

NVIDIA ConnectX-4 InfiniBand/Ethernet Adapter Cards User Manual

Table of contents

Introduction 5

Interfaces 10

Hardware Installation 13

Driver Installation 30

- Linux Driver Installation 30
- Windows Driver Installation 54

Troubleshooting 66

Updating Adapter Firmware 69

Specifications 71

Document Revision History 83

About This Manual

This User Manual describes NVIDIA® ConnectX®-4 InfiniBand/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.

Ordering Part Numbers

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

NVIDIA SKU	Legacy OPN	Marketing Description
900-9X426-0014-ST0	MCX453A-FCAT	ConnectX®-4 InfiniBand/Ethernet adapter card, FDR IB 40/56GbE, single-port QSFP28, PCIe3.0 x8, tall bracket
900-9X426-0054-ST0	MCX454A-FCAT	ConnectX®-4 InfiniBand/Ethernet adapter card, FDR IB and 40/56GbE, dual-port QSFP28, PCIe3.0 x8, tall bracket
900-9X4AC-0014-ST0	MCX455A-FCAT	ConnectX®-4 InfiniBand/Ethernet adapter card, FDR IB and 40/56GbE, single-port QSFP28, PCIe3.0 x16, tall bracket
900-9X4AC-0054-ST0	MCX456A-FCAT	ConnectX®-4 InfiniBand/Ethernet adapter card, FDR IB and 40/56GbE, dual-port QSFP28, PCIe3.0 x16, tall bracket
900-9X4AC-0016-ST0	MCX455A-ECAT	ConnectX®-4 InfiniBand/Ethernet adapter card, EDR IB (100Gb/s) and 100GbE, single-port QSFP28, PCIe3.0 x16, tall bracket
900-9X4AC-0056-ST3	MCX456A-ECAT	ConnectX®-4 InfiniBand/Ethernet adapter card, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, PCIe3.0 x16, tall bracket

Intended Audience

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

Technical Support

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

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

MLNX_OFED for Linux User Manual and Release Notes	User Manual describing OFED features, performance, band diagnostic, tools content and configuration. See MLNX_OFED for Linux Documentation.
WinOF-2 for Windows User Manual and Release Notes	User Manual describing WinOF-2 features, performance, Ethernet diagnostic, tools content and configuration. See WinOF-2 for Windows Documentation.
NVIDIA VMware for Ethernet User Manual	User Manual and release notes describing the various components of the NVIDIA ConnectX® NATIVE ESXi stack. See VMware® ESXi Drivers Documentation.
NVIDIA Firmware Utility (mlxup) User Manual and Release Notes	NVIDIA firmware update and query utility used to update the firmware. Refer to Firmware Utility (mlxup) Documentation.
NVIDIA Firmware Tools (MFT) User Manual	User Manual describing the set of MFT firmware management tools for a single node. See MFT User Manual.
InfiniBand Architecture Specification Release 1.2.1	InfiniBand Specifications
IEEE Std 802.3 Specification	IEEE Ethernet Specifications

PCI Express 3.0 Specifications	Industry Standard PCI Express Base and Card Electromechanical Specifications. Refer to PCI-SIG Specifications.
LinkX Interconnect Solutions	LinkX InfiniBand cables and transceivers are designed to maximize the performance of High-Performance Computing networks, requiring high-bandwidth, low-latency connections between compute nodes and switch nodes. NVIDIA offers one of the industry's broadest portfolios of QDR/FDR10 (40Gb/s), FDR (56Gb/s), EDR/HDR100 (100Gb/s), HDR (200Gb/s), and NDR (400Gb/s) cables, including Direct Attach Copper cables (DACs), copper splitter cables, Active Optical Cables (AOCs) and transceivers in a wide range of lengths from 0.5m to 10km. In addition to meeting IBTA standards, NVIDIA tests every product in an end-to-end environment ensuring a Bit Error Rate of less than 1E-15.Read more at LinkX Cables and Transceivers.

Document Conventions

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

Introduction

Product Overview

This is the User Guide for InfiniBand/Ethernet adapter cards based on the NVIDIA® ConnectX®-4 integrated circuit device. These adapters connectivity provide the highest performing and most flexible interconnect solution for PCI Express Gen 3.0/4.0 servers used in Enterprise Data Centers, High-Performance Computing, and Embedded environments.

The following table provides the ordering part number, port speed, number of ports, and PCI Express speed. Each adapter comes with two bracket heights - short and tall.

ConnectX-4 InfiniBand/Ethernet FDR Adapter Cards

Part Number	MCX453A-FCAT	MCX455A-FCAT	MCX454A-FCAT	MCX456A-FCAT
Data Rate	InfiniBand: SDR/DDR/QDR/FDR Ethernet: 1/10/25/40/56 Gb/s			
Network Connector Type	Single-port QSFP28		Dual-port QSFP28	
PCI Express Connectors	PCIe 3.0 x8 @ 8GT/s	PCIe 3.0 @ 8GT/s x16	PCIe 3.0 x8 @ 8GT/s	PCIe 3.0 @ 8GT/s x16
Dimensions	2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile			
RoHS	RoHS Compliant			
Adapter IC Part Number	MT27704A0-FDCF-FV			
Device ID (decimal)	4115 for Physical Function (PF) 4116 for Virtual Function (VF)			

ConnectX-4 InfiniBand/Ethernet EDR Adapter Cards

Part Number	MCX455A-ECAT	MCX456A-ECAT
Data Rate	InfiniBand: SDR/DDR/QDR/FDR/EDR Ethernet: 10/25/40/50/100 Gb/s	
Network Connector Type	Single-port QSFP28	Dual-port QSFP28
PCI Express Connectors	PCIe Gen 3.0 x16; SerDes @ 8.0GT/s	
Dimensions	2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile	
RoHS	RoHS Compliant	
Adapter IC Part Number	MT27704A0-FDCF-EV	
Device ID (decimal)	4115 for Physical Function (PF) 4116 for Virtual Function (VF)	

For more detailed information see Specifications.

Features and Benefits

> ⚠ **Warning**
>
> 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 or x16 edge connector. Gen 1.1 and 2.0 compatible.
EDR InfiniBa	A standard InfiniBand data rate, where each lane of a 4X port runs a bit rate of 25.78125Gb/s with a 64b/66b encoding, resulting in an effective bandwidth

Feature	Description
nd	of 100Gb/s.
100Gb/s Adapter	ConnectX-4 offers the highest throughput InfiniBand/Ethernet adapter, supporting EDR 100Gb/s InfiniBand and 100Gb/s Ethernet and enabling any standard networking, clustering, or storage to operate seamlessly over any converged network leveraging a consolidated software stack.
InfiniBand Architecture Specification v1.3 compliant	ConnectX-4 delivers low latency, high bandwidth, and computing efficiency for performance-driven server and storage clustering applications. ConnectX-4 is InfiniBand Architecture Specification v1.3 compliant.
Up to 100 Gigabit Ethernet	NVIDIA adapters comply with the following IEEE 802.3 standards: • 100GbE/ 50GbE / 40GbE / 25GbE / 10GbE / 1GbE • 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 802.3ae 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	• SPI - includes one16MB SPI Flash device(M25PX16-VMN6P device by ST Microelectronics) • FRU EEPROM - Stores the parameters and personality of the card. The EEPROM capacity is 128Kbit. FRU I2C address is (0x50) and is accessible

Feature	Description
	through the PCIe SMBus. (Note: Address 0x58 is reserved.)
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 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, 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 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 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 provides dedicated adapter resources and guaranteed isolation and protection for virtual machines within the server.
Storage Accelera	A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file

Feature	Description
tion	access protocols can leverage InfiniBand RDMA for high-performance storage access. • NVMe over Fabric offloads for the target machine
SR-IOV	ConnectX-4 SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server.

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

Interfaces

InfiniBand Interface

The network ports of the ConnectX®-4 adapter cards are compliant with the InfiniBand Architecture Specification, Release 1.3. InfiniBand traffic is transmitted through the cards' QSFP28 connectors.

Ethernet QSFP28 Interface

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

PCI Express Interface

The ConnectX®-4 adapter card supports PCI Express Gen 3.0 (1.1 and 2.0 compatible) through an x8 or x16 edge connector. The device can be either a master initiating the PCI Express bus operations, or a subordinateresponding to PCI bus operations. The following lists PCIe interface features:

- PCIe 3.0 compliant, 2.0 and 1.1 compatible
- 2.5, 5.0, or 8.0, link rate x8/x16
- Auto-negotiates to x16, 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. LED behavior is described below for Ethernet and InfiniBand port configurations.

Physical and Logical Link Indications - Ethernet Protocol:

<table>
<tr><th>LED Color and State</th><th>Description</th></tr>
<tr><td>Off</td><td>A link has not been established</td></tr>
<tr><td>Beacon command for locating the adapter card</td><td>1Hz blinking Yellow</td></tr>
<tr><td>Error</td><td>4Hz blinking Yellow Indicates an error with the link. The error can be one of the following:

<table><tr><th>Error Type</th><th>Description</th><th>LED Behavior</th></tr><tr><td>I^2C</td><td>I^2C access to the networking ports fails</td><td>Blinks until error is fixed</td></tr><tr><td>Over-current</td><td>Over-current condition of the networking ports</td><td>Blinks until error is fixed</td></tr></table></td></tr>
<tr><td>Solid green</td><td>Indicates a valid link with no active traffic</td></tr>
<tr><td>Blinking green</td><td>Indicates a valid logical link with active traffic</td></tr>
</table>

Physical and Logical Link Indications - InfiniBand Protocol:

LED Color and State	Description
Off	A physical link has not been established
Solid amber	Indicates an active physical link
Beacon command for locating the adapter card	1Hz blinking Yellow
Error	4Hz blinking Yellow Indicates an error with the link. The error can be one of the following:

<table>
<tr><th>LED Color and State</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><th>Error Type</th><th>Description</th><th>LED Behavior</th></tr>
<tr><td>I^2C</td><td>I^2C access to the networking ports fails</td><td>Blinks until error is fixed</td></tr>
<tr><td>Over-current</td><td>Over-current condition of the networking ports</td><td>Blinks until error is fixed</td></tr>
</table>
</td></tr>
<tr><td>Solid green</td><td>Indicates a valid logical (data activity) link with no active traffic</td></tr>
<tr><td>Blinking green</td><td>Indicates a valid logical link with active traffic</td></tr>
</table>

Hardware Installation

Installation and initialization of ConnectX-4 adapter cards require attention to the mechanical attributes, power specification, and precautions for electronic equipment.

Safety Warnings

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

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

Note that not all warnings are relevant to all models.

	**General Installation Instructions** Read all installation instructions before connecting the equipment to the power source.
	Jewelry Removal Warning Before you install or remove equipment that

	is connected to power lines, remove jewelry such as bracelets, necklaces, rings, watches, and so on. Metal objects heat up when connected to power and ground and can meltdown, causing serious burns and/or welding the metal object to the terminals.
	**Over-temperature** This equipment should not be operated in an area with an ambient temperature exceeding the maximum recommended: 55°C (131°F). An airflow of 200LFM at this maximum ambient temperature is required for HCA cards and NICs. To guarantee proper airflow, allow at least 8cm (3 inches) of clearance around the ventilation openings.
	**During Lightning - Electrical Hazard** During periods of lightning activity, do not work on the equipment or connect or dis-connect cables.
	Copper Cable Connecting/Disconnecting Some copper cables are heavy and not flexible, as such, they should be carefully

	attached to or detached from the connectors. Refer to the cable manufacturer for special warnings and instructions.
	**Equipment Installation** This equipment should be installed, replaced, or serviced only by trained and qualified personnel.
	**Equipment Disposal** The disposal of this equipment should be in accordance to all national laws and regulations.

	**Local and National Electrical Codes** This equipment should be installed in compliance with local and national electrical codes.
	**Hazardous Radiation Exposure** • Caution – Use of controls or adjustment or performance of procedures other than those specified herein may result in hazardous radiation exposure.For products with optical ports. • CLASS 1 LASER PRODUCT and reference to the most recent laser standards: IEC 60 825-1:1993 + A1:1997 + A2:2001 and EN 60825-1:1994+A1:1996+ A2:20

Installation Procedure Overview

The installation procedure of ConnectX-4 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

Step	Procedure	Direct Link
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 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 adapter card in the system.	Refer to Identifying the Card in Your System

System Requirements

Hardware Requirements

Important

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 or x16 edge connector is required for installing the card.

Important

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 adapter cards are offered with two airflow patterns: from the heatsink to the network ports, and vice versa.

Please refer to the Specifications chapter for airflow numbers for each specific card model.

Important

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

Important

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:
    1. ConnectX-4 Adapter Card, with a tall bracket assembled on the card.
    2. Short bracket (accessory)
2. Check the parts for visible damage that may have occurred during shipping.

Important

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.



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 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.

Important

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 card.

Important

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.

Important

Do not force the bracket onto the adapter card.

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

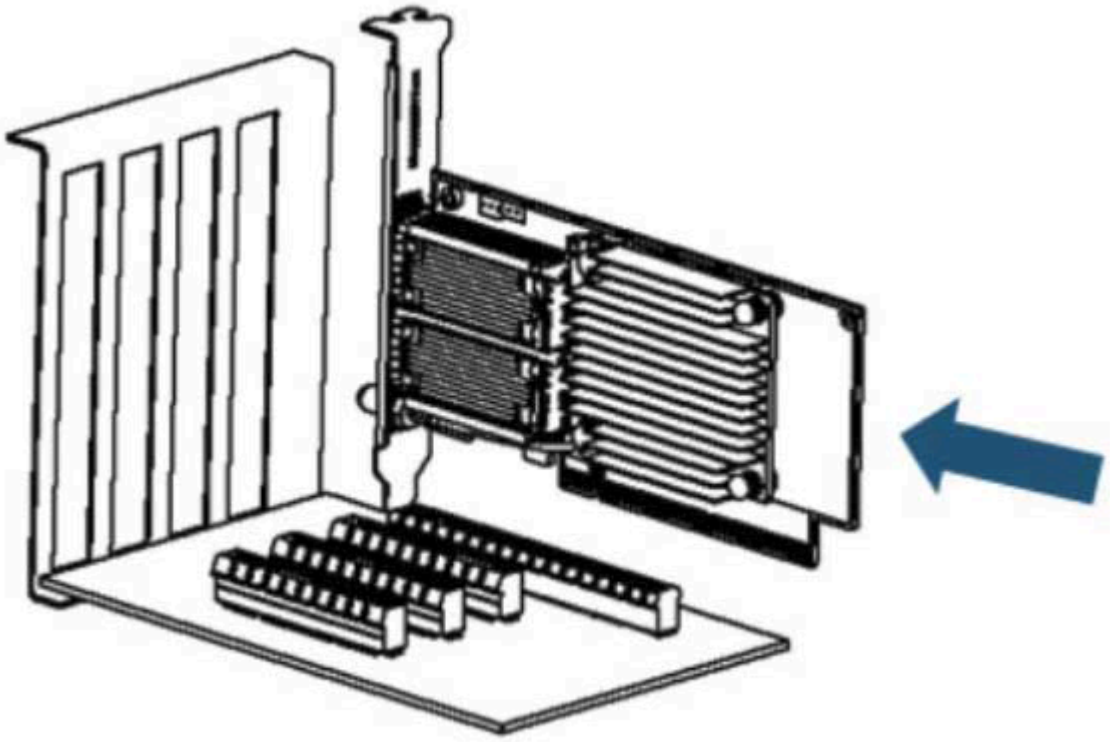
> **Important**
>
> Use a torque driver to apply up to 2 lbs-in torque on the screws.

Installation Instructions

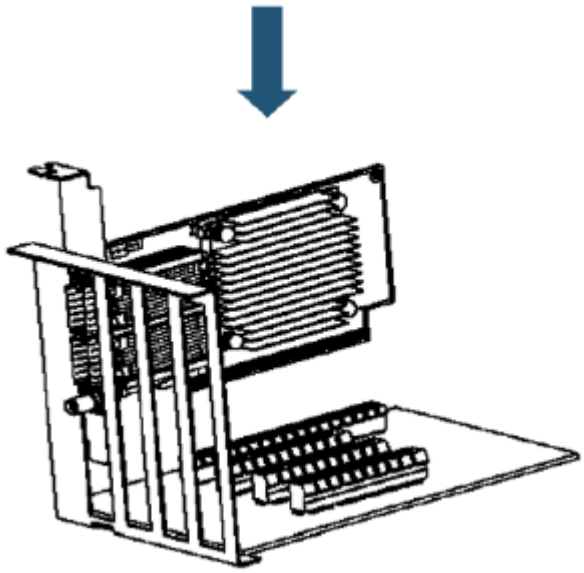
This section provides detailed instructions on how to install your adapter card in a system.

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

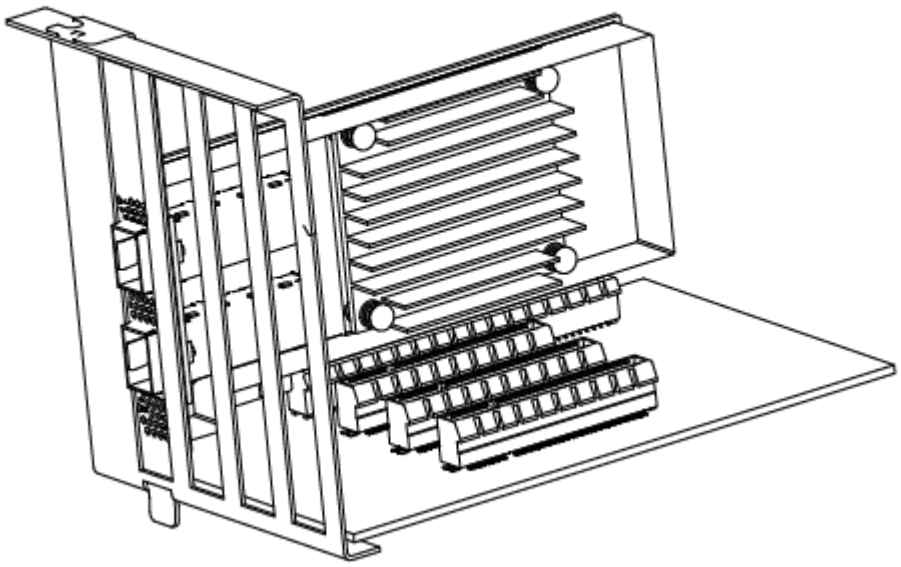
1. Before installing the card, make sure that the system is off and the power cord is not connected to the server. Please follow proper electrical grounding procedures.
2. Open the system case.
3. Place the adapter in an available PCI Express slot.



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.

> **Note**
>
> 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.
    1. 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.

2. 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.

3. 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.

4. Make sure that the connector locks in place.

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

> ⓘ **Important**
>
> 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 MT4115 Family [ConnectX-4]
```

On Windows

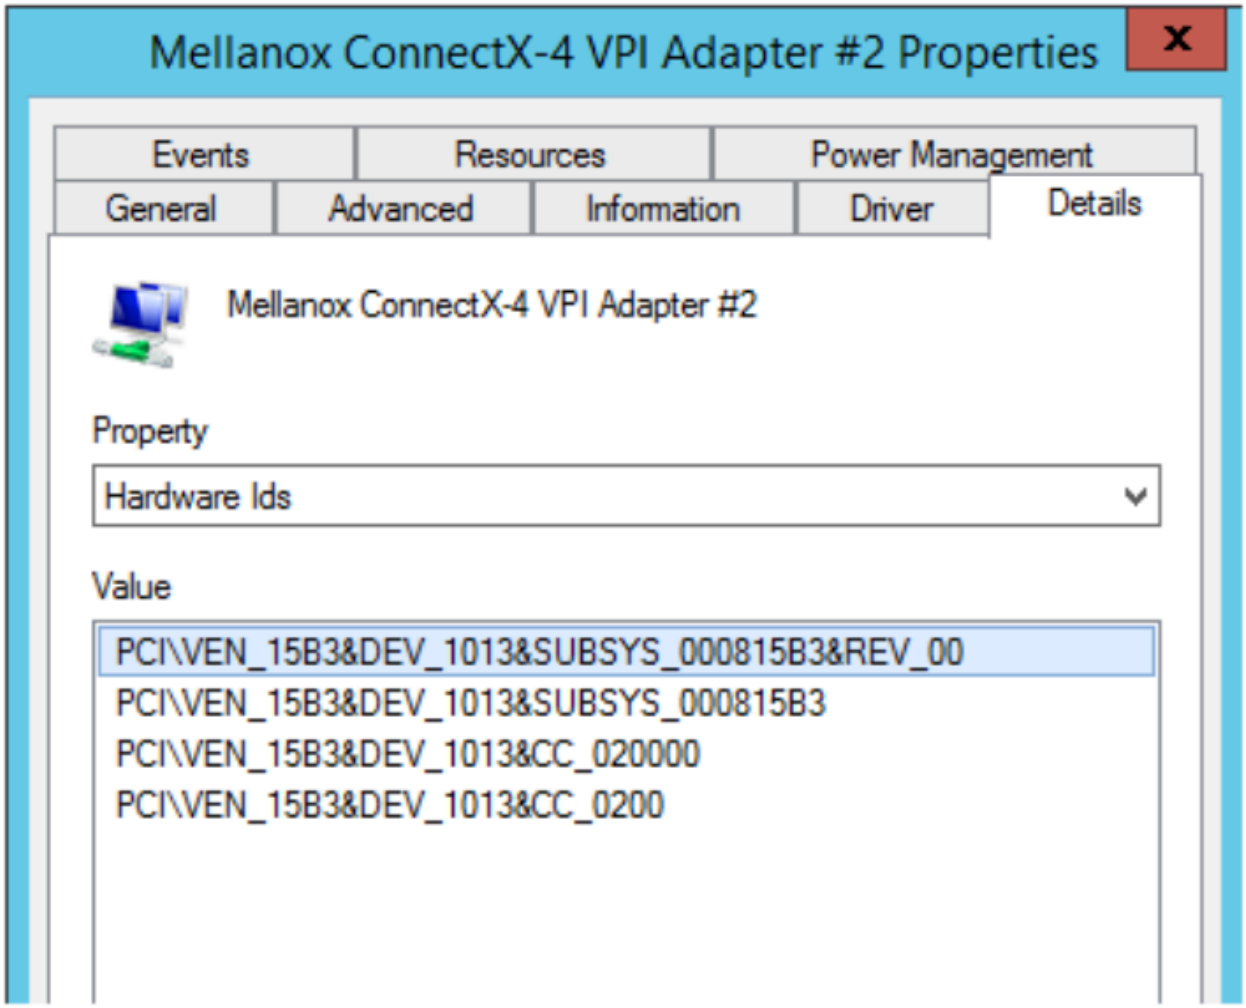
1. Open Device Manager on the server. Click **Start** => **Run**, and then enter **devmgmt.msc**.

2. Expand **System Devices** and locate your NVIDIA ConnectX-4 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 which is the Vendor ID of NVIDIA; and DEV is equal to 1018 (for ConnectX-4) which is a valid NVIDIA PCIe Device ID.

> ⚠ **Warning**
>
> If the PCIe device does not have an NVIDIA adapter ID, return to Step 2 to check another device.

> ⚠ **Warning**

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.

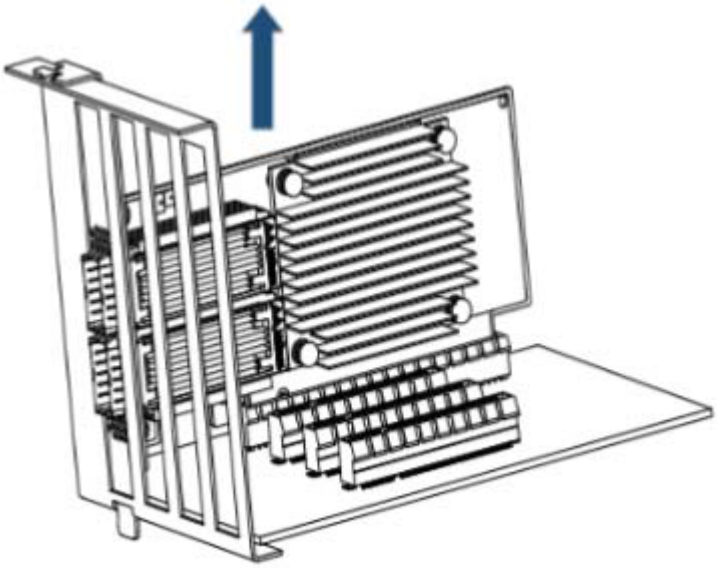
Card Removal

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

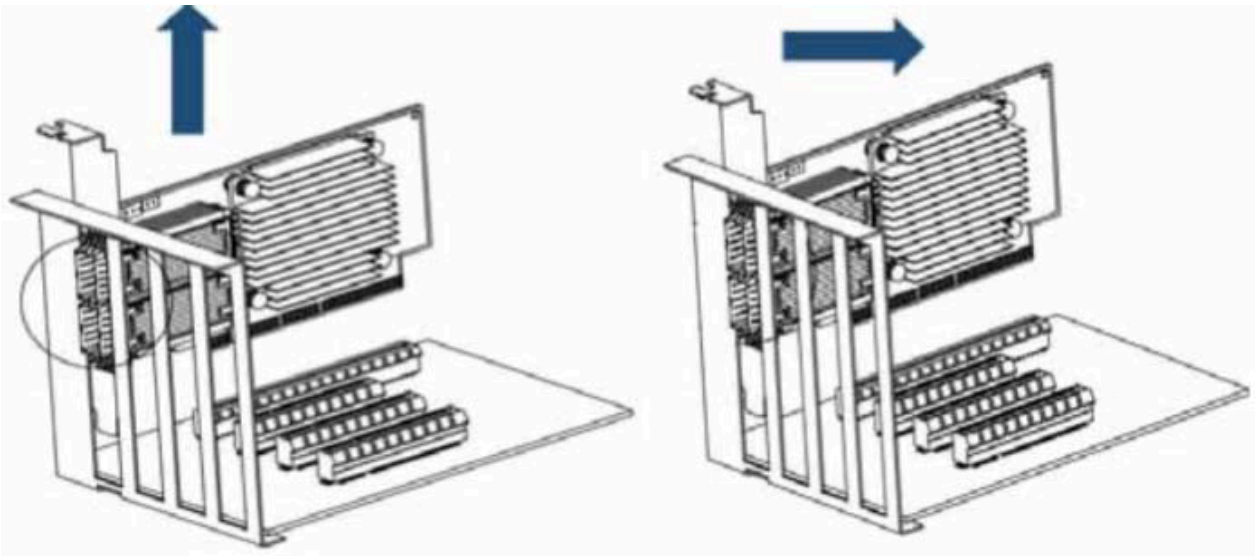
Note

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

This section describes how to install and test the MLNX_OFED for Linux package on a single server with an ConnectX-4 adapter card installed.

Prerequisites

Requirements	Description
Platforms	A server platform with one ConnectX-4 (VPI, IB, EN) (firmware: fw-ConnectX4)
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*
Installer Privileges	The installation requires administrator (root) privileges on the target machine.

Downloading MLNX_OFED

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

```
# lspci -v | grep Mellanox
86:00.0 Network controller [0207]: Mellanox Technologies MT27620 Family
```

```
        Subsystem: Mellanox Technologies Device 0014
86:00.1 Network controller [0207]: Mellanox Technologies MT27620 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

1.

    1. Scroll down to the Download wizard, and click the Download tab.
    2. Choose your relevant package depending on your host operating system.
    3. Click the desired ISO/tgz package.
    4. 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
  ```

> ⚠ **Warning**
>
> **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.

> ⚠ **Warning**

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.

⚠ Warning

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.

⚠ Warning

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.

⚠ Warning

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 ...
```

⚠ Warning

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

⚠ Warning

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 --forceFor example:./mlnxofedinstall --without-dkms --add-kernel-support --kernel 3.13.0-85-generic --without-fw-update --forceNote that the path to kernel sources (--kernel-sources) should be added if the sources are not in their default location.

⚠ Warning

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-XPart Number: MCXXXX-XXXPSID: MT_2190110032PCI Device Name: 0b:00.0Base MAC: 0000e41d2d5cf810Versions: Current AvailableFW 12.14.0114 12.14.0114Status: Up to date

⚠ Warning

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-XPart Number: MCXXXX-XXXPSID: MT_2190110032PCI Device Name: 0b:00.0Base MAC: 0000e41d2d5cf810Versions: Current AvailableFW 12.14.0114 N/AStatus: 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.

> ⚠ **Warning**
>
> 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.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-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!
```

> ⚠ **Warning**
>
> 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
```

> ⚠ **Warning**
>
> 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.

    1. To go to the Start menu, position your mouse in the bottom-right corner of the Remote Desktop of your screen.

    2. Open a CMD console (Click Task Manager-->File --> Run new task and enter CMD).

    3. Enter the following command.

        ```
        echo %PROCESSOR_ARCHITECTURE%
        ```

        > **⚠ Warning**
        >
        > 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.

    > **⚠ Warning**
    >
    > 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]"
   ```

   Example:

   ```
   MLNX_WinOF2-2_10_50000_All_x64.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.

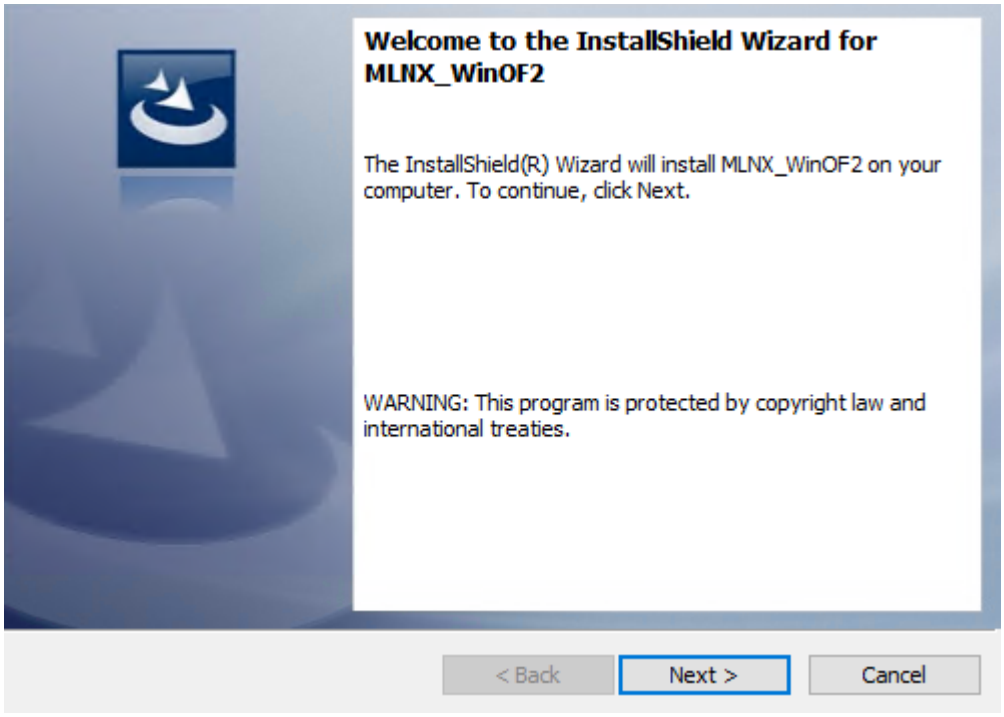
```
MLNX_WinOF2_<revision_version>_All_Arch.exe /v" MT_DISABLE_RSHIM_INSTALL=1"
```

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

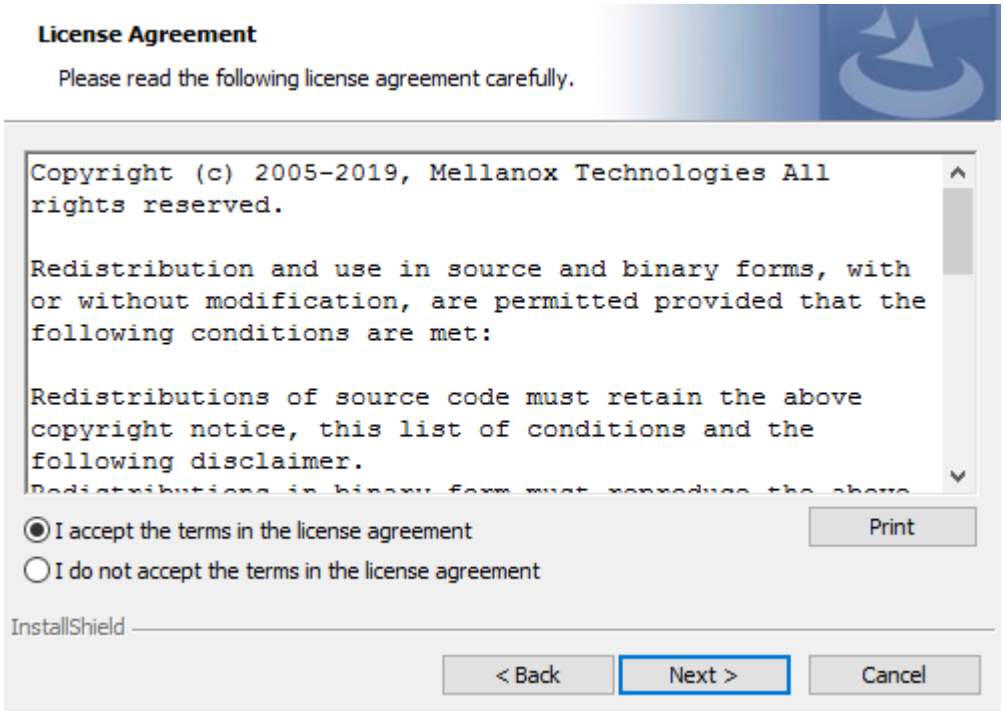
5. **[Optional]** If you want to skip the check for unsupported devices, run.

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

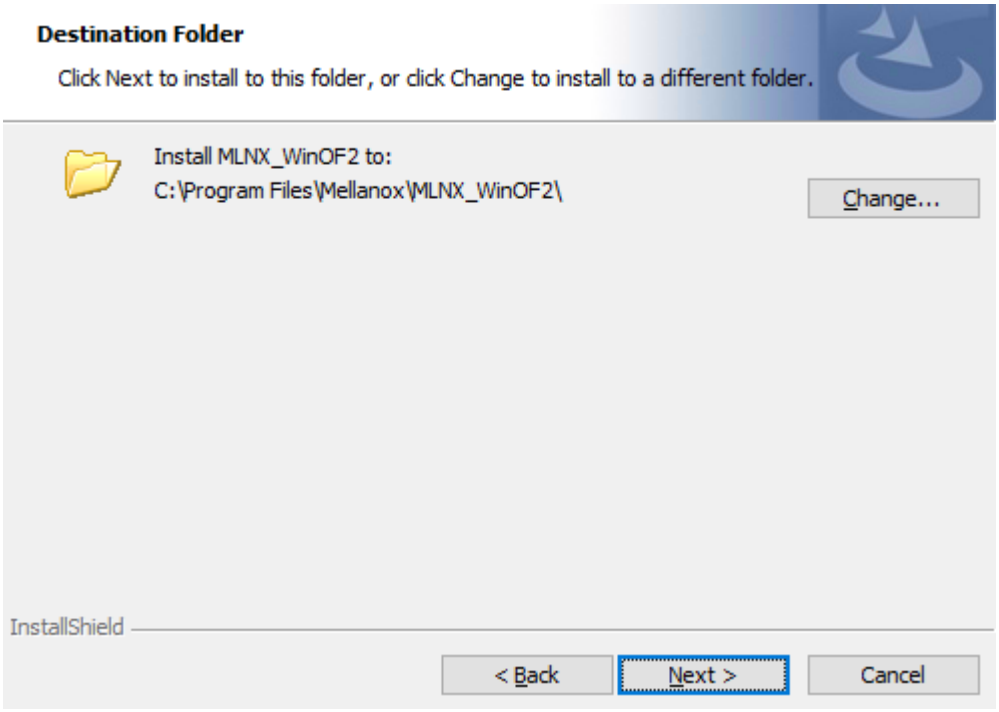
6. Click Next in the Welcome screen.



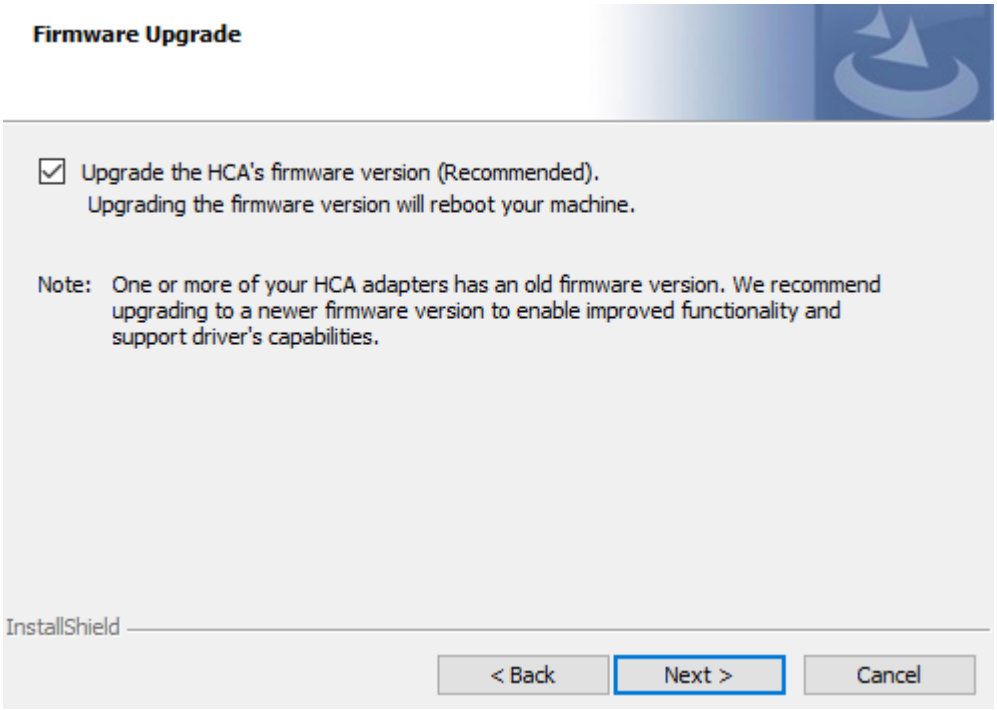
7. Read and accept the license agreement and click Next.



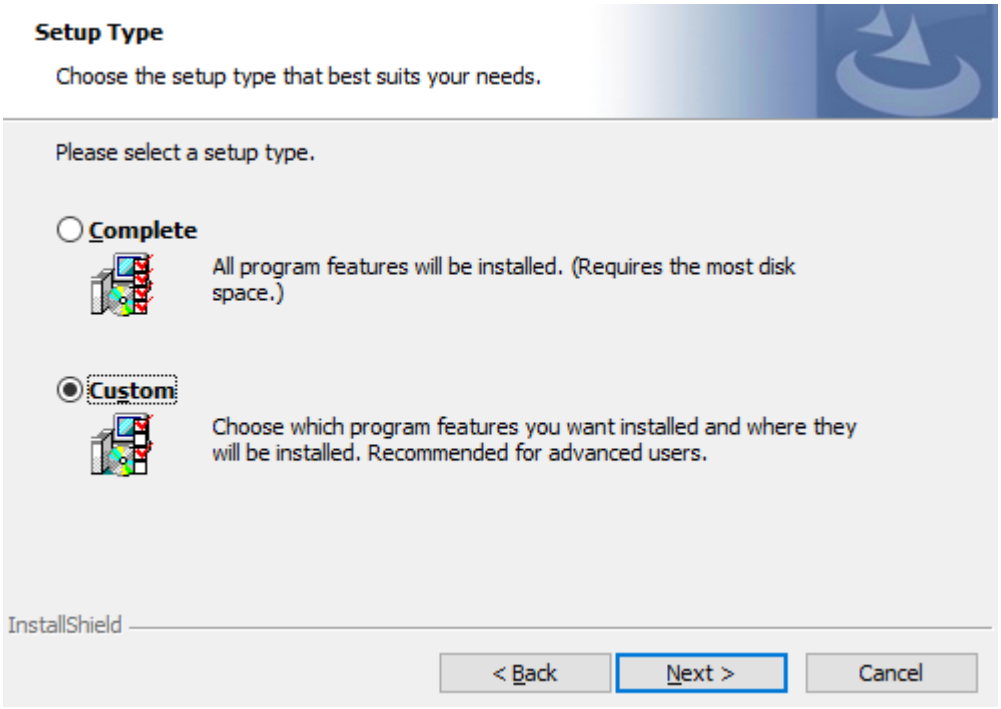
8. Select the target folder for the installation.



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



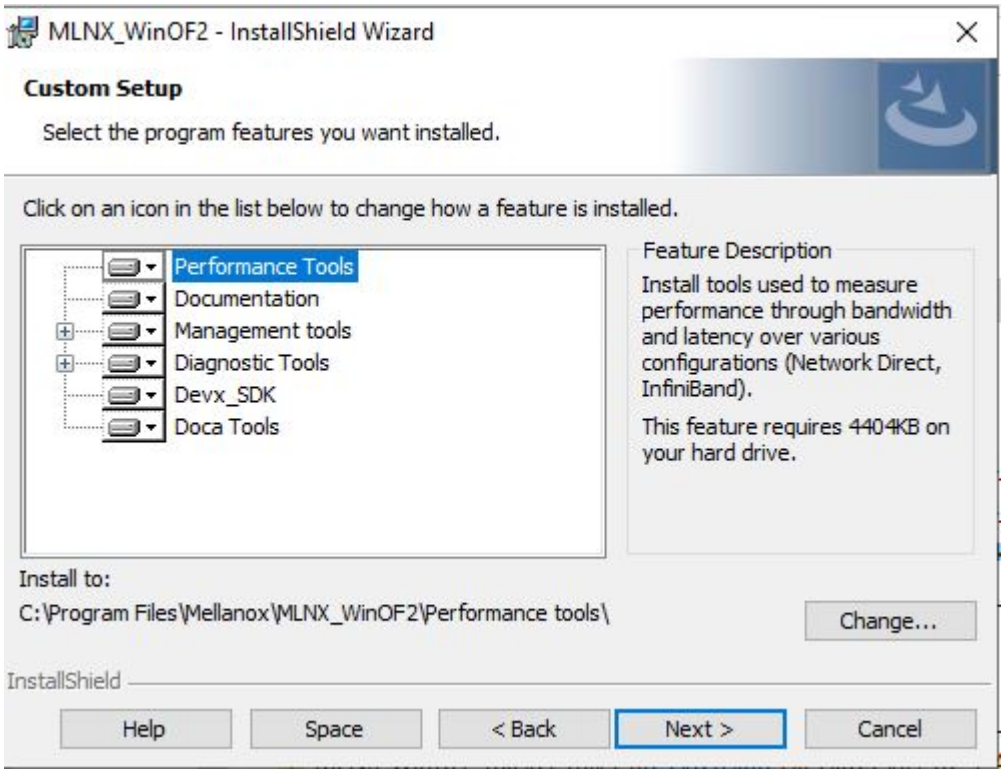
10. Select a Complete or Custom installation, follow Step a onward.



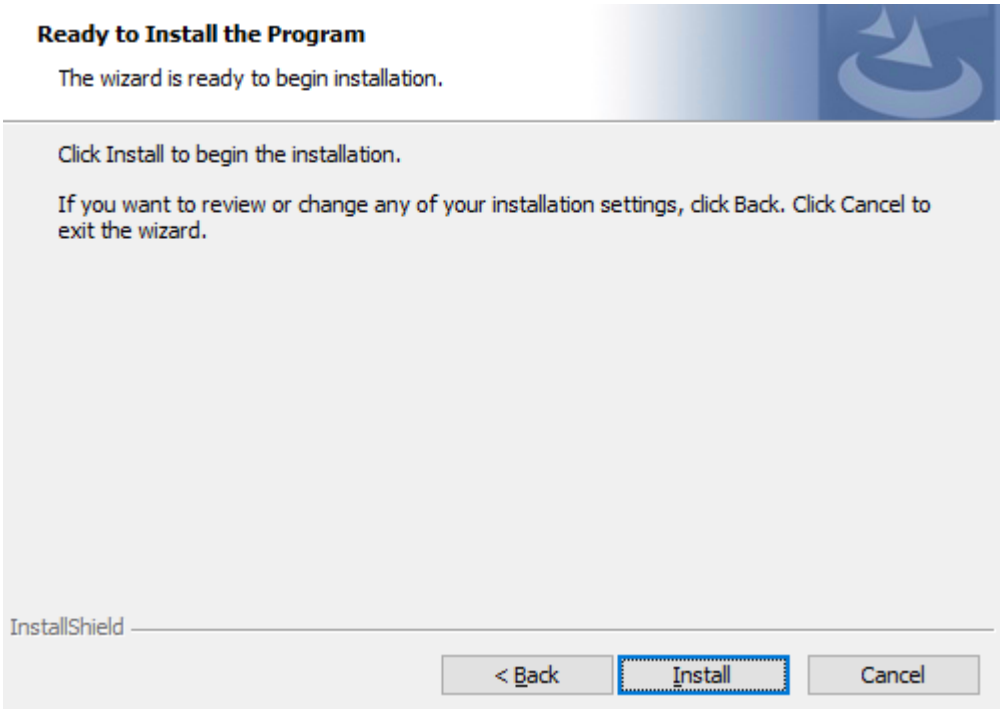
1. Select the desired feature to install:

   - 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

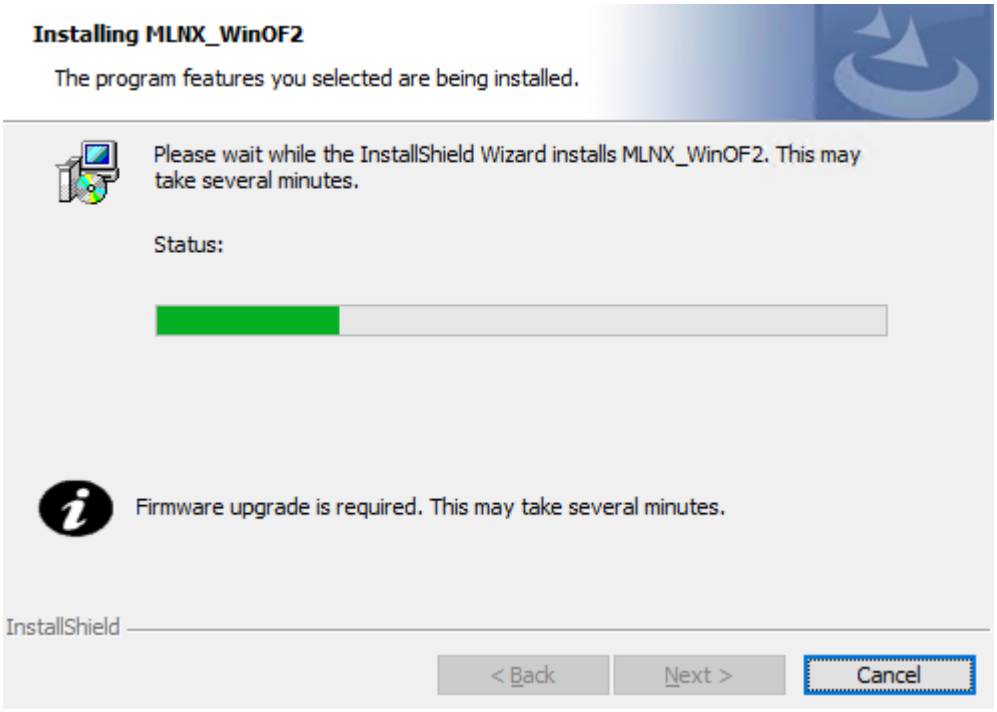
2. Click Next to install the desired tools.



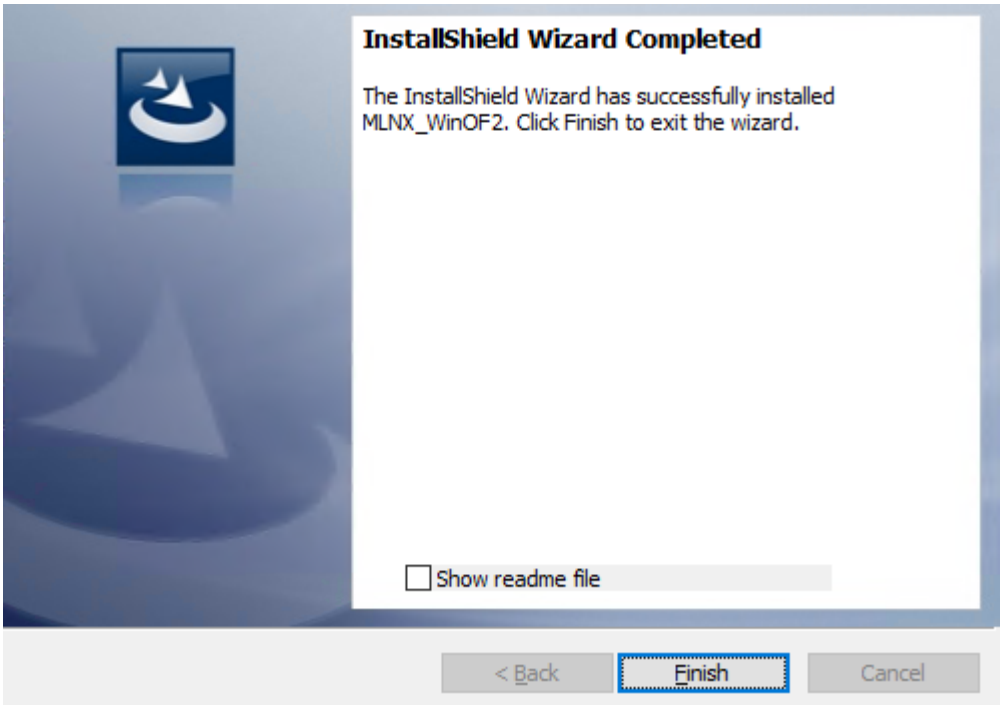
11. Click Install to start the installation.



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



13. Click Finish to complete the installation.



Unattended Installation

> ⚠ **Warning**
>
> 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"
```

> ⚠ **Warning**
>
> 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.

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

Updating Adapter Firmware

Each 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 inmlxup - Update and Query Utility .

Below is an example of a firmware update using the mlxup utility.

```
[server1]# ./mlxup
Querying Mellanox devices firmware ...
Device Type:     ConnectX-4
Part Number:     MCX453A-FCAT
Description:     ConnectX®-4 VPI adapter card, FDR IB 40/56GbE, 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                       16.00.0000     16.00.0000

Status:           Up to date



Device Type:     ConnectX-4
Part Number:     MCX453A-FCAT
Description:     ConnectX®-4 VPI adapter card, FDR IB 40/56GbE, 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               16.00.0000  16.00.1000
Status:          Update required

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
```

Specifications

MCX453A-FACT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile				
	Connector: Single QSFP28 InfiniBand (copper and optical)				
Protocol Support	InfiniBand[(a)]: IBTA v1.3 Auto-Negotiationa : 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane) port				
	Ethernet: 56GBASE-R4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 1000BASE-CX, 1000BASEKX, 10GBASE-SR, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII				
	Data Rate	InfiniBand	SDR/DDR/QDR/FDR		
		Ethernet	1/10/25/40/56 Gb/s		
	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**		**56GbE**	
	Typical Power[(b)]	Passive Cables		10.0W	
	Maximum Power	Passive Cables		12.5W	
		1.5W Active Cables		14.2W	
	Maximum power available through QSFP28 port: 3.5W				
	Airflow (LFM)	**Heatsink to Port**		**Port to Heatsink**	
		Passive Cables	Active 1.5W Cables	Passive Cables	Active 1.5W Cables
		400LFM	450LFM	400LFM	400LFM

Environmental	**Temperature**	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Regulatory	Safety	CB / cTUVus / CE	
	EMC	CE / FCC / VCCI / ICES / RCM	
	RoHS	RoHS compliant	

Notes: a. The ConnectX-4 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.
b. Typical power for ATIS traffic load.
c. For both operational and non-operational states.

MCX454A-FCAT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile		
	Connector: Single QSFP28 InfiniBand (copper and optical)		
Protocol Support	InfiniBand[(a)]: IBTA v1.3 Auto-Negotiationa : 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane) port		
	Ethernet: 56GBASE-R4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 1000BASE-CX, 1000BASEKX, 10GBASE-SR, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII		
	Data Rate:	InfiniBand	SDR/DDR/QDR/FDR

		Ethernet	1/10/25/40/56 Gb/s		
	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**		**FDR/56GbE**	
	Typical Power[(b)]	Passive Cables		12.4W	
	Maximum Power	Passive Cables		14.9W	
		1.5W Active Cables		18.3W	
	Maximum power available through QSFP28 port: 3.5W				
	Airflow (LFM)	**Heatsink to Port**		**Port to Heatsink**	
		Passive Cables	Active 1.5W Cables	Passive Cables	Active 1.5W Cables
		400LFM	450LFM	400LFM	400LFM
Environmental	**Temperature**	Operational	0°C to 55°C		
		Non-operational	-40°C to 70°C		
	Humidity	Operational	10% to 85% relative humidity		
		Non-operational	10% to 90% relative humidity		
	Altitude (Operational)	3050m			
Regulatory	Safety	CB / cTUVus / CE			
	EMC	CE / FCC / VCCI / ICES / RCM			
	RoHS	RoHS compliant			

Notes: a. The ConnectX-4 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.
b. Typical power for ATIS traffic load.
c. For both operational and non-operational states.

MCX455A-FCAT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile				
	Connector: Single QSFP28 InfiniBand (copper and optical)				
Protocol Support	InfiniBand[(a)]: IBTA v1.3 Auto-Negotiationa : 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane) port				
	Ethernet: 56GBASE-R4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 1000BASE-CX, 1000BASEKX, 10GBASE-SR, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII				
	Data Rate:	InfiniBand	SDR/DDR/QDR/FDR		
		Ethernet	1/10/25/40/56 Gb/s		
	PCI Express Gen3: SERDES @ 8.0GT/s, 16 lanes (2.0 and 1.1 compatible)				
Power and Airflow	Voltage: 12V, 3.3V				
	Power	**Cable Type**		**FDR/56GbE**	
	Typical Power[(b)]	Passive Cables		13.51W	
	Maximum Power	Passive Cables		16W	
		1.5W Active Cables		17.7W	
	Maximum power available through QSFP28 port: 3.5W				
	Airflow (LFM)	**Heatsink to Port**		**Port to Heatsink**	
		Passive Cables	Active 1.5W Cables	Passive Cables	Active 1.5W Cables
		450LFM	400LFM	400LFM	400LFM
Environmental	**Temperature**	Operational	0°C to 55°C		
		Non-operational	-40°C to 70°C		
	Humidity	Operational	10% to 85% relative humidity		

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

Notes: a. The ConnectX-4 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.
b. Typical power for ATIS traffic load.
c. For both operational and non-operational states.

MCX456A-FCAT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile		
	Connector: Single QSFP28 InfiniBand (copper and optical)		
Protocol Support	InfiniBand[(a)]: IBTA v1.3 Auto-Negotiationa : 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane) port		
	Ethernet: 56GBASE-R4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 1000BASE-CX, 1000BASEKX, 10GBASE-SR, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII		
	Data Rate:	InfiniBand	SDR/DDR/QDR/FDR
		Ethernet	1/10/25/40/56 Gb/s
	PCI Express Gen3: SERDES @ 8.0GT/s, 16 lanes (2.0 and 1.1 compatible)		
Power and	Voltage: 12V, 3.3V		
	Power	**Cable Type**	**FDR/56GbE**

Airflow	Typical Power[b]	Passive Cables		15.9W	
	Maximum Power	Passive Cables		18.5W	
		1.5W Active Cables		21.8W	
	Maximum power available through QSFP28 port: 3.5W				
	Temperature	Operational		0°C to 55°C	
		Non-operational		-40°C to 70°C	
	Humidity	90% relative humidity[c]			
	Altitude (Operational)	3050m			
	Airflow (LFM)	**Heatsink to Port**		**Port to Heatsink**	
		Passive Cables	Active 1.5W Cables	Passive Cables	Active 1.5W Cables
		400LFM	500LFM	400LFM	400LFM
Environmental	**Temperature**	Operational	0°C to 55°C		
		Non-operational	-40°C to 70°C		
	Humidity	Operational	10% to 85% relative humidity		
		Non-operational	10% to 90% relative humidity		
	Altitude (Operational)	3050m			
Regulatory	Safety	CB / cTUVus / CE			
	EMC	CE / FCC / VCCI / ICES / RCM			
	RoHS	RoHS compliant			

Notes: a. The ConnectX-4 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.

b. Typical power for ATIS traffic load.
c. For both operational and non-operational states.

Notes: a. The ConnectX-4 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.
b. Typical power for ATIS traffic load.
c. For both operational and non-operational states.

MCX455A-ECAT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile		
	Connector: Single QSFP28 InfiniBand (copper and optical)		
Protocol Support	InfiniBand[(a)]: IBTA v1.3 Auto-Negotiationa : 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane) port		
	Ethernet: 56GBASE-R4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 1000BASE-CX, 1000BASEKX, 10GBASE-SR, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII		
	Data Rate:	InfiniBand	SDR/DDR/QDR/FDR/EDR
		Ethernet	1/10/25/40/50/100 Gb/s
	PCI Express Gen3: SERDES @ 8.0GT/s, 16 lanes (2.0 and 1.1 compatible)		
Power and Airflow	Voltage: 12V, 3.3V		
	Power	**Cable Type**	**FDR/56GbE**
	Typical Power[(b)]	Passive Cables	14W
	Maximum Power	Passive Cables	16.6W
		1.5W Active Cables	18.3W
		2.5W Active Cables	19.4W
	Maximum power available through QSFP28 port: 3.5W		
	Airflow (LFM)	**Heatsink to Port**	**Port to Heatsink**

		Passive Cables	Active 3.5W Cables	Passive Cables	Active 3.5W Cables
		400LFM	700LFM	400LFM	600LFM
Environmental	**Temperature**	Operational	0°C to 55°C		
		Non-operational	-40°C to 70°C		
	Humidity	Operational	10% to 85% relative humidity		
		Non-operational	10% to 90% relative humidity		
	Altitude (Operational)	3050m			
Regulatory	Safety	CB / cTUVus / CE			
	EMC	CE / FCC / VCCI / ICES / RCM			
	RoHS	RoHS compliant			

Notes: a. The ConnectX-4 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.
b. Typical power for ATIS traffic load.
c. For both operational and non-operational states.

MCX456A-ECAT Specifications

Physical	Size: 2.71 in. x 5.6 in. (68.90mm x 142.24 mm) – low profile
	Connector: Dual QSFP28 InfiniBand (copper and optical)
Protocol Support	InfiniBand[(a)]: IBTA v1.3 Auto-Negotiationa : 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane) port
	Ethernet: 56GBASE-R4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-

	KR2, 1000BASE-CX, 1000BASEKX, 10GBASE-SR, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII				
	Data Rate:	InfiniBand	SDR/DDR/QDR/FDR/EDR		
		Ethernet	1/10/25/40/50/100 Gb/s		
	PCI Express Gen3: SERDES @ 8.0GT/s, 16 lanes (2.0 and 1.1 compatible)				
Power and Airflow	Voltage: 12V, 3.3V				
	Power	**Cable Type**		**FDR/56GbE**	
	Typical Power[(b)]	Passive Cables		16.3W	
	Maximum Power	Passive Cables		19W	
		1.5W Active Cables		22.3W	
		2.5W Active Cables		24.5W	
	Maximum power available through QSFP28 port: 3.5W				
	Airflow (LFM)	**Heatsink to Port**		**Port to Heatsink**	
		Passive Cables	Active 3.5W Cables	Passive Cables	Active 3.5W Cables
		400LFM	800LFM	400LFM	600LFM
Environmental	**Temperature**	Operational	0°C to 55°C		
		Non-operational	-40°C to 70°C		
	Humidity	Operational	10% to 85% relative humidity		
		Non-operational	10% to 90% relative humidity		
	Altitude (Operational)	3050m			
Regulatory	Safety	CB / cTUVus / CE			
	EMC	CE / FCC / VCCI / ICES / RCM			
	RoHS	RoHS compliant			

Notes: a. The ConnectX-4 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.
b. Typical power for ATIS traffic load.
c. For both operational and non-operational states.

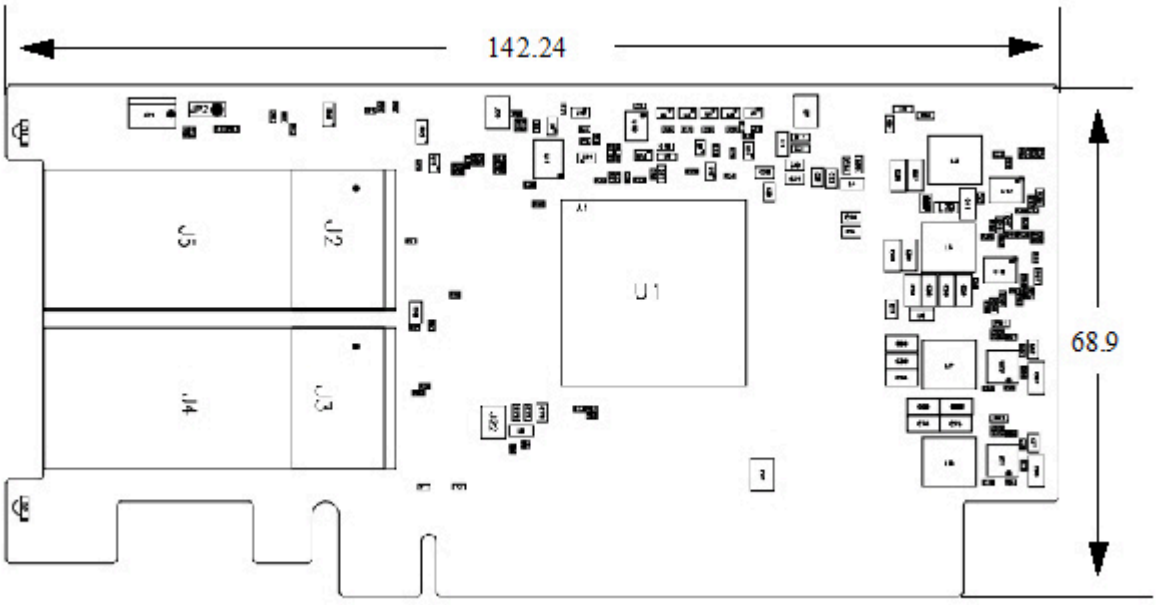
LED Specifications

See LED Interface.

Board Mechanical Drawings and Dimensions

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

Adapter Card Mechanical Drawing (Component Side)

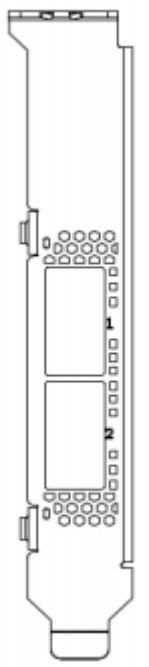


e

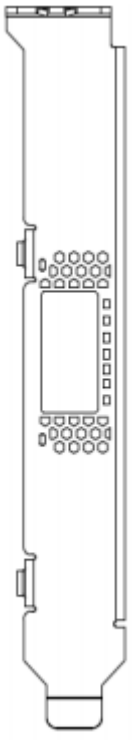
Bracket Mechanical Drawing

> ⚠ **Warning**
>
> All dimensions are in millimeters. The bracket mechanical tolerance is +/- 0.25mm.

Dual-port Bracket



Single-port Bracket



Document Revision History

Date	Comments/Changes
Sep. 2022	Added a note on FRU EEPROM under the Features and Benefits table.
Feb. 2020	Added altitude criteria to "Specifications".
Oct. 2019	Updated the supported Ethernet protocols in "Specifications".
Jun. 2019	Migrated to online format; minor reorganization
Jul. 2018	Updated all specifications tables
Oct. 2017	Updated the following "Specifications Tables"
Jun. 2017	• Updated FDR adapters product overview • Updated "MCX456A-ECAT Specifications"
Dec. 2016	Added "Windows Driver"
Sep. 2016	• Updated "Adapter Card Installation Instructions" • Added "Adapter Card Un-installation Instructions" • Updated "Troubleshooting"
Aug. 2016	• Updated "Adapter Card Installation Instructions" • Updated "Cable Installation" • Updated "Airflow Specifications" • Added mechanical drawings • Updated "Safety Warnings"
Apr. 2016	• Updated "Cable Installation" • Updated "Airflow Specifications"
Mar. 2016	Removed I2C-Compatible Interface

Feb. 2016	• Updated speed specifications in the following sections:
	  ◦ 'Product Overview"
	  ◦ "Features and Benefits"

Date	Comments/Changes
Feb. 2016	Added 56GbE support to relevant sections.
Dec. 2015	Added port numbers to the specification tables
Dec. 2015	Updated specifications tables
Nov.2015	Updated specifications tables
Jun. 2015	First Release

© Copyright 2024, NVIDIA. PDF Generated on 08/26/2024