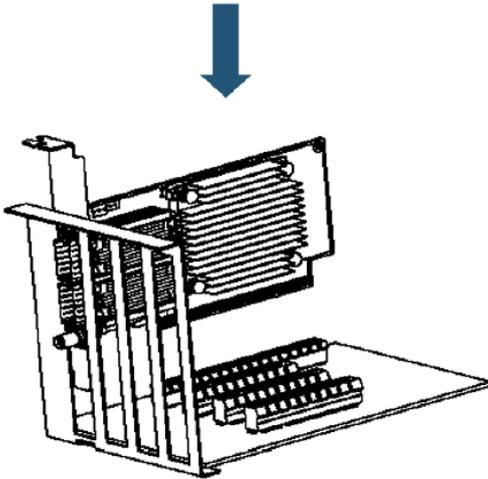
5. When the adapter is properly seated, the port connectors are aligned with the slot opening, and the adapter faceplate is visible against the system chassis.
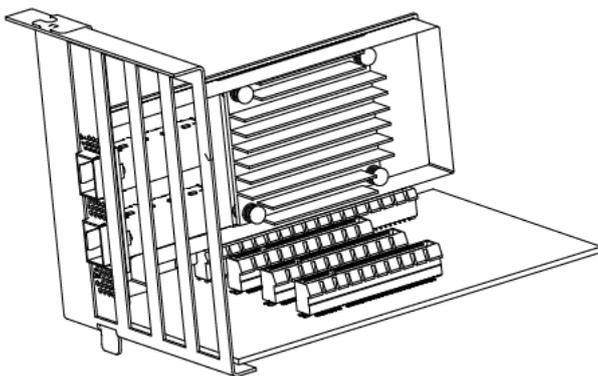
6. Secure the adapter with the adapter clip or screw.

> ⓘ   To uninstall the adapter card, see Uninstalling the 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 Operations 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 Operations 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. The LED indicator will turn off when the cable is unseated.

Identifying the Card in Your System

On Linux

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

```
lspci |grep -i Mellanox
Network controller: Mellanox Technologies MT27710 Family [ConnectX-4 Lx]
```

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-4 Lx 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)
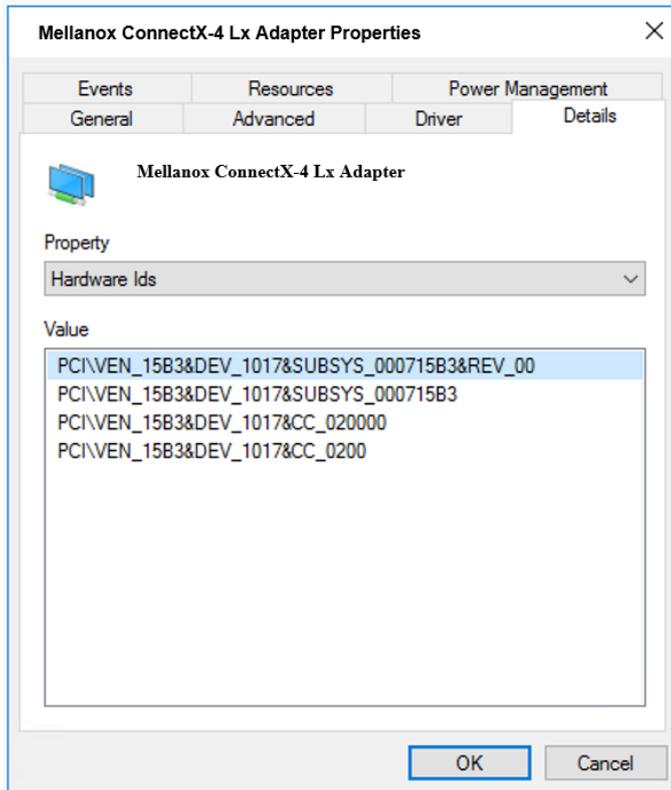
5. In the Value display box, check the fields VEN and DEV (fields are separated by '&'). In the display example above, notice the sub-string "PCI\VEN_15B3&DEV_1003": VEN is equal to 0x15B3 – this is the Vendor ID; and DEV is equal to 1018 (for ConnectX-4 Lx) – this is a valid PCI Device ID.

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

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

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.

> ⓘ 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 adapter card out of the PCI Express slot.
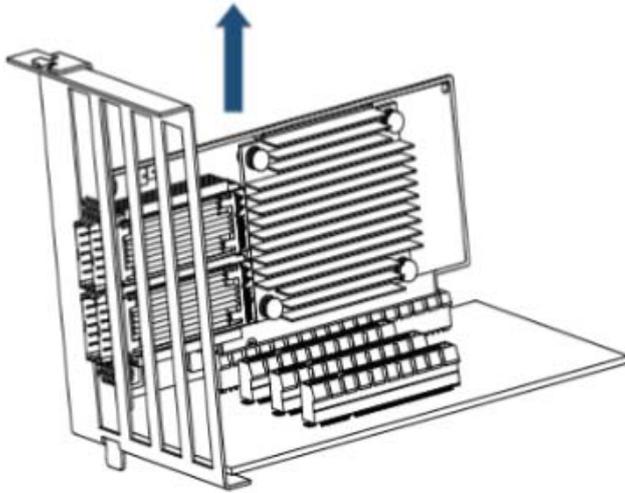
5. When the port connectors reach the top of the chassis window, gently pull the adapter card in parallel to the motherboard.
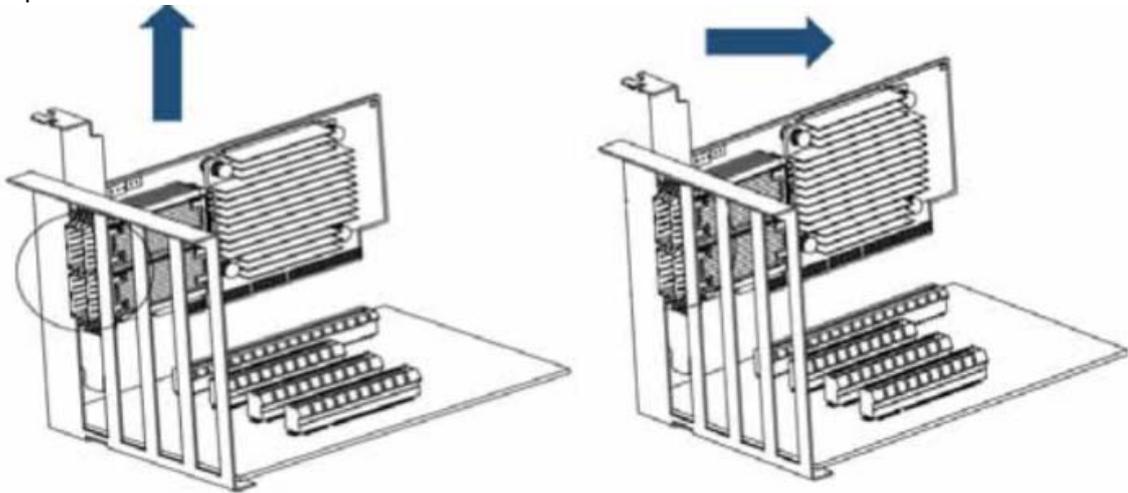
Driver Installation

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 ConnectX-5 adapter card installed.

Prerequisites

Requirements	Description
Platforms	A server platform with one ConnectX®-4 Lx (EN) (firmware: fw-ConnectX4Lx)
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 *MLNX_OFED Release Notes* file.
Installer Privileges	The installation requires administrator (root) privileges on the target machine.

Downloading NVIDIA OFED

1. Verify that the system has a network adapter installed.
   The following example shows a system with an installed adapter card:

```
# lspci -v | grep Mellanox
86:00.0 Network controller [0207]: Mellanox Technologies MT27710 Family
        Subsystem: Mellanox Technologies Device 0014
86:00.1 Network controller [0207]: Mellanox Technologies MT27710 Family
        Subsystem: Mellanox Technologies Device 0014
```

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 (suchaspdsh).

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

Installation Logging

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

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

Example:

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

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

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

Installation Logging

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

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

Example:

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

Uninstalling MLNX_OFED

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

Additional Installation Procedures

Installing MLNX_OFED Using YUM

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

Setting up MLNX_OFED YUM Repository

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.x8
6_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.fc2
1.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.fc2
1.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.x8
6_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-
01.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.x8
6_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-
01.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.x8
6_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-isert = 1.0-OFED.3.1.0.1.2.1.g832a737.rhel7u1 for package: mlnx-ofed-
all-3.1-0.1.2.noarch
..................
..................
  qperf.x86_64 0:0.4.9-9
  rds-devel.x86_64 0:2.0.7-1.12
  rds-tools.x86_64 0:2.0.7-1.12
  sdpnetstat.x86_64 0:1.60-26
  srptools.x86_64 0:1.0.2-12

Complete!
```

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

Installing MLNX_OFED Using apt-get

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

Setting up MLNX_OFED apt-get Repository

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)	

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

Downloading WinOF-2 Driver

➤ *To download the .exe file according to your Operating System, please follow the steps below:*
1. Obtain the machine architecture.
    a. To go to the Start menu, position your mouse in the bottom-right corner of the Remote Desktop of your screen.
    b. Open a CMD console (Click Task Manager-->File --> Run new task and enter CMD).
    c. Enter the following command.

    ```
    echo %PROCESSOR_ARCHITECTURE%
    ```

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

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

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

Installing WinOF-2 Driver

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

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

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

Attended Installation

The following is an example of an installation session.

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

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

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

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

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

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

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

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

```
MLNX_WinOF2_<revision_version>_All_Arch.exe /v" SKIPUNSUPPORTEDDEVCHECK=1"
```
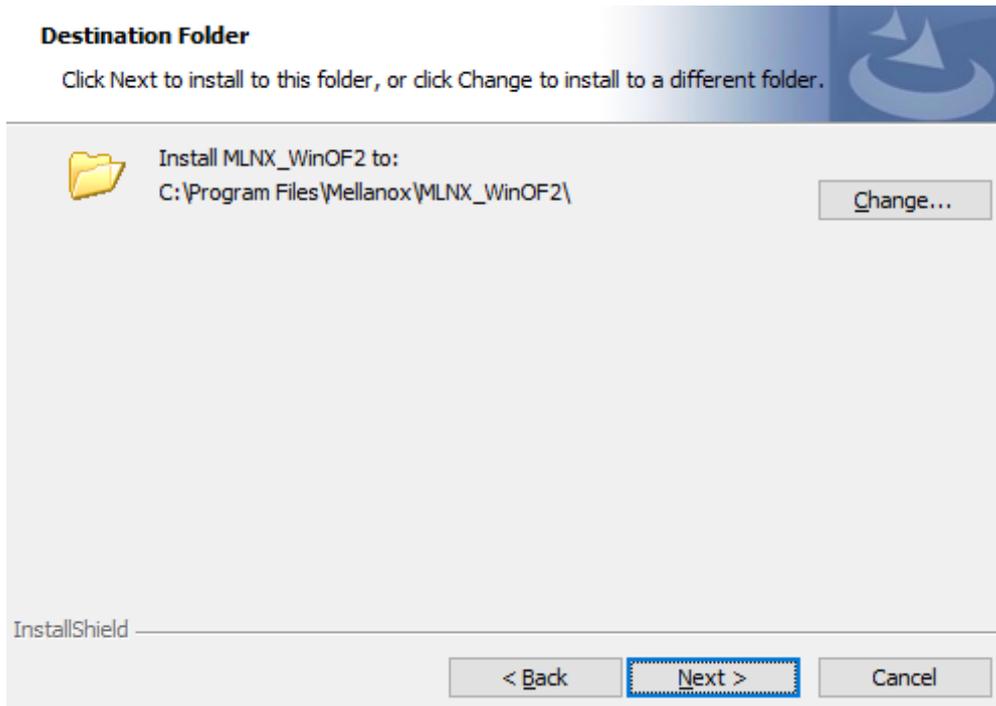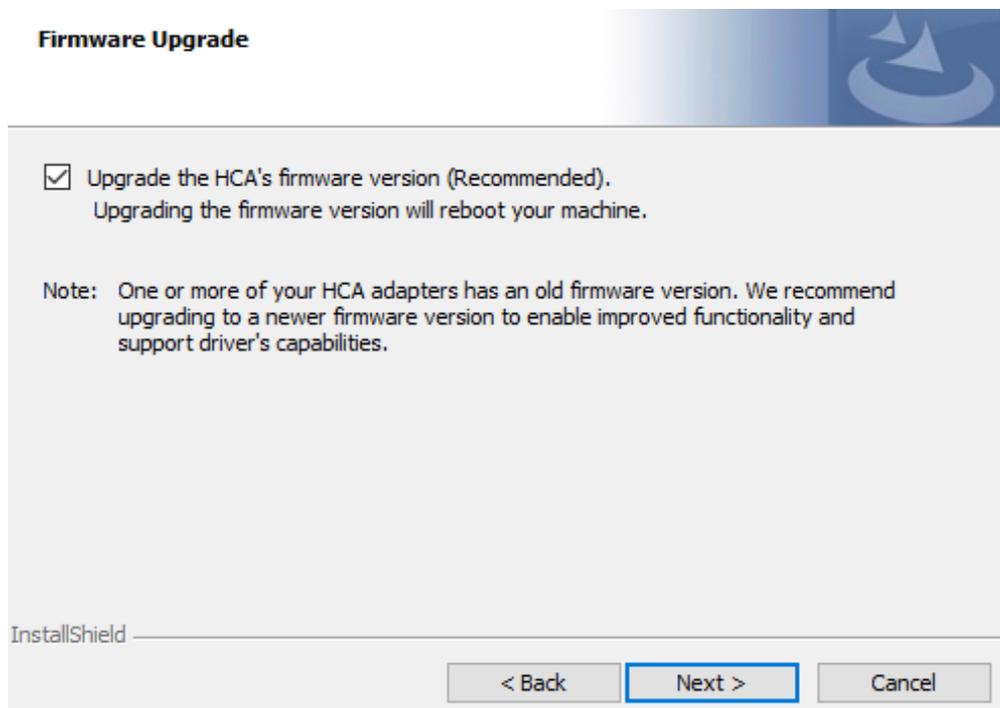
6. Click Next in the Welcome screen.
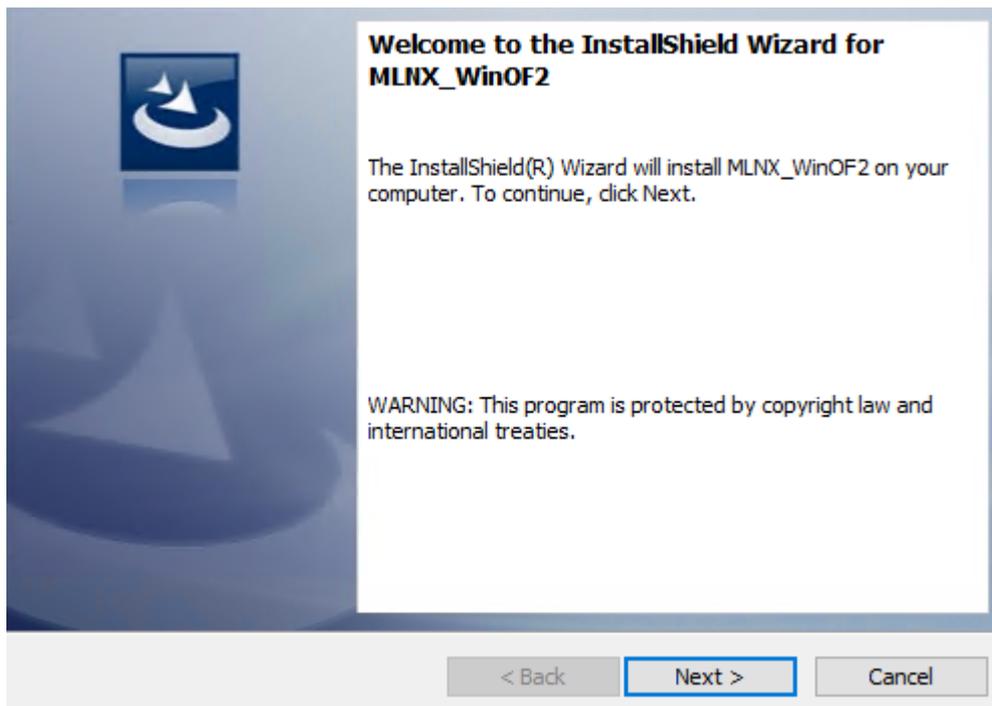
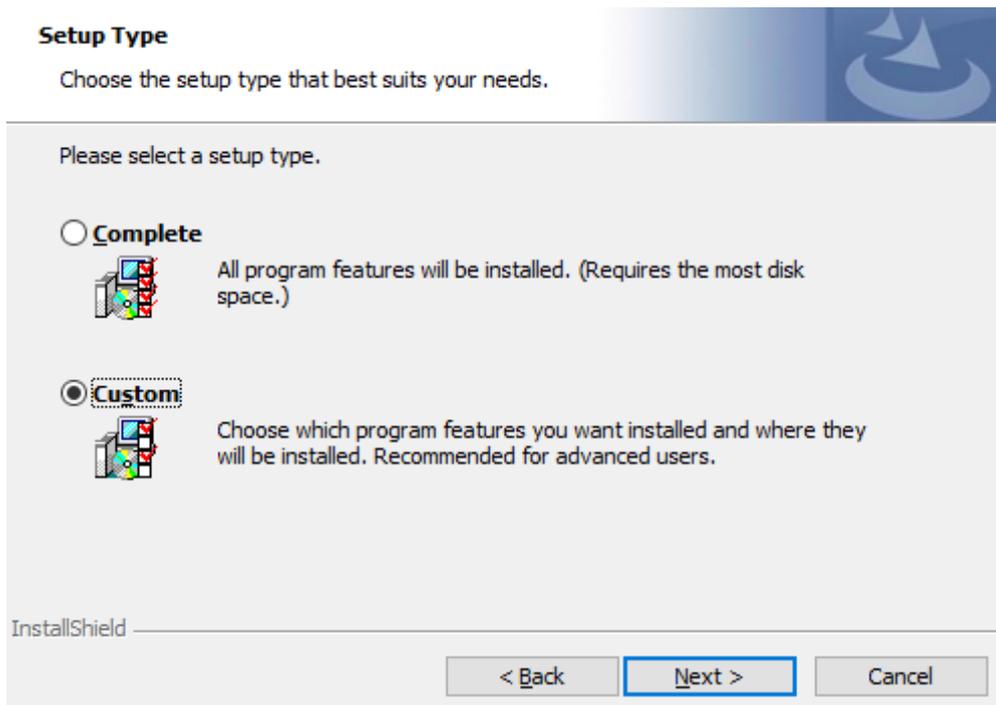7. Read and accept the license agreement and click Next.
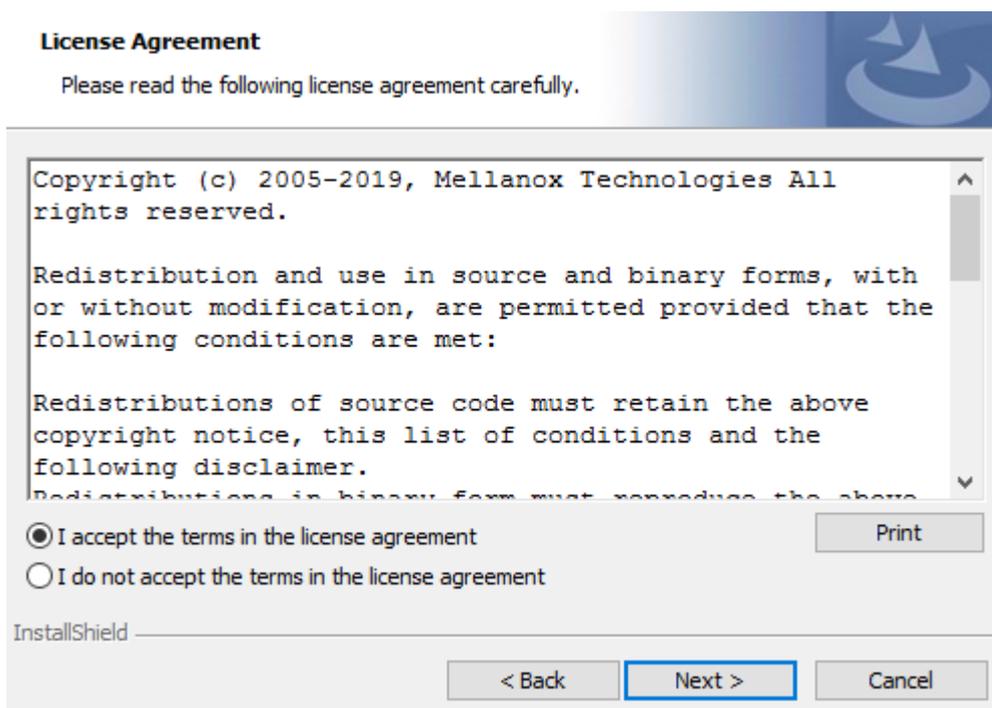
8. Select the target folder for the installation.

Destination Folder
Click Next to install to this folder, or click Change to install to a different folder.

Install MLNX_WinOF2 to:
C:\Program Files\Mellanox\MLNX_WinOF2\

Change...

InstallShield

< Back   Next >   Cancel

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.

Firmware Upgrade

☑ Upgrade the HCA's firmware version (Recommended).
    Upgrading the firmware version will reboot your machine.

Note: One or more of your HCA adapters has an old firmware version. We recommend upgrading to a newer firmware version to enable improved functionality and support driver's capabilities.

InstallShield

< Back   Next >   Cancel

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

a. Select the desired feature to install:
   - Performances tools - install the performance tools that are used to measure performance in user environment
   - Documentation - contains the User Manual and Release Notes
   - Management tools - installation tools used for management, such as mlxstat
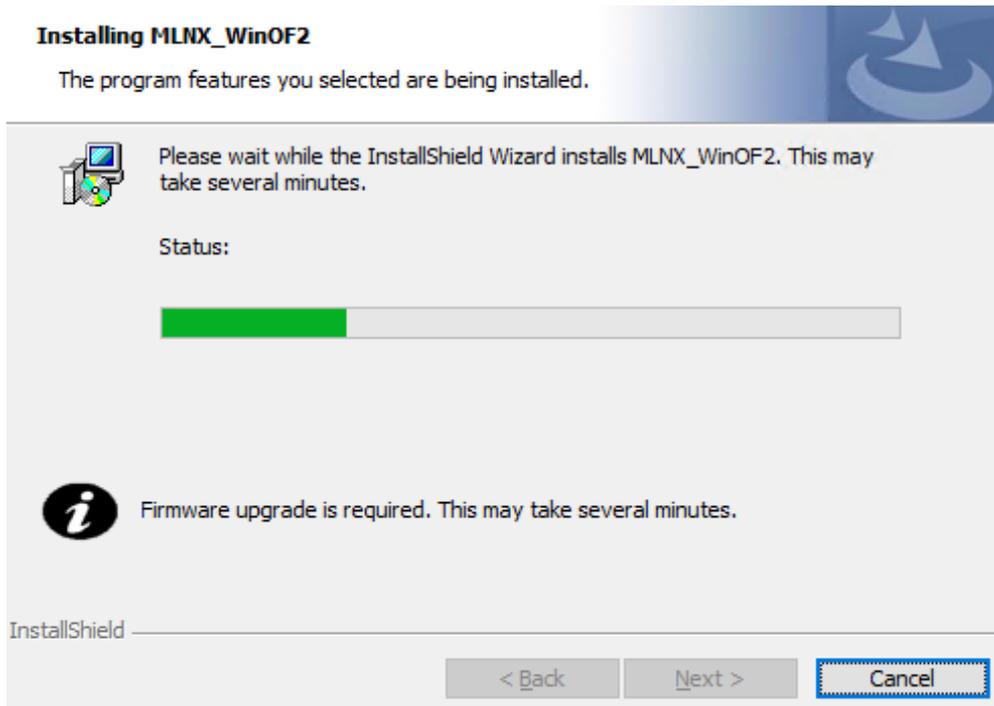   - Diagnostic Tools - installation tools used for diagnostics, such as mlx5cmd
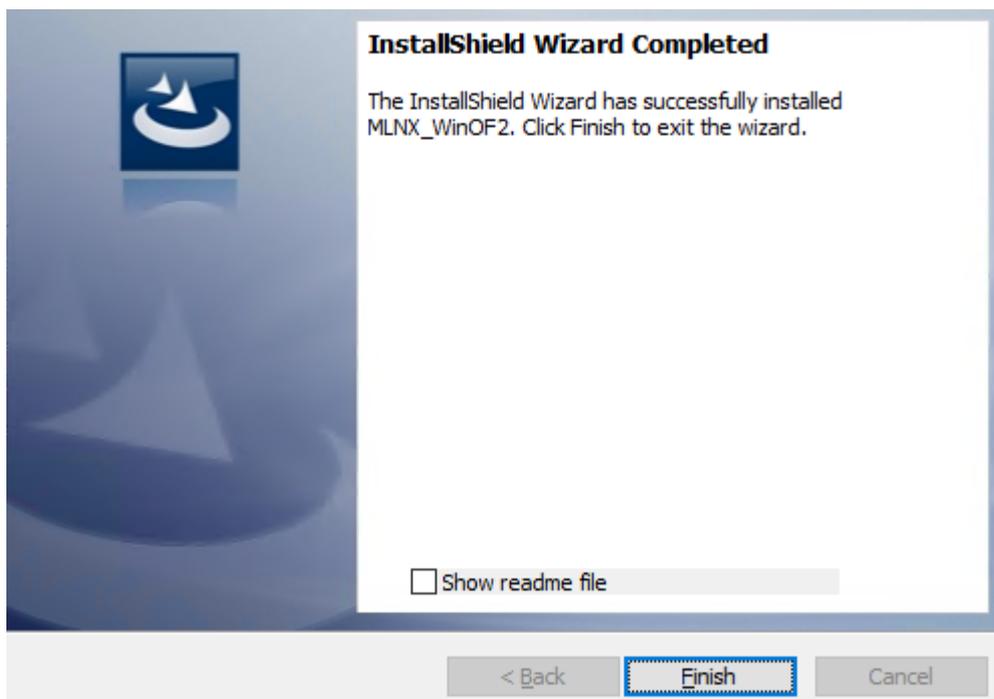
b. Click Next to install the desired tools.
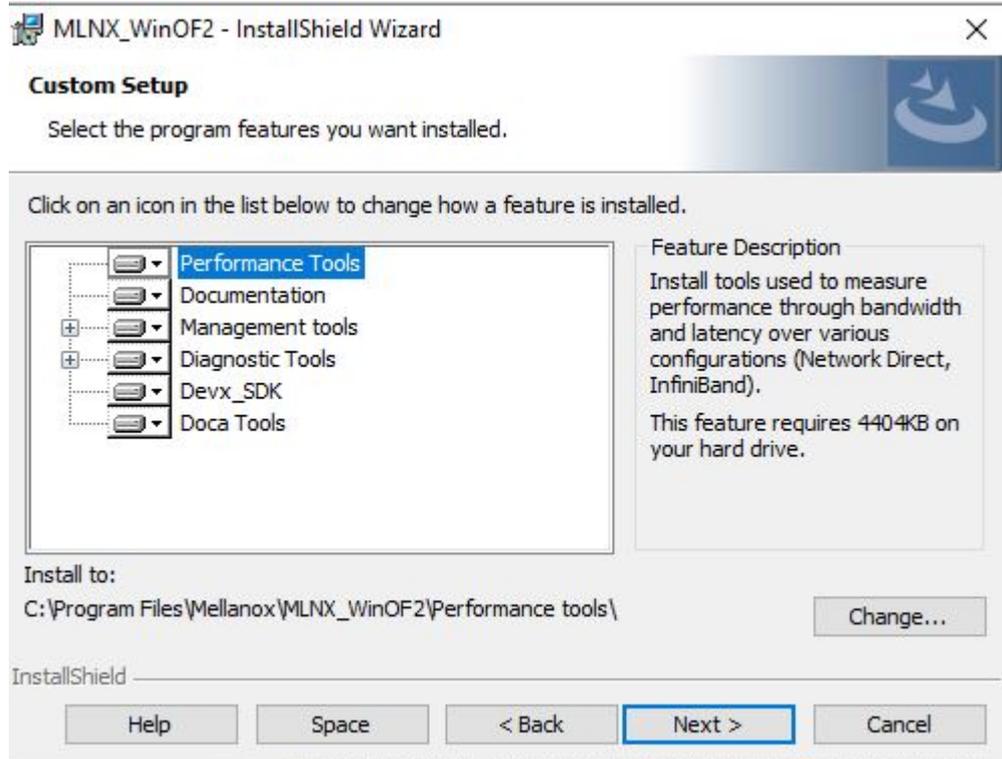
11. Click Install to start the installation.
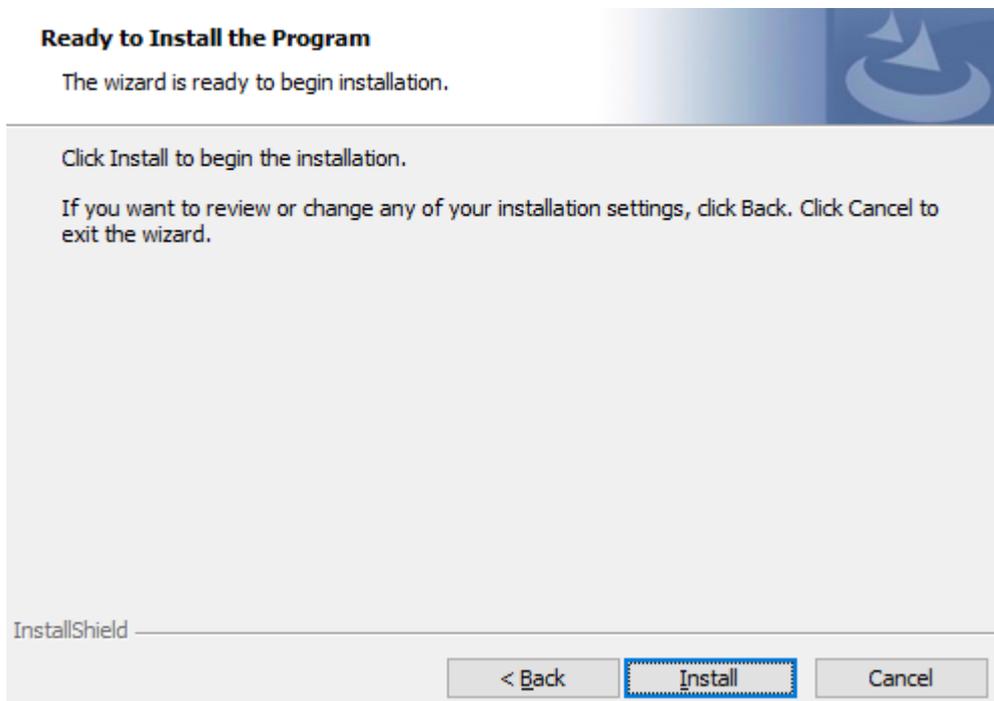
12. In case firmware upgrade option was checked in Step 7, you will be notified if a firmware upgrade is required (see 🛈).

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.

InstallShield Wizard Completed

The InstallShield Wizard has successfully installed MLNX_WinOF2. Click Finish to exit the wizard.

☐ Show readme file

< Back   Finish   Cancel

Unattended Installation

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

The following is an example of an unattended installation session.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Hardware and Software Requirements

Requirement	Description
Platforms	A server platform with an adapter card based on ConnectX®-4 Lx (EN) (firmware: fw-ConnectX4Lx)
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        4.16.8.8-1OEM.650.0.0.4240417    MEL    PartnerSupported 2017-01-31
nmlx5-rdma        4.16.8.8-1OEM.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-4 Lx
Part Number:       MCX4131A-GCAT
Description:       ConnectX®-4 Lx EN network interface card, 50GbE single-port QSFP28, PCIe3.0 x8, 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-4 Lx
Part Number:       MCX4131A-GCAT
Description:       ConnectX®-4 Lx EN network interface card, 50GbE single-port QSFP28, PCIe3.0 x8, 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	• Ensure that the adapter is placed correctly • Make sure the adapter slot and the adapter are compatible Install the adapter in a different PCI Express slot • Use the drivers that came with the adapter or download the latest • Make sure your motherboard has the latest BIOS • Try to reboot the server
The adapter no longer works	• Reseat the adapter in its slot or a different slot, if necessary • Try using another cable • Reinstall the drivers for the network driver files may be damaged or deleted • Reboot the server
Adapters stopped working after installing another adapter	• Try removing and re-installing all adapters • Check that cables are connected properly • Make sure your motherboard has the latest BIOS
Link indicator light is off	• Try another port on the switch • Make sure the cable is securely attached • Check you are using the proper cables that do not exceed the recommended lengths • Verify that your switch and adapter port are compatible
Link light is on, but with no communication established	• Check that the latest driver is loaded • Check that both the adapter and its link are set to the same speed and duplex settings
Event message received of insufficient power	• 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 slow") • 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	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
Card Detection	lspci | grep -i Mellanox
Mellanox Firmware Tool (MFT)	Download and install MFT: MFT Documentation Refer to the User Manual for installation instructions. Once installed, run: mst start mst status flint -d <mst_device> q
Ports Information	ibstat ibv_devinfo
Firmware Version Upgrade	To download the latest firmware version, refer to the NVIDIA Update and Query Utility.
Collect Log File	cat /var/log/messages dmesg >> system.log journalctl (Applicable on new operating systems) cat /var/log/syslog

Windows Troubleshooting

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

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

Specifications

MCX4111A-XCAT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile		
	Connector: Single SFP28 (copper and optical)		
Protocol Support	Ethernet: 10GBASE-SR, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, SGMII, 1000BASE-X		
	Data Rate: 1/10 Gb/s Ethernet		
	PCI Express Gen3: SERDES @ 8.0GT/s, 8 lanes (2.0 and 1.1 compatible)		
Power and Airflow	Voltage: 12V, 3.3V		
	Power	**Cable Type**	
	Typical Power[a]	Passive Cables	8.7W
	Maximum Power	Passive Cables	10.4W
		1W Active Cables	11.6W
	Maximum power available through SFP28 port: 1.5W		
	Airflow	See Airflow Specifications.	
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	
	RoHS	RoHS compliant	

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

MCX4121A-XCAT/MCX4121A-XCHT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile		
	Connector: Single SFP28 (copper and optical)		
Protocol Support	Ethernet: 10GBASE-SR, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, SGMII, 1000BASE-X		
	Data Rate: 1/10 Gb/s Ethernet		
	PCI Express Gen3: SERDES @ 8.0GT/s, 8 lanes (2.0 and 1.1 compatible)		
Power	Voltage: 12V, 3.3V		
	Power	**Cable Type**	
	Typical Power[a]	Passive Cables	9.3W
	Maximum Power	Passive Cables	11.1W
		1W Active Cables	13.3W
	Maximum power available through SFP28 port: 1.5W		

	Airflow	See Airflow Specifications.	
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	
	RoHS	RoHS compliant	

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

MCX4111A-ACAT and MCX4111A-ACUT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile		
	Connector: Single SFP28 (copper and optical)		
Protocol Support	Ethernet: 25GBASE-R, 20GBASE-KR2, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII		
	Data Rate: 1/10/25 Gb/s Ethernet		
	PCI Express Gen3: SERDES @ 8.0GT/s, 8 lanes (2.0 and 1.1 compatible)		
Power	Voltage: 12V, 3.3V		
	Power	**Cable Type**	
	Typical Power[a]	Passive Cables	8.9W
	Maximum Power	Passive Cables	10.7W
		1.5W Active Cables	12.3W
	Maximum power available through SFP28 port: 1.5W		
	Airflow	See Airflow Specifications.	
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	
	RoHS	RoHS compliant	

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

MCX4121A-ACAT / MCX4121A-ACHT / MCX4121A-ACST / MCX4121A-ACUT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile
	Connector: Dual SFP28 (copper and optical)
Protocol Support	Ethernet: 25GBASE-R, 20GBASE-KR2, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII
	Data Rate: 1/10/25 Gb/s Ethernet
	PCI Express Gen3: SERDES @ 8.0GT/s, 8 lanes (2.0 and 1.1 compatible)

Power and Airflow	Voltage: 12V, 3.3V		
	Power	**Cable Type**	
	Typical Power[a]	Passive Cables	9.6W
	Maximum Power	Passive Cables	11.6W
		1W Active Cables	15.0W
	Maximum power available through SFP28 port: 1.5W		
	Airflow	See Airflow Specifications.	

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
	RoHS	RoHS compliant

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

MCX4131A-BCAT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile
	Connector: Single QSFP28 (copper and optical)
Protocol Support	Ethernet: 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-CR/CR-S, 25GBASE-KR/KR-S, 25GBASE-SR, 25GBASE-LR, 25G-R (EC), 20GBASE-KR2, 10GBASE-SR, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 1000BASE-KX, SGMII / 1000BASE-X
	Data Rate: 1/10/25/40 Gb/s Ethernet
	PCI Express Gen3: SERDES @ 8.0GT/s, 8 lanes (2.0 and 1.1 compatible)

Power and Airflow	Voltage: 12V, 3.3V		
	Power	**Cable Type**	
	Typical Power[a]	Passive Cables	10.6W
	Maximum Power	Passive Cables	12.6W
		1.5W Active Cables	14.2W
	Maximum power available through SFP28 port: 1.5W		
	Airflow	See Airflow Specifications.	
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	
	RoHS	RoHS compliant	

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

MCX4131A-GCAT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile		
	Connector: Single QSFP28 (copper and optical)		
Protocol Support	Ethernet: 50G-R2 (EC), 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-CR/CR-S, 25GBASE-KR/KR-S, 25GBASE-SR, 25GBASE-LR, 25G-R (EC), 20GBASE-KR2, 10GBASE-SR, 10GBASE-LR, 10GBASE-ER,10GBASE-CX4, 1000BASE-KX, SGMII / 1000BASE-X		
	Data Rate: 1/10/25/40 Gb/s Ethernet		
	PCI Express Gen3: SERDES @ 8.0GT/s, 8 lanes (2.0 and 1.1 compatible)		
Power and Airflow	Voltage: 12V, 3.3V		
	Power	**Cable Type**	
	Typical Power[a]	Passive Cables	9.6W
	Maximum Power	Passive Cables	11.6
		3.5W Active Cables	15.4W
	Maximum power available through SFP28 port: 1.5W		
	Airflow	See Airflow Specifications.	
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
	RoHS	RoHS compliant

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

Airflow Specifications

10Gb/s Adapter Cards Airflow

	Airflow (LFM)			
Airflow Direction	Heatsink to Port		Port to Heatsink	
OPN	Passive	Active 1W	Passive	Active 1W
MCX4111A-XCAT	300	300	300	300
MCX4121A-XCAT	300	300	300	300

25Gb/s Adapter Cards Airflow

	Airflow (LFM)			
Airflow Direction	Heatsink to Port		Port to Heatsink	
OPN	Passive	Active 1.5W	Passive	Active 1.5W
MCX4111A-ACAT	300	400	300	300
MCX4121A-ACAT	300	400	300	300
MCX4111A-ACUT	300	400	300	300
MCX4121A-ACUT	300	400	300	300
MCX4121A-ACST	300	400	300	300
MCX4121A-ACHT	300	400	300	300

40/50Gb/s Adapter Cards Airflow[a]

	Airflow (LFM)				
Airflow Direction	Heatsink to Port			Port to Heatsink	
OPN	Passive	Active 1.5W	Active 3.5W	Passive	Active 1.5W
MCX413A-BCAT	300	350	-	300	300
MCX413A-GCAT	300	-	500	300	-

a. Cells marked with (–) are not applicable

Board Mechanical Drawing and Dimensions

⚠ All PSB dimensions are in millimeters. All the mechanical tolerances are +/- 0.13mm.

Single-port 10/25 Gb/s	Dual-port 10/25 Gb/s	Single-port 40/50 Gb/s
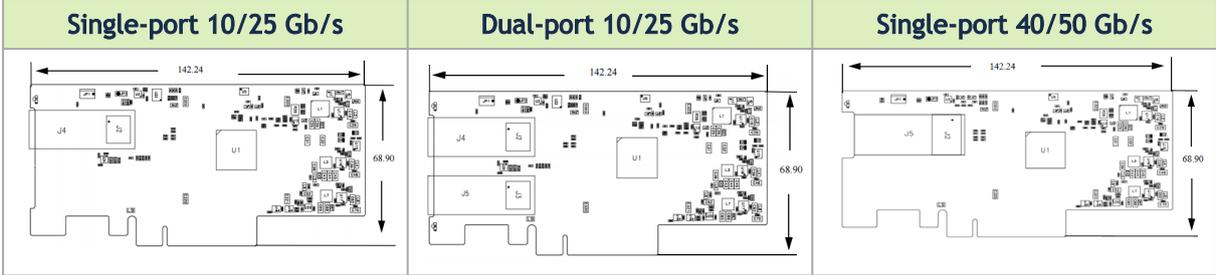
Bracket Mechanical Drawing

⚠ All bracket dimensions are in millimeters. All the mechanical tolerances are +/- 0.25mm.

Single-port Tall Bracket	Dual-port Tall Bracket
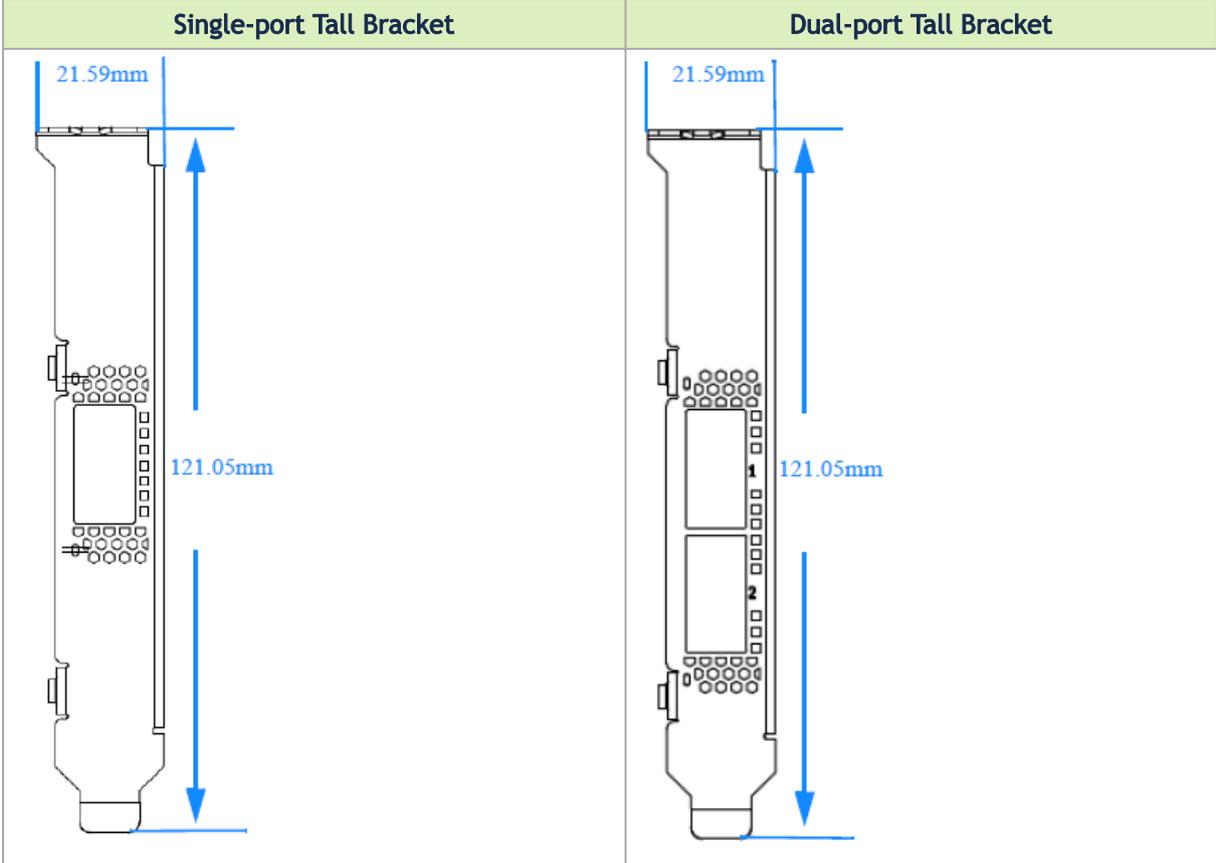
Finding the MAC on the Adapter Card

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.

MCX4111A-XCAT Boarding Label (Example)

```
Model No: CX4111A

ConnectX-4 Lx 10GbE
P/N: MCX4111A-XCAT
                              2019-05-06
|||||||||||||||||||||||||       Rev: AE
S/N: MT0806X01504

MAC: 00 02 C9 27 05 1C

Made in Israel
```

MCX4111A-ACAT Board Label (Example)

```
Model No: CX4111A

ConnectX-4 Lx 25GbE
P/N: MCX4111A-ACAT
                              2019-05-06
|||||||||||||||||||||||||       Rev: AG
S/N: MT0806X01504

MAC: 00 02 C9 27 05 1C

Made in Israel
```

MCX4121A-ACAT Board Label (Example)

Model No: CX4121A

ConnectX-4 Lx 25GbE
P/N: MCX4121A-ACAT

S/N: MT0806X01504

MAC: 00 02 C9 27 05 1C

Made in Israel

2019-05-06
Rev: AG

MCX4121A-XCHT Board Label (Example)

TBD

MCX4111A-ACUT Board Label (Example)

Model No: CX4111A

ConnectX-4 Lx 25GbE
P/N: MCX4111A-ACUT

2019-05-06
Rev: A1

S/N: MT0806X01504

MAC: 00 02 C9 27 05 1C

Made in Israel

MCX4121A-ACUT Board Label (Example)

Model No: CX4121A

ConnectX-4 Lx 25GbE
P/N: MCX4121A-ACUT

2019-05-06
Rev: A1

S/N: MT0806X01504

MAC: 00 02 C9 27 05 1C

Made in Israel

MCX4121A-ACST Board Label (Example)

Model No: CX4121A

ConnectX-4 Lx 25GbE
P/N: MCX4121A-ACST

2019-05-06
Rev: A3

S/N: MT0806X01504

MAC: 00 02 C9 27 05 1C

Made in Israel

MCX4131A-BCAT Board Label (Example)

Model No: CX4131A

ConnectX-4 LX 40GbE
P/N: MCX4131A-BCAT

2019-05-06
Rev: AG

S/N: MT0806X01504

MAC: 00 02 C9 27 05 1C

Made in Israel

MCX4131A-GCAT Board Label (Example)

Model No: CX4131A

ConnectX-4 LX 50GbE
P/N: MCX4131A-GCAT

S/N: MT0806X01504

MAC: 00 02 C9 27 05 1C

Made in Israel

2019-05-05
Rev: AJ

Document Revision History

Date	Description of Changes
Jan. 2024	Added a note that the ConnectX-4 Lx product line has moved to end-of-life.
May. 2023	Added non-operational storage temperature specifications.
Nov. 2020	Added MCX4121A-ACHT support across the document.
Feb. 2020	Added altitude criteria to "Specifications".
Oct. 2019	Updated Ethernet supported protocols in "Specifications".
Sep. 2019	Added MCX4121A-XCHT support across the document.
May. 2019	Updated bracket and PCB tolerance in "Specifications".
May. 2019	Updated board labels and added a note to "Finding the MAC and Serial Number on the Adapter Card"
Dec. 2018	- Migrated to on-line format; minor reorganization. - Added the following OPNs to relevant sections: - MCX4111A-ACUT - MCX4121A-ACUT - MCX4121A-ACST - MCX4131A-BCAT - MCX4131A-GCAT - Updated Features
Nov. 2017	- Updated Features on page 15. - Added a note to System Requirements on page 19.
Jul. 2017	- Updated Adapter LEDs Operations on page 62. - Updated "Cable Installation" on page 24.
Sep. 2016	- Updated Adapter Card Installation Instructions on page 19. - Added Card Un-installation on page 25. - Updated the following tables: - MXC4111A-XCAT Specifications Table on page 60. - MCX4121A-XCAT Specifications Table on page 58. - MCX4111A-AC[A/U]T Specifications Table on page 62. - MCX4121A-AC[S/U/A]T Specifications on page 63.
Aug. 2016	Added MCX512F-ACAT and MCX512A-ACAT support across the document
Jun. 2016	Added a note to System Requirements.
Apr. 2016	Updated Physical and Logical Link Indications
Mar. 2016	- Removed I2C-compatible Interface from Interfaces - Updated Features and Benefits
Feb. 2016	Updated Bracket Installation Instructions on page 21
Feb. 2016	- Updated Linux Driver - Updated Windows Driver - Updated VMware Driver - Added a note to Finding the MAC and Serial Number on the Adapter Card
Dec. 2015	- Updated Product Overview - Updated speed specifications in the following sections: - Features and Benefits - Connectivity - Updated Specifications

Date	Description of Changes
Dec. 2015	• Added mlxup to Documents List on page 9 and Updating Adapter Card Firmware • Updated Device ID in Product Overview • Added IEEE 802.3by support. See "Features and Benefits" • Updated PCI Device • Added VMware support. See VMware Driver • Updated Specifications • Added Single-port Bracket
Nov. 2015	First release