NVIDIA ConnectX-6 Lx Ethernet Adapter Cards for OCP Spec 3.0 User Manual

Exported on Dec/12/2022 02:32 PM

Table of Contents

Introduction 6
Products Overview 6
ConnectX-6 Lx 25GbE Adapter Cards for OCP 3.0 6
ConnectX-6 Lx 50GbE Adapter Cards for OCP 3.0 7
Features and Benefits 7
Operating Systems/Distributions 9
Connectivity 9
Interfaces 11
ConnectX-6 Lx IC 12
Encryption 12
Ethernet SFP28/QSFP28 Interfaces 12
PCI Express Interface 12
Networking Ports LEDs Interface 12
FRU EEPROM 13
SMBus Interface 13
Voltage Regulators 14
CPLD Interface 14
Heatsink 14
Hardware Installation 15
Safety Warnings 15
Installation Procedure Overview 15
System Requirements 15
Hardware Requirements 16
Airflow Requirements 16
Software Requirements 16
Safety Precautions 16
Pre-Installation Checklist 17
OCP 3.0 Bracket Replacement Instructions 17
OCP 3.0 Adapter Card Installation Instructions 17
Cables and Modules 17
Identifying the Card in Your System 18
Extraction Instructions 19

Driver Installation 20
Windows Driver Installation 20
Software Requirements 20
Downloading WinOF-2 Driver 20
Installing WinOF-2 Driver 21
Firmware Upgrade 28
Linux Driver Installation 28
Prerequisites 28
VMware Driver Installation 41
Hardware and Software Requirements 41
Installing NATIVE ESXi Driver for VMware vSphere 41
Removing Earlier NVIDIA Drivers 42
Firmware Programming 42
Updating Adapter Firmware 43
Troubleshooting 44
General Troubleshooting 44
Linux Troubleshooting 45
Windows Troubleshooting 45
Specifications 47
MCX631432AC-ADAB / MCX631432AE-ADAB / MCX631432AN-ADAB / MCX631432AS-ADAB Specifications 47
MCX631435AC-GDAB / MCX631435AE-GDAB / MCX631435AN-GDAB Specifications 48
Board Mechanical Drawing and Dimensions 49
Bracket Mechanical Drawings and Dimensions 49
Monitoring 51
Thermal Sensors 51
Heatsink 51
Finding the MAC and Serial Number on the Adapter Card 52
Document Revision History 53

About This Manual

This is the User Guide for Ethernet adapter cards based on the NVIDIA® ConnectX®-6 Lx integrated circuit device for Open Compute Project Spec 3.0. These adapters connectivity provide the highest performing low latency and most flexible interconnect solution for servers supporting OCP spec 3.0 adapter cards used in Enterprise Data Centers and High-Performance Computing environments.

Important Note:

In some of the OCP 3.0 prototype samples, the insertion force that is required to install the card into a particular PCI connector may exceed the maximum insertion force that is allowed by the connector's spec. NVIDIA is updating the OCP 3.0 cards to ensure the proper insertion process.

Ordering Part Numbers

The table below provides the ordering part numbers (OPN) for the available ConnectX-6 Lx Ethernet adapter cards for OCP Spec 3.0.

NVIDIA SKU	Legacy OPN	Marketing Description
900-9X625-0073-SB1	MCX631432AS-ADAB	ConnectX®-6 Lx EN adapter card, **25GbE** OCP3.0, With **Host management**, Dual-port **SFP28**, PCIe 4.0 **x8**, **Secure Boot**, No Crypto, Thumbscrew (Pull Tab) Bracket
900-9X625-0053-SB0	MCX631432AN-ADAB	ConnectX®-6 Lx EN adapter card, **25GbE** OCP3.0, with **Host Management**, **Dual**-port **SFP28**, PCIe 4.0 **x8**, No Crypto, Thumbscrew (Pull Tab) Bracket
900-9X625-0083-SB0	MCX631432AC-ADAB	ConnectX®-6 Lx EN adapter card, **25GbE** OCP3.0, With **Host Management**, **Dual**-port **SFP28**, PCIe 4.0 **x8**, **Crypto Enabled** and **Secure Boot**, Thumbscrew (Pull Tab) Bracket
900-9X659-0015-SB1	MCX631435AN-GDAB	ConnectX®-6 Lx EN adapter card, **50GbE**, OCP3.0, with **Host Management**, **Single**-port **QSFP28**, PCIe 4.0 **x8**, No Crypto, Thumbscrew (Pull Tab) Bracket
900-9X659-0045-SB0	MCX631435AC-GDAB	ConnectX®-6 Lx EN adapter card, **50GbE**, OCP3.0, With **Host Management**, **Single**-port **QSFP28**, PCIe 4.0 **x8**, **Crypto Enabled** and **Secure Boot**, Thumbscrew (Pull Tab) Bracket

Legacy (EOL) Ordering Part Numbers

NVIDIA SKU	Legacy OPN	Marketing Description
900-9X625-0063-SB0	MCX631432AE-ADAB	ConnectX®-6 Lx EN adapter card, **25GbE** OCP3.0, with **Host Management**, **Dual**-port **SFP28**, PCIe 4.0 **x8**, **Crypto Enabled**, No Secure Boot, Thumbscrew (Pull Tab) Bracket
900-9X625-0073-SI0	MCX631432AS-ADAI	ConnectX-6 Lx EN adapter card, 25GbE OCP3.0, With Host management , Dual-port SFP28, PCIe 4.0 x8, Secure Boot, No Crypto, Internal Lock Bracket
900-9X659-0025-SB0	MCX631435AE-GDAB	ConnectX®-6 Lx EN adapter card, **50GbE**, OCP3.0, with **Host Management**, **Single**-port **QSFP28**, PCIe 4.0 **x8**, **Crypto Enabled**, No Secure Boot, Thumbscrew (Pull Tab) 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 MLNX_OFED for Linux User Manual and Release Notes	User Manual and Release Notes describing MLNx_OFED features, performance, band diagnostic, tools content and configuration. See NVIDIA 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 and Release Notes	User Manual describing the various components of the NVIDIA ConnectX® NATIVE ESXi stack. See VMware® ESXi Documentation.
NVIDIA Firmware Update	NVIDIA firmware update and query utility used to update the firmware. See NVIDIA 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.
IEEE Std 802.3 Specification	IEEE Ethernet Specifications
PCI Express Specifications	Industry Standard PCI Express Base and Card Electromechanical Specifications.
LinkX Interconnect Solutions	LinkX Ethernet cables and transceivers are designed to maximize the performance of High-Performance Computing networks, requiring high-bandwidth, low-latency connections between compute nodes and switch nodes. NVIDIA offers one of the industry's broadest portfolio of 40GbE, 56GbE, 100GbE, 200GbE and 400GbE cables, including Direct Attach Copper cables (DACs), copper splitter cables, Active Optical Cables (AOCs) and transceivers in a wide range of lengths from 0.5m to 10km. In addition to meeting Ethernet standards, NVIDIA tests every product in an end-to-end environment ensuring a Bit Error Rate of less than 1E-15. Read more at LinkX Cables and Transceivers.
Open Compute Project 3.0 Specification	https://www.opencompute.org/

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

This is the *User Guide* for Ethernet adapter cards based on the ConnectX®-6 Lx integrated circuit device for OCP Spec 3.0. These adapters connectivity provide the highest performing and most flexible interconnect solution for servers supporting OCP 3.0 used in Enterprise Data Centers and High-Performance Computing environments.

Providing up to two ports of 25GbE or a single-port of 50GbE connectivity, and PCIe Gen 3.0/4.0 x8 host connectivity, ConnectX-6 Lx is a member of NVIDIA's world-class, award-winning, ConnectX family of network adapters. Continuing NVIDIA's consistent innovation in networking, ConnectX-6 Lx provides agility and efficiency at every scale. ConnectX-6 Lx delivers cutting edge 25GbE performance and security for uncompromising data centers. ConnectX-6 Lx offers IPsec inline encryption/decryption acceleration (in Crypto enabled products) and connection-tracking hardware offload to accelerates L4 firewall performance as part of NVIDIA ASAP^2.

ConnectX-6 Lx also delivers supply chain protection with hardware Root-of-Trust (RoT) for Secure Boot (in Secure Boot enabled products) as well as Secure Firmware Update using RSA cryptography and cloning protection, via a device-unique key, to guarantee firmware authenticity

ConnectX-6 Lx SmartNICs are available in several form factors including low-profile PCIe and OCP 3.0 cards with SFP28 connectors for 10/25GbE applications, or QSFP28 for 50GbE applications. OCP3.0 cards are available with either a pull tab or an internal lock bracket. For the low-profile PCIe cards hardware user manual, please refer to ConnectX-6 Lx HHHL Cards UM.

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

> ⚠ Please note the following OPNs include a thumbscrew (pull tab) bracket. For other retention mechanisms brackets, please contact NVIDIA.

Products Overview

ConnectX-6 Lx 25GbE Adapter Cards for OCP 3.0

OPN	MCX631432AN-ADAB	MCX631432AE-ADAB	MCX631432AC-ADAB	MCX631432AS-ADAB
Network Connector Type	Dual-port SFP28	Dual-port SFP28	Dual-port SFP28	Dual-port SFP28
Ethernet Data Rate	10/25 Gb/s			
PCI Express Connectors	PCIe Gen 4.0 x8 SERDES @ 16.0GT/s			
Dimensions	Small Form Factor (SFF) OCP 3.0 4.52 in. x 2.99 in (115.00mm x 76.00mm)			
Retention Mechanism	**Thumbscrew** (Pull Tab) Bracket			
RoHS	RoHS Compliant	RoHS Compliant	RoHS Compliant	RoHS Compliant
IC Part Number	MT28942A0-YCCF-AE	MT28942A0-YCCF-AE	MT28942A0-YCCF-AE	MT28942A0-YCCF-AE
Host Management	✓	✓	✓	✓
Crypto	-	✓	✓	-

OPN	MCX631432AN-ADAB	MCX631432AE-ADAB	MCX631432AC-ADAB	MCX631432AS-ADAB
Secure Boot	-	-	✓	✓
Device ID	4127 for Physical Function (PF) 4128 for Virtual Function (VF)			

ConnectX-6 Lx 50GbE Adapter Cards for OCP 3.0

OPN	MCX631435AN-GDAB	MCX631435AE-GDAB	MCX631435AC-GDAB
Network Connector Type	Single-port QSFP28	Single-port QSFP28	Single-port QSFP28
Ethernet Data Rate	10/25/50 Gb/s		
PCI Express Connectors	PCIe Gen 4.0 x8 SERDES @ 16.0GT/s		
Dimensions	Small Form Factor (SFF) OCP 3.0 4.52 in. x 2.99 in (115.00mm x 76.00mm)		
Retention Mechanism	Thumbscrew (Pull Tab) Bracket		
RoHS	RoHS Compliant	RoHS Compliant	RoHS Compliant
IC Part Number	MT28942A0-YCCF-AE	MT28942A0-YCCF-AE	MT28942A0-YCCF-AE
Host Management	✓	✓	✓
Crypto	-	✓	✓
Secure Boot	-	-	✓
Device ID	4127 for Physical Function (PF) 4128 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 4.0 (16GT/s) through x8 edge connector.

Feature	Description
Up to 200 Gigabit Ethernet	NVIDIA adapters comply with the following IEEE 802.3 standards: • 200GbE / 100GbE / 50GbE / 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 802.3by 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 Components	• EEPROM - The EEPROM capacity is 32Kbit. FRU I^2C address is (0x50) and is accessible through the PCIe SMBus (Note: Address 0x58 is reserved.) • SPI Quad - includes 256Mbit SPI Quad Flash device (MX25L25645GXDI-08G device by Macronix)
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®-6 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®-6 Lx, utilizing 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®-6 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®-6 Lx advanced acceleration technology enables higher cluster efficiency and scalability to tens of thousands of nodes.
CPU Offload	Adapter functionality enables reduced CPU overhead leaving more CPU resources available for computation tasks. Open vSwitch (OVS) offload using $ASAP^{2(TM)}$ • Flexible match-action flow tables • Tunneling encapsulation/decapsulation
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®-6 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

Feature	Description
SR-IOV	ConnectX®-6 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
Host Management	NVIDIA host management sideband implementations enable remote monitor and control capabilities using RBT, MCTP over SMBus, and MCTP over PCIe – Baseboard Management Controller (BMC) interface, supporting both NC-SI and PLDM management protocols using these interfaces. NVIDIA OCP 3.0 adapters support these protocols to offer numerous Host Management features such as PLDM for Firmware Update, network boot in UEFI driver, UEFI secure boot, and more.
Secure Boot	Hardware Root-of-Trust (RoT) Secure Boot and secure firmware update using RSA cryptography, and cloning-protection, via a device-unique secret key.
Crypto	Crypto – IPsec and TLS data-in-motion inline encryption and decryption offload and AES-XTS block-level data-at-rest encryption and decryption offload.
Wake-on-LAN (WoL)	The adapter supported Wake-on-LAN (WoL), a computer networking standard that allows an adapter to be turned on or awakened by a network message. In STBY mode, only port0 is available.
Reset-on-LAN (RoL)	Supported
Connection tracking	ConnectX-6 Lx ASAP² connection-tracking hardware offload accelerates L4 firewall performance.
IPSec Offload (in selected OPNs)	ConnectX-6 Lx adapters offer a range of advanced built-in capabilities that bring infrastructure security down to every endpoint with unprecedented performance and scalability. ConnectX-6 Lx offers IPsec inline encryption/decryption acceleration to offload the CPU and run the entire data path encryption or decryption in the NIC hardware while maintaining the line rate.
Secure Boot (Hardware Root Of Trust)	ConnectX-6 Lx delivers supply chain protection with hardware Root-of-Trust (RoT) for Secure Boot as well as Secure Firmware Update using RSA cryptography and cloning protection, via a device-unique key, to guarantee firmware authenticity

Operating Systems/Distributions

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

Connectivity

- Interoperable with 1/10/25/50 Gb/s Ethernet switches
- Passive copper cable with ESD protection

- Powered connectors for optical and active cable support

Interfaces

The below figures show the component side of the NVIDIA ConnectX-6 Lx adapter card. Each numbered interface that is referenced in the figures is described in the following table with a link to detailed information.

> ⚠ The ConnectX-6 Lx adapter card includes special circuits to protect from ESD shocks to the card/server when plugging copper cables

> ⚠ The below figures are for illustration purposes only and might not reflect the current revision of the adapter card.

NVIDIA ConenctX-6 Lx for OCP 3.0 Adapter Card Interfaces

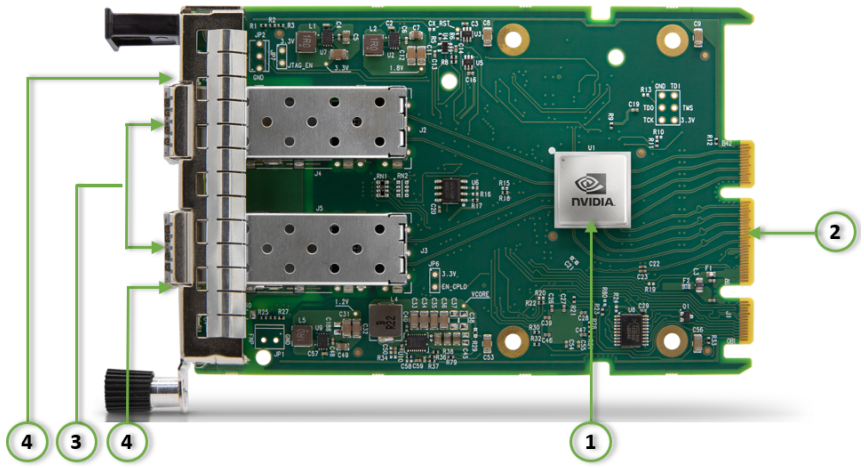

Call out	Item	Description
1	“ConnectX-6 Lx IC”	NVIDIA ConnectX-6 Lx IC on the board.
2	“PCI Express Interface"	PCIe Gen 4.0 through an x8 edge connector.
3	“Ethernet SFP28/ QSFP28 Interface”	Ethernet traffic is transmitted through the adapter's SFP28/QSFP28 connectors. The networking connectors allow for the use of modules, optical and passive cable interconnect solutions.
4	“Networking Ports LEDs Interface”	There are two I/O LEDs, LED0 and LED1, per port to indicate speed and link status.
	“FRU EEPROM”	FRU EEPROM capacity 4Kb
	“SMBus Interface”	Allows BMC connectivity using MCTP over SMBus or MCTP over PCIe protocols.
	“Voltage Regulators”	Voltage supply pins that feed onboard regulators.
	“CPLD Interface”	Controls the networking port LEDs and implements the OCP 3.0 host scan chain.

ConnectX-6 Lx IC

The ConnectX®-6 Lx EN family of adapter IC devices delivers two ports of 10/25Gb/s or a single-port of 50Gb/s Ethernet connectivity paired with best-in-class hardware capabilities that accelerate and secure cloud and data-center workloads.

Encryption

⚠ Applies to Crypto-enabled OPNs.

ConnectX-6 Lx brings security to every end-point, including:

- Purpose-built inline acceleration engines that offload IPsec and TLS data-in-motion and XTS-AES data-at-rest cryptographic operations.
- Stateful firewall solution acceleration, powered by Open vSwitch connection tracking and NVIDIA's ASAP2 technology.
- Embedded hardware root-of-trust and support for RSA-based secure firmware update and secure boot, providing guaranteed integrity of the network adapter.

Ethernet SFP28/QSFP28 Interfaces

The network ports of the ConnectX®-6 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®-6 Lx adapter cards.

Supported PCIe Interface	Features
PCIe Gen 4.0 (1.1 and 2.0 compatible) through x8 edge connectors	Link Rates: 2.5. 5.0, 8.0 or 16GT/s. Auto Negotiation to: x16, x8, x4, x2 or x1. Support for MSI/MSI-X mechanisms.

Networking Ports LEDs Interface

There are two I/O LEDs, LED0 and LED1, per port to indicate speed and link status. LED0 is a bicolor (yellow and green) LED and LED1 is a single color (green) LED.

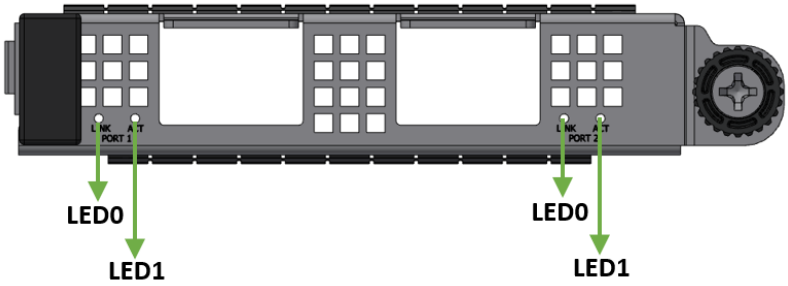

Link Indications

LED and State	Description
1Hz blinking Yellow	Beacon command for locating the adapter card
4Hz blinking Yellow	Indicates an error with the link. The error can be one of the following: Error Type / Description / LED Behavior I^2C / I^2C access to the networking ports fails / Blinks until the error is fixed Over-current / Over-current condition of the networking ports / Blinks until the error is fixed
LED0 - Link Speed	• A constant Green indicates a link with the maximum networking speed. • A constant Yellow indicates a link with less than the maximum networking speed • If LED0 is off, then the link has not been established.
LED1 - Activity	• A blinking Green indicates a valid link with data transfer. • If LED1 is off, then there is no activity

FRU EEPROM

FRU EEPROM allows the baseboard to identify different types of Mezzanine cards. FRU EEPROM is accessible through SMCLK and SMDATA. FRU EEPROM address is defined according to SLOT_ID0 and SLOT_ID1 and its capacity is 4Kb.

SMBus Interface

ConnectX®-6 Lx technology maintains support for manageability through a BMC. ConnectX®-6 Lx OCP 3.0 adapter can be connected to a BMC using MCTP over SMBus or MCTP over PCIe protocols as if it is a standard NVIDIA OCP 3.0 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 OCP 3.0 edge connector 12V and 3.3V supply pins. These voltage supply pins feed onboard regulators that provide the necessary power to the various components on the card.

CPLD Interface

The adapter card incorporates a CPLD device that controls the networking port LEDs, as well as an implementation of the OCP 3.0 host scan chain. It draws its power supply from 3.3V_EDGE.

Heatsink

A heatsink is attached to the ConnectX®-6 Lx IC in order to dissipate the heat. It is attached either by using two spring-loaded push pins that insert into four mounting holes.
ConnectX®-6 Lx IC has a thermal shutdown safety mechanism that automatically shuts down the ConnectX®-6 Lx card in case of a high-temperature event, improper thermal coupling or heatsink removal.

Hardware Installation

Installation and initialization of ConnectX-6 Lx adapter cards for OCP Spec 3.0 require attention to the mechanical attributes, power specification, and precautions for electronic equipment.

Safety Warnings

Safety warnings are provided here in the English language. For safety warnings in other languages, refer to the Adapter Installation Safety Instructions document available on nvidia.com.

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

Installation Procedure Overview

The installation procedure 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	Follow the safety precautions	Refer to Safety Precautions
4	Follow the pre-installation checklist	Refer to Pre-Installation Checklist
5	(Optional) Replace the assembled OCP 3.0 bracket with the desired form factor bracket	Refer to OCP 3.0 Bracket Replacement Instructions
6	Install ConnectX-6 Lx adapter card for OCP spec 3.0 in the system	Refer to Cards with Internal Lock Bracket Installation
		Refer to Cards with Internal Lock Bracket Installation
7	Connect cables or modules to the card	Refer to Cables and Modules
8	Identify ConnectX-6 Lx adapter card in the system	Refer to Identify the Card in Your System

System Requirements

Unless otherwise specified, NVIDIA products are designed to work in an environmentally controlled data center with low levels of gaseous and dust (particulate) contamination. The operating environment should meet severity level G1 as per ISA 71.04 for gaseous contamination and ISO 14644-1 class 8 for cleanliness level.

Hardware Requirements

> 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.

A system with a PCI Express x16 slot for OCP spec 3.0 is required for installing the card.

Airflow Requirements

ConnectX-6 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.

Hot Aisle Cooling Heatsink-to-port Airflow Direction	Cold Aisle Cooling Port-to-heatsink Airflow Direction
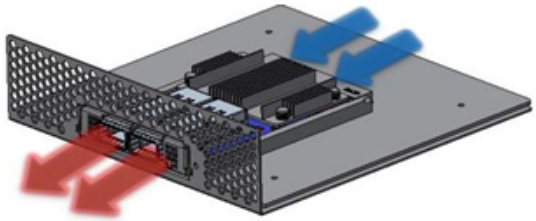	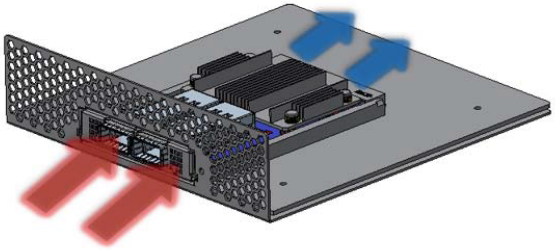

Software Requirements

- See Operating Systems/Distributions section under the Introduction section.
- Software Stacks - NVIDIA OpenFabric software package MLNX_OFED for Linux, WinOF-2 for Windows, and VMware. See the Driver Installation section.

Safety Precautions

> The adapter is being installed in a system that operates with voltages that can be lethal. Before opening the case of the system, observe the following precautions to avoid injury and prevent damage to system components.

1. Remove any metallic objects from your hands and wrists.
2. Make sure to use only insulated tools.

3. Verify that the system is powered off and is unplugged.
4. It is strongly recommended to use an ESD strap or other antistatic devices.

Pre-Installation Checklist

1. Unpack the ConnectX-6 Lx adapter card.
   Unpack and remove the ConnectX-6 Lx card. Check the parts for visible damage that may have occurred during shipping. Please note that the cards must be placed on an antistatic surface.

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

2. Shut down your system if active:
   Turn off the power to the system, and disconnect the power cord. Refer to the system documentation for instructions. Before you install the ConnectX-6 Lx card, make sure that the system is disconnected from power.

OCP 3.0 Bracket Replacement Instructions

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

OCP 3.0 Adapter Card Installation Instructions

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

Cables and Modules

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

1. All cables can be inserted or removed with the unit powered on.
2. To insert a cable, press the connector into the port receptacle until the connector is firmly seated.
   a. Support the weight of the cable before connecting the cable to the adapter card. Do this by using a cable holder or tying the cable to the rack.
   b. Determine the correct orientation of the connector to the card before inserting the connector. Do not try and insert the connector upside down. This may damage the adapter card.
   c. Insert the connector into the adapter card. Be careful to insert the connector straight into the cage. Do not apply any torque, up or down, to the connector cage in the adapter card.
   d. Make sure that the connector locks in place.

When installing cables make sure that the latches engage.

Always install and remove cables by pushing or pulling the cable and connector in a straight line with the card.

3. After inserting a cable into a port, the Yellow or Green LED0 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 Adapter Card LED Operations.
4. After plugging in a cable, lock the connector using the latching mechanism particular to the cable vendor. When a logical connection is made, Green LED1 will light. When data is being transferred, Green LED1 will blink.
5. Care should be taken as not to impede the air exhaust flow through the ventilation holes. Use cable lengths which allow for routing horizontally around to the side of the chassis before bending upward or downward in the rack.
6. To remove a cable, disengage the locks and slowly pull the connector away from the port receptacle. LED indicator will turn off when the cable is unseated.

To uninstall the adapter card, see Uninstalling the Card.

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 “NVIDIA Technologies”:

```
lspci |grep -i Mellanox
Network controller: Mellanox Technologies MT2894 Family [ConnectX-6 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-6 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)

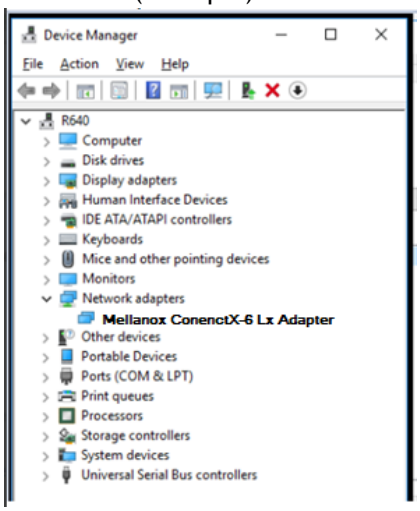

5. In the Value display box, check the fields VEN and DEV (fields are separated by '&'). In the display example above, notice the sub-string "PCI\VEN_15B3&DEV_1003": VEN is equal to 0x15B3 – this is the Vendor ID of NVIDIA Technologies; and DEV is equal to 1018 (for ConnectX-6 Lx) – this is a valid NVIDIA Technologies 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 Technologies PCI Device IDs can be found in the PCI ID repository at http://pci-ids.ucw.cz/read/PC/15b3.

Extraction Instructions

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

Driver Installation

Please use the relevant driver installation section.

- Windows Driver Installation
- Linux Driver Installation
- VMware Driver Installation

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]"
   ```

   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"
   ```

   > ⚠ 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.

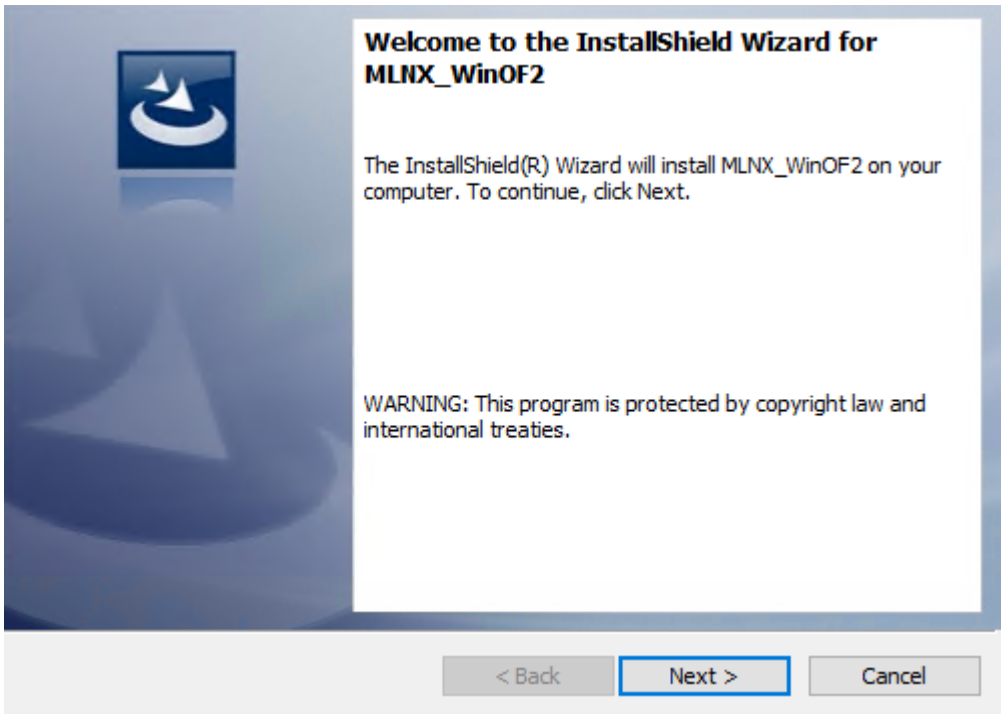

7. Read and accept the license agreement and click Next.

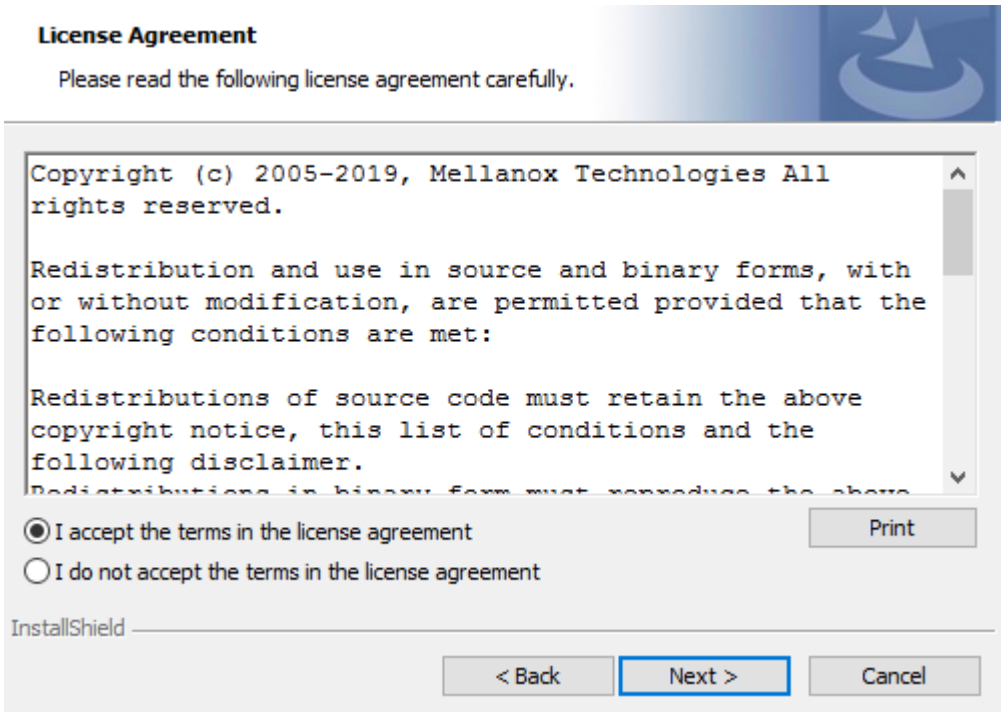

8. Select the target folder for the installation.

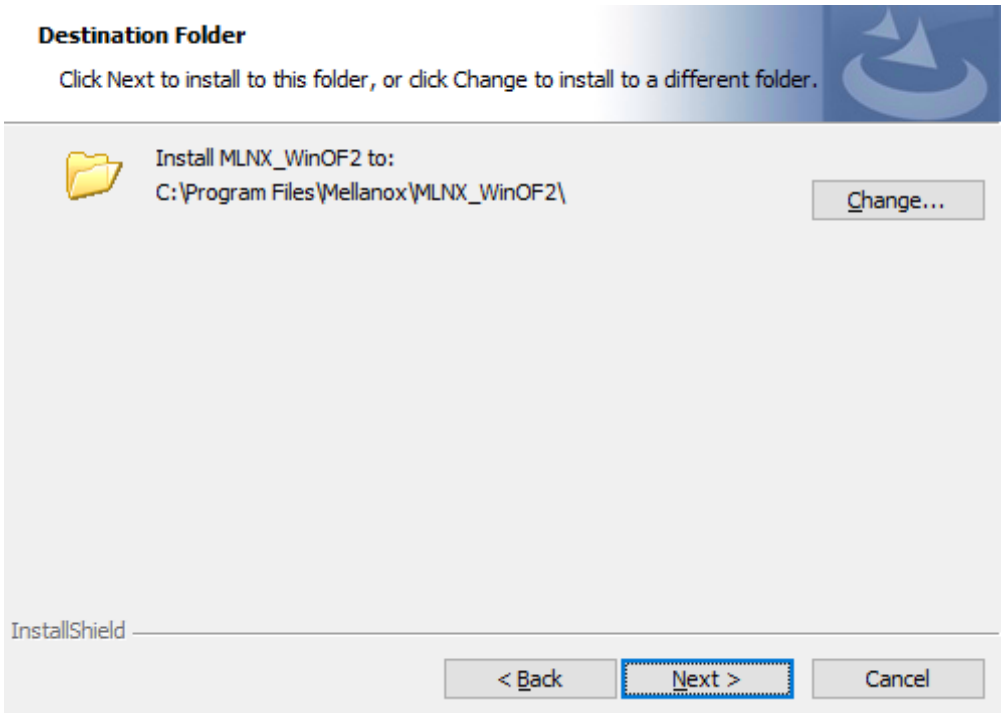

9. The firmware upgrade screen will be displayed in the following cases:
   - If the user has an OEM card. In this case, the firmware will not be displayed.
   - If the user has a standard NVIDIA® card with an older firmware version, the firmware will be updated accordingly. However, if the user has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

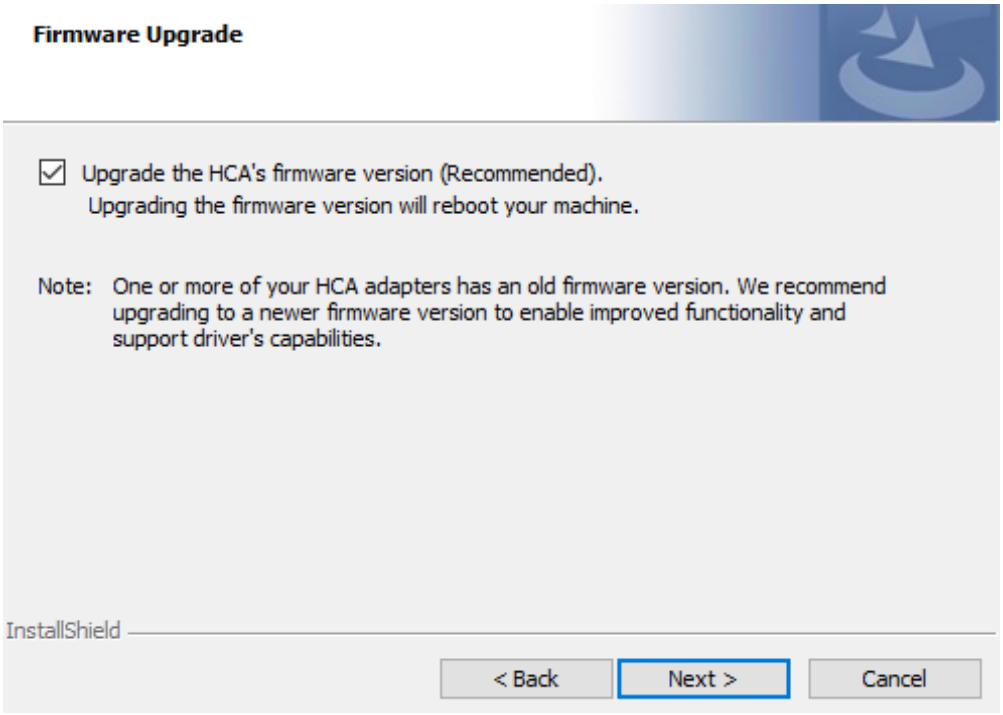

10. Select a Complete or Custom installation, follow Step a onward.

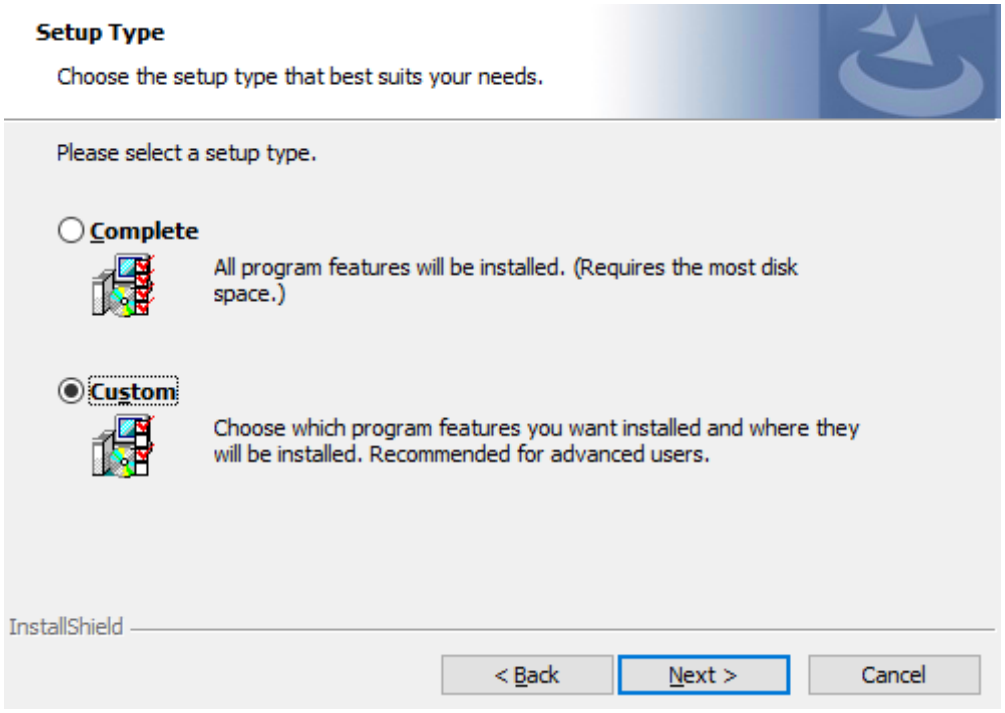

a. Select the desired feature to install:
   - Performances tools - install the performance tools that are used to measure performance in user environment
   - Documentation - contains the User Manual and Release Notes

- 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.

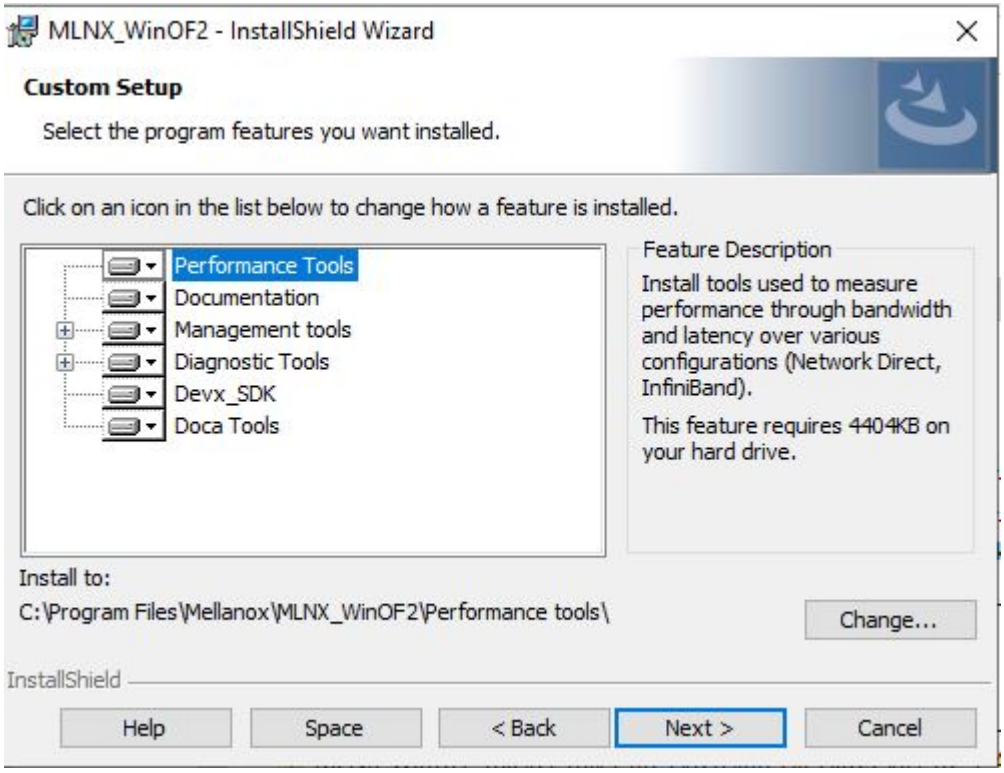

11. Click Install to start the installation.

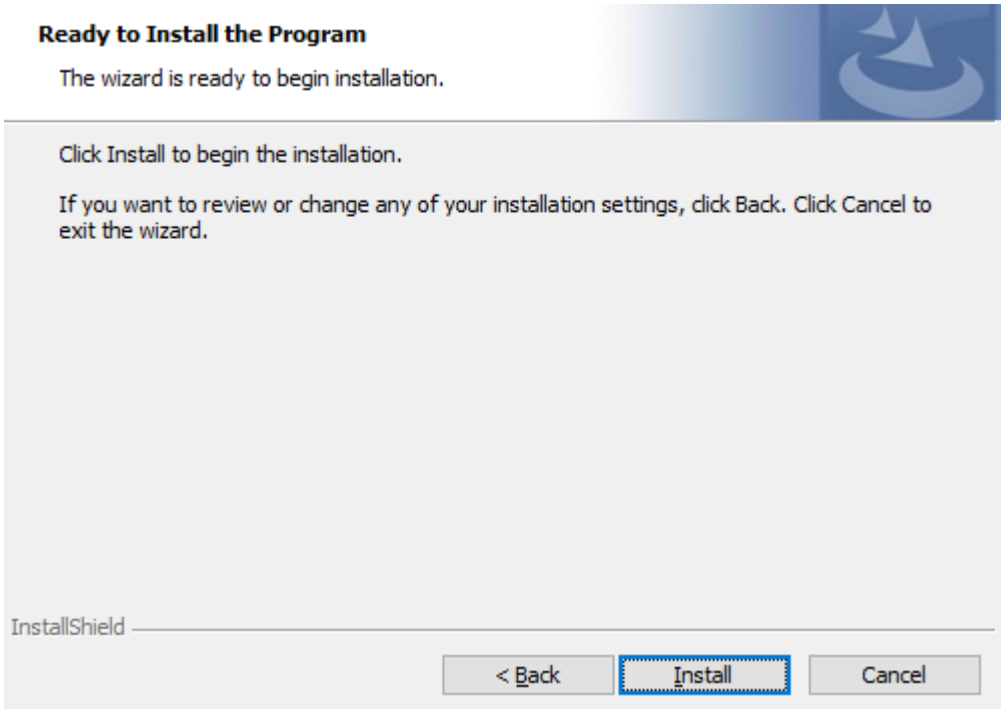

12. In case firmware upgrade option was checked in Step 7, you will be notified if a firmware upgrade is required (see).

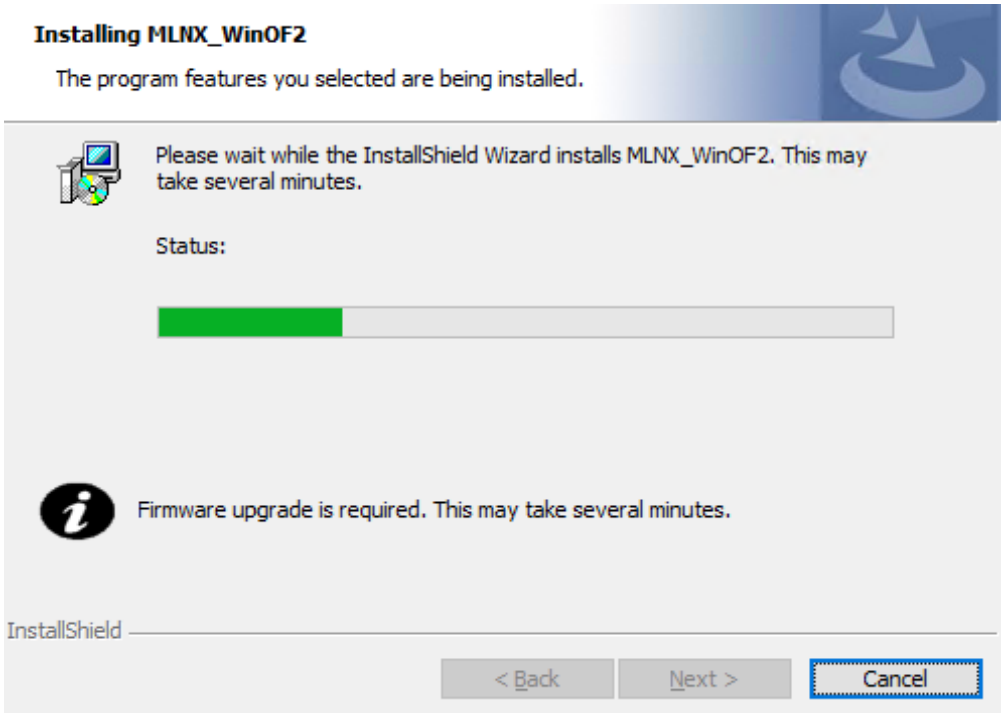

13. Click Finish to complete the installation.

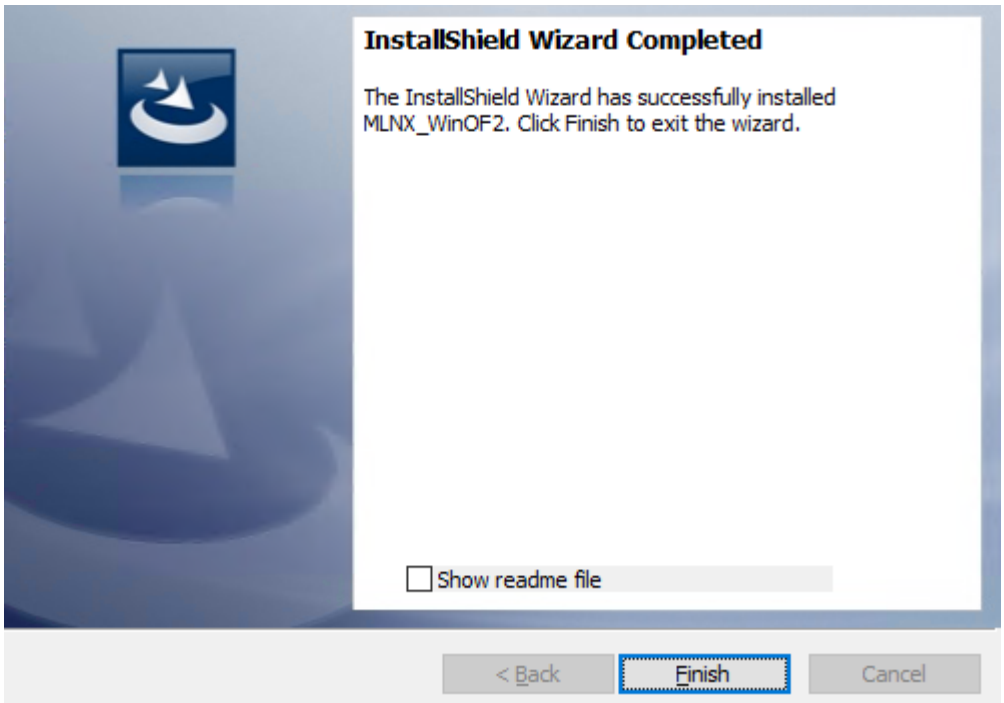

Unattended Installation

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

The following is an example of an unattended installation session.

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

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

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

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

Linux Driver Installation

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

Prerequisites

Requirements	Description
Platforms	A server platform with ConnectX®-6 Lx (EN) (firmware: fw-ConnectX6 Lx)
Required Disk Space for Installation	1GB
Device ID	For the latest list of device IDs, please visit the NVIDIA website at http://www.nvidia.com/page/firmware_HCA_FW_identification.
Operating System	Linux operating system. For the list of supported operating system distributions and kernels, please refer to the *NVIDIA OFED Release Notes* file.
Installer Privileges	The installation requires administrator (root) privileges on the target machine.

2. Download the ISO image to your host.
   The image's name has the format `MLNX_OFED_LINUX-<ver>-<OS label><CPU arch>.iso`.
   You can download and install the latest OpenFabrics Enterprise Distribution (OFED) software package available via the NVIDIA web site at Please visit 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_2190110032
PCI Device Name: 0b:00.0
Base MAC: 0000e41d2d5cf810
Versions: Current Available
FW 12.14.0114 12.14.0114
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_2190110032
PCI Device Name: 0b:00.0
Base MAC: 0000e41d2d5cf810
Versions: Current Available
FW 12.14.0114 N/A
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.

VMware Driver Installation

This section describes VMware Driver Installation.

Hardware and Software Requirements

Requirement	Description
Platforms	A server platform with an adapter card based on one of the following NVIDIA Technologies’ devices: • ConnectX®-6 Lx (EN) (firmware: fw-ConnectX6 Lx)
Device ID	For the latest list of device IDs, please visit NVIDIA's website.
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.

```
[server1]# ./mlxup
Querying Mellanox devices firmware ...
Device Type:      ConnectX-6 Lx
Part Number:      MCX631435AN-GDAB Description: ConnectX®-6 Lx EN adapter card, 50GbE, OCP3.0, With Host
management , Single-port QSFP28, PCIe 4.0 x8, No Crypto, Thumbscrew (Pull Tab) 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-6 Lx
Part Number:      MCX631435AN-GDAB Description: ConnectX®-6 Lx EN adapter card, 50GbE, OCP3.0, With Host
management , Single-port QSFP28, PCIe 4.0 x8, No Crypto, Thumbscrew (Pull Tab) Bracket

PCI Device Name: 0000:07:00.0
Base MAC:         0000e41d2da206d4
Versions:         Current          Available
                  FW 16.20.1000    16.24.1000

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
```

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

⚠ Please make sure to install the ConnectX-6 Lx OCP 3.0 card in a PCIe slot that is capable of supplying the power levels mentioned below for the appropriate use case

⚠ In Standby mode only port0 is available.

MCX631432AC-ADAB / MCX631432AE-ADAB / MCX631432AN-ADAB / MCX631432AS-ADAB Specifications

<table>
<tr><td rowspan="3">Physical</td><td colspan="6">Form Factor: Small Form Factor (SFF) OCP 3.0
Size: 4.52 in. x 2.99 in (115.00mm x 76.00mm)</td></tr>
<tr><td colspan="6">Connector: Dual SFP28 Ethernet (copper and optical)</td></tr>
<tr><td colspan="6">Retention Mechanism: Thumbscrew (Pull Tab) Bracket</td></tr>
<tr><td rowspan="3">Protocol Support</td><td colspan="6">Ethernet: 25GBASE-R, 20GBASE-KR2, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-KX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR</td></tr>
<tr><td>Data Rate:</td><td colspan="5">Ethernet 1/10/25 Gb/s</td></tr>
<tr><td colspan="6">PCI Express Gen 4.0: SERDES @ 16.0GT/s, 8 lanes (2.0 and 1.1 compatible)</td></tr>
<tr><td rowspan="9">Power and Airflow</td><td colspan="6">Voltage: 3.3VAUX, 12V</td></tr>
<tr><td>Power</td><td>Cable Type</td><td colspan="2">Active Mode</td><td colspan="2">Standby Mode</td></tr>
<tr><td>Typical Power(b)</td><td>Passive Cables</td><td colspan="2">11.8W</td><td colspan="2">TBD</td></tr>
<tr><td rowspan="2">Maximum Power</td><td>Passive Cables</td><td colspan="2">13.4W</td><td colspan="2">9.2W</td></tr>
<tr><td>1.5W Active Cables</td><td colspan="2">16.7W</td><td colspan="2">10.9W</td></tr>
<tr><td colspan="6">Maximum power available through SFP28 port: 1.5W</td></tr>
<tr><td rowspan="4">Airflow</td><td rowspan="2">Cable Type</td><td colspan="2">Hot Aisle
HSK to Port</td><td colspan="2">Cold Aisle @35C
Port to HSK</td></tr>
<tr><td>Active Mode @55C</td><td>Standby Mode @45C</td><td>Active Mode</td><td>Standby Mode</td></tr>
<tr><td>Passive Cable</td><td>250LFM</td><td>100LFM</td><td>100LFM</td><td>100LFM</td></tr>
<tr><td>Active 0.8W NVIDIA Cable</td><td>300LFM</td><td>100LFM</td><td>100LFM</td><td>100LFM</td></tr>
<tr><td rowspan="4">Environmental</td><td rowspan="2">Temperature</td><td>Operational</td><td colspan="4">0°C to 55°C</td></tr>
<tr><td>Non-operational</td><td colspan="4">-40°C to 70°C</td></tr>
<tr><td rowspan="2">Humidity</td><td>Operational</td><td colspan="4">10% to 85% relative humidity</td></tr>
<tr><td>Non-operational</td><td colspan="4">10% to 90% relative humidity</td></tr>
</table>

	Altitude (Operational)	3050m
Regulatory	**Safety**	CB / cTUVus / CE
	EMC	CE / FCC / VCCI / ICES / RCM
	RoHS	RoHS compliant

a. ConnectX-6 Lx adapters supplement the auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA Ethernet product.
b. Typical power for ATIS traffic load.
c. For both operational and non-operational states.
d. Airflow is measured in wind tunnel.

MCX631435AC-GDAB / MCX631435AE-GDAB / MCX631435AN-GDAB Specifications

Physical	**Form Factor:** Small Form Factor (SFF) OCP 3.0 **Size**: 4.52 in. x 2.99 in (115.00mm x 76.00mm)					
	Connector: Single QSFP28 Ethernet (copper and optical)					
	Retention Mechanism: Thumbscrew (Pull Tab) Bracket					
Protocol Support	**Ethernet**: 50GBASE-R2, 50GBASE-R4, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR					
	Data Rate:	Ethernet	1/10/25/40/50 Gb/s			
	PCI Express Gen 4.0: SERDES @ 16.0GT/s, 8 lanes (2.0 and 1.1 compatible)					
Power and Airflow	**Voltage**: 3.3VAUX, 12V					
	Power	**Cable Type**	**Active Mode**		**Standby Mode**	
	Typical Power[b]	Passive Cables	11.8W		TBD	
	Maximum Power	Passive Cables	13.4W		9.2W	
		3.5W Active Cables	17.4W		13W	
	Maximum power available through QSFP28 port: 3.5W					
	Airflow[d]	**Cable Type**	**Hot Aisle HSK to Port**		**Cold Aisle @35C Port to HSK**	
			Active Mode @55C	**Standby Mode @45C**	**Active Mode**	**Standby Mode**
		Passive Cable	200LFM	100LFM	100LFM	100LFM
		Active 2W NVIDIA Cable	350LFM	100LFM	100LFM	100LFM

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	

a. ConnectX-6 Lx adapters supplement the auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA Ethernet product.
b. Typical power for ATIS traffic load.
c. For both operational and non-operational states.
d. Airflow is measured in wind tunnel.

Board Mechanical Drawing and Dimensions

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

ConnectX-6 Lx Dual-Port SFP56 OCP 3.0	ConnectX-6 Lx Single-Port QSFP28 OCP 3.0
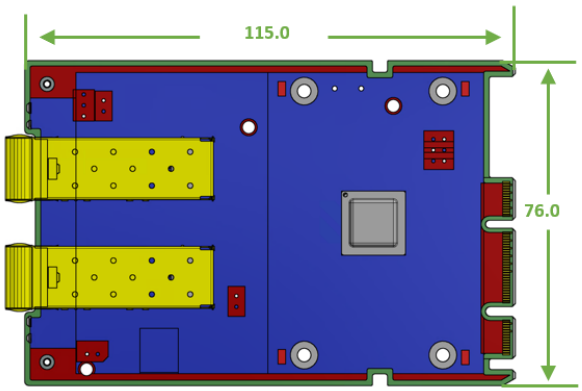	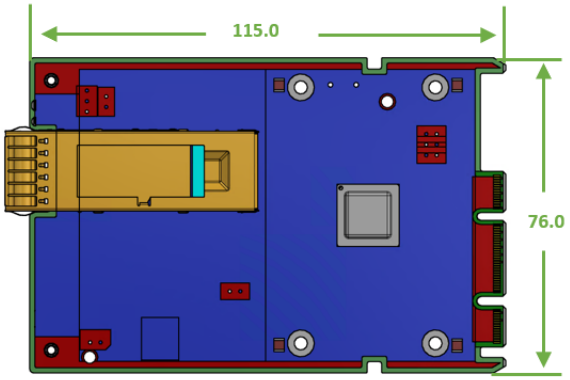

Bracket Mechanical Drawings and Dimensions

⚠ All dimensions are in millimeters. The bracket's mechanical tolerance is +/- 0.25mm.

Dual-port SFP28 Thumbscrew Bracket	Single-port QSFP28 Thumbscrew Bracket
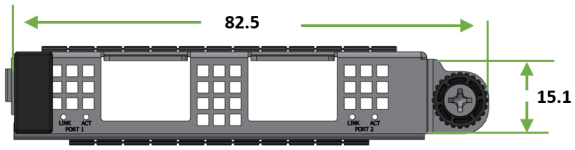	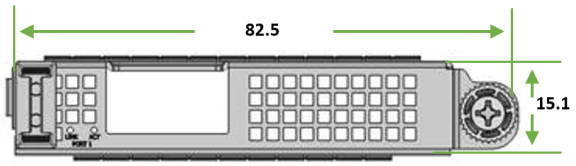

Monitoring

Thermal Sensors

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

Three thermal threshold definitions impact the overall system operation state:

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

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

Heatsink

A heatsink is attached to the ConnectX®-6 Lx IC in order to dissipate the heat. It is attached either by using two spring-loaded push pins that insert into four mounting holes.

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

Finding the MAC and Serial Number 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.

MCX631432AN-ADAB Board Label Example



Document Revision History

Date	Revision	Description of Changes
Nov. 2022	1.7	Updated Specifications tables
Oct. 2022	1.6	Updated Finding the MAC and Serial Number on the Adapter Card
Sep. 2022	1.5	Added MCX631432AS-ADAB to the relevant sections
Feb. 2022	1.4	Added table "Legacy (EOL) Ordering Part Numbers"
Jan. 2021	1.3	Updated Specifications tables
Oct. 2020	1.2	• Updated airflow numbers in "Specifications" • Updated networking ports LEDs indications
Oct. 2020	1.1	• Updated Specifications tables • Added board labels
Sep. 2020	1.0	First release

Notice

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

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

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

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

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

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

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

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

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

Trademarks

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



Copyright
© 2022 NVIDIA Corporation & affiliates. All Rights Reserved.

