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

Exported on Sep/19/2022 12:36 PM

Table of Contents

Introduction 9
Products Overview 9
ConnectX-6 Dx 25GbE Adapter Cards for OCP 3.0 9
ConnectX-6 Dx 100GbE Adapter Cards for OCP 3.0 10
ConnectX-6 Dx 200GbE Adapter Cards for OCP 3.0 11
Features and Benefits 12
Operating Systems/Distributions 14
Connectivity 14
Interfaces 15
ConnectX-6 Dx IC 16
Encryption 16
NVIDIA Multi-HostTM Support 16
Ethernet Networking Interfaces 18
PCI Express Interface 18
Networking Ports LEDs Interface 18
SMBus Interface 19
Voltage Regulators 19
CPLD Interface 19
Heatsink 20
Thermal Sensors 20
Hardware Installation 21
Safety Warnings 21
Installation Procedure Overview 22
System Requirements 22
Hardware Requirements 23
Airflow Requirements 23
Software Requirements 23
Safety Precautions 24
Pre-Installation Checklist 24
OCP 3.0 Bracket Replacement Instructions 24
OCP 3.0 Adapter Card Installation Instructions 24
Cards with Thumbscrew (Pull-tab) Bracket 25

Cards with Internal Lock Bracket 26
Cards with Ejector-Latch Bracket 26
Cables and Modules 27
Identifying the Card in Your System 28
Card Extraction Instructions 29
Safety Precautions 24
Cards Extraction Instructions 30
Driver Installation 32
Windows Driver Installation 32
Hardware and Software Requirements 32
Downloading NVIDIA WinOF-2 Driver 32
Installing NVIDIA WinOF-2 Driver 33
Installation Results 38
Extracting Files Without Running Installation 38
Uninstalling NVIDIA WinOF-2 Driver 40
Firmware Upgrade 40
Deploying the Driver on a Nano Server 41
Linux Driver Installation 42
Prerequisites 42
Downloading NVIDIA OFED 42
Installing NVIDIA OFED 43
Installing MLNX_OFED Using YUM 48
Installing MLNX_OFED Using apt-get Tool 51
Updating Firmware After Installation 52
UEFI Secure Boot 53
Performance Tuning 55
VMware Driver Installation 55
Hardware and Software Requirements 55
Installing NVIDIA NATIVE ESXi Driver for VMware vSphere 55
Removing Earlier NVIDIA Drivers 56
Firmware Programming 56
Updating Adapter Firmware 57
Troubleshooting 58
General Troubleshooting 58

Linux Troubleshooting 58
Windows Troubleshooting 59
Specifications 60
MCX623432AE-ADAB / MCX623432AN-ADAB / MCX623432AC-ADAB / MCX623432AS-ADAB Specifications 60
MCX623432AC-GDAB / MCX623432AN-GDAB / MCX623432AS-GDAI Specifications 61
MCX623435MN-CDAB Specifications 62
MCX623439MC-CDAB / MCX623439MN-CDAB Specifications 63
MCX623436AC-CDAB / MCX623436AN-CDAB / MCX623436AS-CDAB / MCX623436MC-CDAB / MCX623436MN-CDAB / MCX623436MS-CDAB / MCX623436AS-CDAI Specifications 64
MCX623430MS-CDAB Specifications 65
MCX623435AC-CDAB / MCX623435AE-CDAB / MCX623435AN-CDAB / MCX623435AC-VDAB / MCX623435AE-VDAB / MCX623435AN-VDAB / MCX623435AS-VDAI Specifications 66
Board Mechanical Drawing and Dimensions 68
Bracket Mechanical Drawings and Dimensions 69
Dual-port SFP28/SFP56 Thumbscrew Bracket 69
Dual-port QSFP56 Thumbscrew Bracket 69
Single-port QSFP56 Thumbscrew Bracket 70
Finding the MAC and Serial Number on the Adapter Card 71
Document Revision History 73

About This Manual

This is the User Guide for Ethernet NVIDIA® ConnectX®-6 Dx adapter cards for Open Compute Project Spec 3.0. These adapters' connectivity provides 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 Dx Ethernet adapter cards for OCP Spec 3.0.

OPN	Marketing Description
MCX623432AS-ADAB	ConnectX®-6 Dx EN adapter card, **25GbE**, OCP3.0, With Host management, **Dual**-port **SFP28**, PCIe 4.0 **x16**, Secure Boot, No Crypto, **Thumbscrew *(Pull Tab)* Bracket**
MCX623432AC-GDAB	ConnectX®-6 Dx EN adapter card, **50GbE**, OCP3.0, With Host Management, **Dual**-port **SFP56**, PCIe 4.0 **x16**, **Crypto and Secure Boot, Thumbscrew *(Pull Tab)* Bracket**
MCX623435AN-CDAB	ConnectX®-6 Dx EN adapter card, **100GbE**, OCP3.0, With Host management, **Single**-port **QSFP56**, PCIe 4.0 **x16**, No Crypto, **Thumbscrew (*Pull Tab*) Bracket**
MCX623435AC-CDAB	ConnectX-6 Dx EN adapter card, **100GbE,** OCP3.0, With Host Management, **Single**-port **QSFP56**, PCIe 4.0 **x16**, **Crypto and Secure Boot, Thumbscrew *(Pull Tab)* Bracket**
MCX623436MN-CDAB	ConnectX®-6 Dx EN adapter card, **100GbE**, OCP3.0, With Host management, **Dual**-port **QSFP56**, **Multi-Host or Socket Direct**, PCIe 4.0 **x16**, No Crypto, **Thumbscrew (*Pull Tab*) Bracket**
MCX623436MS-CDAB	ConnectX®-6 Dx EN adapter card, **100GbE**, OCP3.0, With Host management, **Dual**-port **QSFP56**, **Multi-Host or Socket Direct**, PCIe 4.0 **x16**, Secure Boot, No Crypto, **Thumbscrew (*Pull Tab*) Bracket**
MCX623430MS-CDAB	ConnectX-6 Dx EN adapter card, **100GbE** OCP3.0, With Host Management, **Dual**-port **DSFP**, **Multi-Host or Socket Direct**, PCIe 4.0 **x16**, Secure Boot, No Crypto, **Thumbscrew (Pull Tab) Bracket**
MCX623436AC-CDAB	ConnectX®-6 Dx EN adapter card, **100GbE**, OCP3.0, With Host Management, **Dual**-port **QSFP56**, PCIe 4.0 **x16**, **Crypto and Secure Boot, Thumbscrew *(Pull Tab)* Bracket**
MCX623436AN-CDAB	ConnectX®-6 Dx EN adapter card, **100GbE**, OCP3.0, With Host Management, **Dual**-port **QSFP56**, PCIe 4.0 **x16**, No Crypto, **Thumbscrew *(Pull Tab)* Bracket**
MCX623435AC-CDAI	ConnectX-6 Dx EN adapter card, **100GbE**, OCP3.0, With Host Management, **Single**-port **QSFP56,** PCIe 4.0 x**16**, **Crypto and Secure Boot**, **Internal Lock Bracket**
MCX623435AN-VDAB	ConnectX®-6 Dx EN adapter card, **200GbE**, OCP3.0, With Host Management, **Single**-port **QSFP56**, PCIe 4.0 **x16**, No Crypto, **Thumbscrew (Pull Tab) Bracket**

Legacy (EOL) Ordering Part Numbers

OPN	Marketing Description
MCX623432AE-ADAB	ConnectX®-6 Dx EN adapter card, **25GbE**, OCP3.0, With Host Management, **Dual**-port **SFP28**, PCIe 4.0 **x16, Crypto**, No Secure Boot, **Thumbscrew *(Pull Tab)* Bracket**
MCX623432AS-GDAI	ConnectX®-6 Dx EN adapter card, **50GbE**, OCP3.0, With Host Management, **Dual**-port **SFP56**, PCIe 4.0 **x16**, Secure Boot, No Crypto, **Internal Lock Bracket**
MCX623435AE-CDAB	ConnectX®-6 Dx EN adapter card, **100GbE**, OCP3.0, With Host Management, **Single**-port **QSFP56**, PCIe 4.0 **x16**, **Crypto**, No Secure Boot, **Thumbscrew (*Pull Tab*) Bracket**
MCX623435AS-VDAI	ConnectX®-6 Dx EN adapter card, **200GbE**, OCP3.0, With Host Management, **Single**-port **QSFP56**, PCIe 4.0 **x16**, Secure Boot, No Crypto, **Internal Lock Bracket**
MCX623436AE-CDAB	ConnectX-6 Dx EN adapter card, 100GbE, OCP3.0, With Host Management, Dual-port QSFP56, PCIe 4.0 x16, Crypto, No Secure Boot, Thumbscrew (Pull Tab) Bracket
MCX623436AS-CDAI	ConnectX®-6 Dx EN adapter card, **100GbE** OCP3.0, With Host Management, **Dual**-port **QSFP56**, PCIe 4.0 **x16**, Secure Boot, No Crypto, **Internal Lock Bracket**
MCX623439MC-CDAB	ConnectX®-6 Dx EN adapter card, **100GbE**, OCP3.0, With Host Management, **Single**-port **QSFP56**, **Multi-Host or Socket Direct**, PCIe 4.0 **x16**, **Crypto and Secure Boot, Thumbscrew *(Pull Tab)* Bracket**
MCX623432AC-ADAB	ConnectX®-6 Dx EN adapter card, **25GbE**, OCP3.0, With Host Management, **Dual**-port **SFP28**, PCIe 4.0 **x16**, **Crypto and Secure Boot, Thumbscrew *(Pull Tab)* Bracket**
MCX623432AN-ADAB	ConnectX®-6 Dx EN adapter card, **25GbE**, OCP3.0, With Host management, **Dual**-port **SFP28**, PCIe 4.0 **x16**, No Crypto, **Thumbscrew *(Pull Tab)* Bracket**
MCX623432AN-GDAB	ConnectX®-6 Dx EN adapter card, **50GbE**, OCP3.0, With Host Management, **Dual**-port **SFP56**, PCIe 4.0 **x16**, No Crypto, **Thumbscrew *(Pull Tab)* Bracket**
MCX623435AC-VDAB	ConnectX®-6 Dx EN adapter card, **200GbE**, OCP3.0, With Host Management, **Single**-port **QSFP56**, PCIe 4.0 **x16, Crypto and Secure Boot, Thumbscrew (Pull Tab) Bracket**
MCX623435MN-CDAB	ConnectX®-6 Dx EN adapter card, **100GbE**, OCP3.0, With Host Management, **Single**-port **QSFP56**, **Multi-Host or Socket Direct**, PCIe 4.0 **x16**, No Crypto, **Thumbscrew *(Pull Tab)* Bracket**

Intended Audience

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

Technical Support

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

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

Customers who purchased NVIDIA Global Support Services, please see your contract for details regarding Technical Support.
Customers who purchased NVIDIA products through an NVIDIA-approved reseller should first seek assistance through their reseller.

Related Documentation

NVIDIA OFED for Linux User Manual and Release Notes	User Manual describing OFED features, performance, band diagnostic, tools content, and configuration. See NVIDIA OFED for Linux Documentation.
WinOF-2 for Windows User Manual and Release Notes	User Manual describing WinOF-2 features, performance, Ethernet diagnostic, tools content, and configuration. See WinOF-2 for Windows Documentation.
NVIDIA VMware for Ethernet User Manual and Release Notes	User Manual describing the various components of the NVIDIA ConnectX® NATIVE ESXi stack. See http://www.nvidia.com Products > Software > Ethernet Drivers > VMware Driver
NVIDIA Firmware Utility (mlxup) User Manual and Release Notes	NVIDIA firmware update and query utility used to update the firmware. See http://www.nvidia.com > Products > Software > Firmware Tools >mlxup Firmware Utility
NVIDIA Firmware Tools (MFT) User Manual	User Manual describing the set of MFT firmware management tools for a single node. See MFT User Manual.
IEEE Std 802.3 Specification	IEEE Ethernet specification at http://standards.ieee.org/
PCI Express Specifications	Industry Standard PCI Express Base and Card Electromechanical Specifications at https://pcisig.com/specifications
Open Compute Project 3.0 Specification	https://www.opencompute.org/
NVIDIA LinkX Interconnect Solutions	NVIDIA LinkX cables and transceivers are designed to maximize the performance of High-Performance Computing networks, requiring high-bandwidth, low-latency connections between compute nodes and switch nodes. NVIDIA offers one of the industry's broadest portfolios of 40Gb/s, 56Gb/s, 100Gb/s, and 200Gb/s cables, including Direct Attach Copper cables (DACs), copper splitter cables, Active Optical Cables (AOCs), and transceivers in a wide range of lengths from 0.5m to 10km. NVIDIA tests every product in an end-to-end environment ensuring a Bit Error Rate of less than 1E-15. Read more at https://www.nvidia.com/products/interconnect/ethernet-overview.php

Document Conventions

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

Introduction

This is the *User Guide* for Ethernet adapter cards based on the ConnectX®-6 Dx integrated circuit device for OCP Spec 3.0. These adapters connectivity provide the highest performing low latency and most flexible interconnect solution for servers supporting OCP 3.0 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.

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

> ⚠ Please note the following OPNs are shipped with an assembled thumbscrew (pull tab) bracket. For other retention mechanism brackets, please contact NVIDIA.

Products Overview

ConnectX-6 Dx 25GbE Adapter Cards for OCP 3.0

OPN	MCX623432AN-ADAB	MCX623432AC-ADAB	MCX623432AE-ADAB	MCX623432AS-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 3.0 / 4.0 x16 SERDES @ 8.0GT/s / 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	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE
Host Management	✓	✓	✓	✓
Crypto	-	✓	✓	-
Secure Boot	✓	✓	-	✓
Device ID	4125 for Physical Function (PF) 4126 for Virtual Function (VF)			

ConnectX-6 Dx 50GbE Adapter Cards for OCP 3.0

OPN	MCX623432AN-GDAB	MCX623432AC-GDAB	MCX623432AS-GDAI
Network Connector Type	Dual-port SFP56	Dual-port SFP56	Dual-port SFP56
Ethernet Data Rate	10/25/50 Gb/s		
PCI Express Connectors	PCIe Gen 3.0 / 4.0 x16 SERDES @ 8.0GT/s / 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	Thumbscrew (Pull Tab) Bracket	Internal Lock Bracket
RoHS	RoHS Compliant	RoHS Compliant	RoHS Compliant
IC Part Number	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE
Host Management	✓	✓	✓
Crypto	-	✓	-
Secure Boot	-	✓	✓
Device ID	4125 for Physical Function (PF) 4126 for Virtual Function (VF)		

ConnectX-6 Dx 100GbE Adapter Cards for OCP 3.0

OPN	MCX623436AN-CDAB	MCX623436AC-CDAB	MCX623435MN-CDAB	MCX623436MN-CDAB	MCX623436MS-CDAB	MCX623436AS-CDAI	MCX623435AC-CDAI	MCX623435AC-CDAB	MCX623435AE-CDAB	MCX623435AN-CDAB	MCX623439MC[/N]-CDAB	MCX623430MS-CDAB
Network Connector Type	Dual-port QSFP56	Dual-port QSFP56	Single-port QSFP56	Dual-port QSFP56	Dual-port QSFP56	Dual-port QSFP56	Single-port QSFP56	Single-port QSFP56	Single-port QSFP56	Single-port QSFP56	Single-port QSFP56	Dual-port QSFP56
Ethernet Data Rate	10/25/40/50/100 Gb/s											
PCI Express Connectors	PCIe Gen 3.0 / 4.0 x16 SERDES @ 8.0GT/s / 16.0GT/s											
Dimensions	Small Form Factor (SFF) OCP 3.0 4.52 in. x 2.99 in (115.00mm x 76.00mm)											

OPN	MCX6234 36AN-CDAB	MCX6234 36AC-CDAB	MCX6234 35MN-CDAB	MCX6234 36MN-CDAB	MCX6234 36MS-CDAB	MCX6234 36AS-CDAI	MCX6234 35AC-CDAI	MCX6234 35AC-CDAB	MCX6234 35AE-CDAB	MCX6234 35AN-CDAB	MCX6234 39MC[/N]-CDAB	MCX6234 30MS-CDAB
Retention Mechanism	Thumbscrew (Pull Tab) Bracket	Thumbscrew (Pull Tab) Bracket	Thumbscrew (Pull Tab) Bracket	Thumbscrew (Pull Tab) Bracket	Thumbscrew (Pull Tab) Bracket	Internal Lock Bracket	Internal Lock Bracket	Thumbscrew (Pull Tab) Bracket	Thumbscrew (Pull Tab) Bracket	Thumbscrew (Pull Tab) Bracket	Thumbscrew (Pull Tab) Bracket	Thumbscrew (Pull Tab) Bracket
RoHS	RoHS Compliant	RoHS Compliant	RoHS Compliant	RoHS Compliant	RoHS Compliant	RoHS Compliant	RoHS Compliant	RoHS Compliant	RoHS Compliant	RoHS Compliant	RoHS Compliant	RoHS Compliant
IC Part Number	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE
Host Management	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Crypto	-	✓	-	-	-	-	✓	✓	✓	-	✓	✓
Secure Boot	-	✓	-	-	✓	✓	✓	✓	-	-	✓	✓
Mutli-Host or Socket Direct	-	-	✓	✓	✓	-	-	-	-	-	✓	-
Device ID	4125 for Physical Function (PF) 4126 for Virtual Function (VF)											

ConnectX-6 Dx 200GbE Adapter Cards for OCP 3.0

OPN	MCX623435AN-VDAB	MCX623435AC-VDAB	MCX623435AS-VDAI
Network Connector Type	Single-port QSFP56	Single-port QSFP56	Single-port QSFP56
Ethernet Data Rate	10/25/40/50/100/200 Gb/s		
PCI Express Connectors	PCIe Gen 3.0 / 4.0 x16 SERDES @ 8.0GT/s / 16.0GT/s		
Dimensions	Small Form Factor (SFF) OCP 3.0 4.52 in. x 2.99 in (115.00mm x 76.00mm)		

OPN	MCX623435AN-VDAB	MCX623435AC-VDAB	MCX623435AS-VDAI
Retention Mechanism	Thumbscrew (Pull Tab) Bracket	Thumbscrew (Pull Tab) Bracket	Internal Lock Bracket
RoHS	RoHS Compliant	RoHS Compliant	RoHS Compliant
IC Part Number	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE	MT28928A0-XCCF-CE
Host Management	✓	✓	✓
Crypto	-	✓	-
Secure Boot	-	✓	✓
Device ID	4125 for Physical Function (PF) 4126 for Virtual Function (VF)		

For more detailed information see Specifications.

Features and Benefits

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

Feature	Description
PCI Express (PCIe)	Uses PCIe Gen 3.0 (8GT/s) or Gen 4.0 (16GT/s) through x16 edge connector.
Up to 200 Gigabit Ethernet	NVIDIA adapters comply with the following IEEE 802.3 standards: • 200GbE / 100GbE / 50GbE / 40GbE / 25GbE / 10GbE / 1GbE • IEEE 802.3bj, 802.3bm 100 Gigabit Ethernet • IEEE 802.3by, Ethernet Consortium25, 50 Gigabit Ethernet, supporting all FEC modes • IEEE 802.3ba 40 Gigabit Ethernet • IEEE 802.3by 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)

Feature	Description
Overlay Networks	In order to better scale their networks, data center operators often create overlay networks that carry traffic from individual virtual machines over logical tunnels in encapsulated formats such as NVGRE and VXLAN. While this solves network scalability issues, it hides the TCP packet from the hardware offloading engines, placing higher loads on the host CPU. ConnectX®-6 Dx 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 Dx, 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 Dx 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 Dx 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 Dx provides dedicated adapter resources and guaranteed isolation and protection for virtual machines within the server.
Storage Acceleration	A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols can leverage RDMA for high-performance storage access. • NVMe over Fabric offloads for the target machine
SR-IOV	ConnectX®-6 Dx 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.

Feature	Description
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
NVIDIA Multi-Host	NVIDIA® Mellanox® Multi-Host technology enables next-generation Cloud, Web 2.0 and high-performance data centers to design and build new scale-out heterogeneous compute and storage racks with direct connectivity between multiple hosts and the centralized network controller. This enables direct data access with the lowest latency to significantly improve densities and maximizes data transfer rates. For more information, please visit NVIDIA Multi-Host Solutions.

Operating Systems/Distributions

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

Connectivity

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

Interfaces

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

> ⚠ The ConnectX-6 Dx 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 Dx for OCP 3.0 Adapter Card Interfaces - Component Side

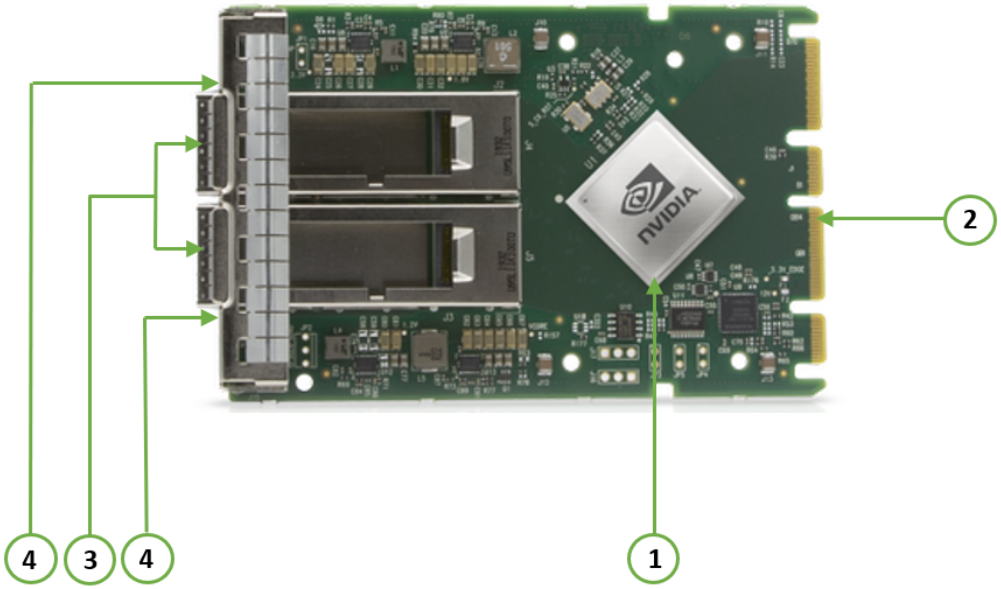

Callout	Item	Description
1	ConnectX-6 Dx IC	NVIDIA ConnectX-6 Dx IC on the board.
2	PCI Express Interface	PCIe Gen 3.0/4.0 through an x16 edge connector.
3	Ethernet Networking Interface	Ethernet traffic is transmitted through the adapter's SFP28/SFP56/QSFP56 connectors. The networking connectors allow 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.

Callout	Item	Description
	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.

ConnectX-6 Dx IC

The ConnectX®-6 Dx EN family of adapter IC devices delivers two ports of 10/25/40/50/100Gb/s or a single port of 200Gb/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 Dx 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, guaranteeing the network adapter's integrity.

NVIDIA Multi-Host™ Support

In addition to building exceptionally high bandwidth to the data center, the ConnectX-6 Dx device enables leveraging this speed across the entire data center utilizing its NVIDIA Multi-Host feature.

Using its 32-lane PCI Express interface, a single ConnectX-6 Dx device can provide 200GbE interconnect for up to four independent hosts without any performance degradation.

The figure below shows a ConnectX-6 Dx device with NVIDIA Multi-Host connected to four separate hosts, each with a PCIe x4 interface, on one side to a switch on the other side.

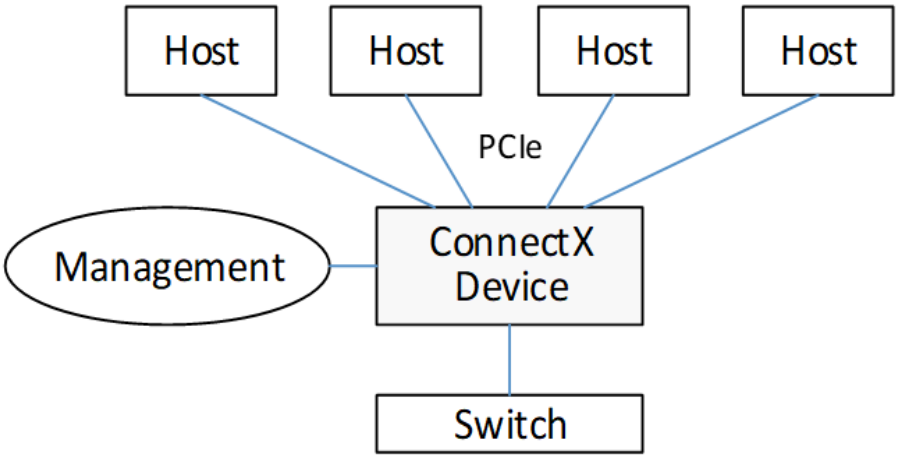

The below bifurcation are optional for the adapter's x16 PCIe interface:

- x1 PCIe x16, x1 PCIe x8, x1 PCIe x4
- x2 PCIe x8, x2 PCIe x4, x2 PCIe x2, x2 PCIe x1
- x4 PCIe x4, x4 PCIe x2, x4 PCIe x1

Multi-host capable cards also support Socket-Direct applications and work as regular Single-Host cards, depending on the type of server they are plugged into, assuming the server complies with the OCP 3.0 spec.

According to the OCP 3.0 spec, the adapter card advertises its capability through the PRSNTB[3:0]# pins. The server determines the configuration through the BIF[2:0]# pins, which it drives to the adapter card.
The NVIDIA OCP3.0 card has an internal logic that uses the BIF[2:0]# data and determines the correct operating mode to boot at. The combination of the PRSNTB[3:0]# and BIF[2:0]# pins deterministically sets the PCIe lane width for a given combination of OCP 3.0 cards and baseboard. The logic and the decoding table can be found in the OCP 3.0 spec (Chapter 3.5 PCIe Bifurcation Mechanism.

For example:
the NVIDIA OCP 3.0 Multi-host adapter drives 0100 on PRSNTB[3:0]# to the server.

If Server Drivers	Adapter PCIe Mode
000	Single-Host Mode: **x1 PCIe x16**
001	Socket Direct Mode: **x2 PCIe x8**
010	Socket Direct Mode: **x4 PCIe x4**
101	Multi-Host Mode: **x2 PCIe x8**
110	Multi-Host Mode: **x4 PCIe x4**

Ethernet Networking Interfaces

The network ports of the ConnectX®-6 Dx adapter card are compliant with the IEEE 802.3 Ethernet standards listed in Features and Benefits. Ethernet traffic is transmitted through the networking 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 Dx adapter cards.

Supported PCIe Interface	Features
PCIe Gen 3.0/4.0 (1.1 and 2.0 compatible) through x16 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.

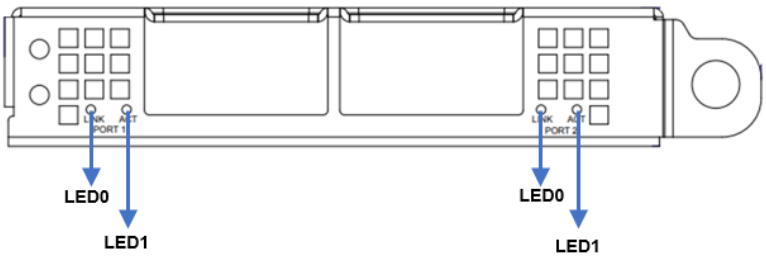

Link Indications

LED and State	Description
1Hz blinking Yellow	Beacon command for locating the adapter card

<table>
<tr><th>LED and State</th><th>Description</th></tr>
<tr><td>4Hz blinking Yellow</td><td>Indicates an error with the link. The error can be one of the following:
<table>
<tr><th>Error Type</th><th>Description</th><th>LED Behavior</th></tr>
<tr><td>I²C</td><td>I²C access to the networking ports fails</td><td>Blinks until the error is fixed</td></tr>
<tr><td>Over-current</td><td>Over-current condition of the networking ports</td><td>Blinks until the error is fixed</td></tr>
</table></td></tr>
<tr><td>LED0 -
Link Speed</td><td>• 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.</td></tr>
<tr><td>LED1 -
Activity</td><td>• A blinking Green indicates a valid link with data transfer.
• If LED1 is off, then there is no activity</td></tr>
</table>

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 Dx technology maintains support for manageability through a BMC. ConnectX®-6 Dx 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 server's specific manageability solution, 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 and the scan chain. It draws its power supply from 3.3V_EDGE.

Heatsink

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

Thermal Sensors

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

There are three thermal threshold definitions which impact the overall system operation state:

- Warning – 105°C: On managed systems only: When the device crosses the 100°C threshold, a Warning Threshold message will be 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 will automatically shut down the device.
- Emergency – 130°C: In case the firmware fails to shut down the device upon crossing the Critical threshold, the device will auto-shutdown 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 in accordance with the readouts and the needs of the above-mentioned IC thermal requirements.

Hardware Installation

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

Safety Warnings

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

Note that not all warnings are relevant to all models.

	**General Installation Instructions** Read all installation instructions before connecting the equipment to the power source.
	**Jewelry Removal Warning** Before you install or remove equipment that is connected to power lines, remove jewelry such as bracelets, necklaces, rings, watches, and so on. Metal objects heat up when connected to power and ground and can meltdown, causing serious burns and/or welding the metal object to the terminals.
	**Over-temperature** This equipment should not be operated in an area with an ambient temperature exceeding the maximum recommended: 55°C (131°F). An airflow of 200LFM at this maximum ambient temperature is required for HCA cards and NICs. To guarantee proper airflow, allow at least 8cm (3 inches) of clearance around the ventilation openings.
	**During Lightning - Electrical Hazard** During periods of lightning activity, do not work on the equipment or connect or disconnect cables.
	**Copper Cable Connecting/Disconnecting** Some copper cables are heavy and not flexible, as such, they should be carefully attached to or detached from the connectors. Refer to the cable manufacturer for special warnings and instructions.
	**Equipment Installation** This equipment should be installed, replaced, or serviced only by trained and qualified personnel.
	**Equipment Disposal** The disposal of this equipment should be in accordance to all national laws and regulations.
	**Local and National Electrical Codes** This equipment should be installed in compliance with local and national electrical codes.

Hazardous Radiation Exposure

- Caution – Use of controls or adjustment or performance of procedures other than those specified herein may result in hazardous radiation exposure.For products with optical ports.
- CLASS 1 LASER PRODUCT and reference to the most recent laser standards: IEC 60 825-1:1993 + A1:1997 + A2:2001 and EN 60825-1:1994+A1:1996+ A2:20

Installation Procedure Overview

The installation procedure of ConnectX-6 Dx adapter cards for OCP Spec 3.0 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 Dx adapter card for OCP spec 3.0 in the system	Refer to Installation Instructions
7	Connect cables or modules to the card	Refer to Cables and Modules
8	Identify the ConnectX-6 Dx 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 gaseous and dust (particulate) contamination levels.

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 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 Dx 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
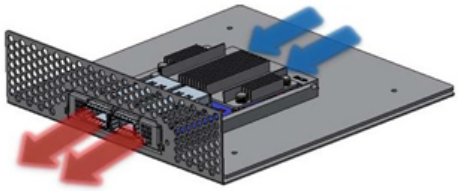	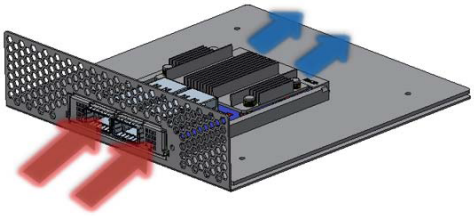

Software Requirements

- 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 Dx adapter card.
   Unpack and remove the ConnectX-6 Dx 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.

   > 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 installing the ConnectX-6 Dx card, ensure 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

This section provides detailed instructions on installing your adapter card in a system. The below table lists the different ConnectX-6 Dx OCP 3.0 retention mechanisms and offers direct links to installation instructions per bracket type.

Retention Mechanism	Installation Instructions
Thumbscrew (Pull-tab) Bracket	Installation Instructions for Cards with Thumbscrew (Pull-tab) Bracket

Retention Mechanism	Installation Instructions
Internal-Lock Bracket	Installation Instructions for Cards with Internal Lock
Ejector-Latch Bracket	Installation Instructions for Cards with Internal-Lock Bracket

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

Cards with Thumbscrew (Pull-tab) Bracket

1. Before installing the card, ensure the system is off, and the power cord is not connected to the server. Please follow proper electrical grounding procedures.
2. Open the system case.
3. Align the card with the system rails.

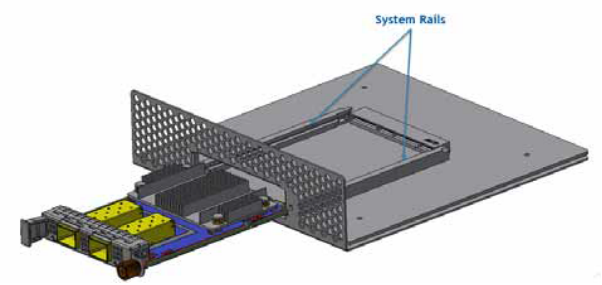

4. Push the card until connectors are in full mate.

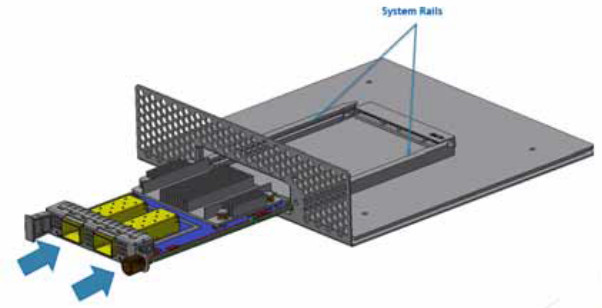

5. Turn the captive screw clockwise until firmly locked.

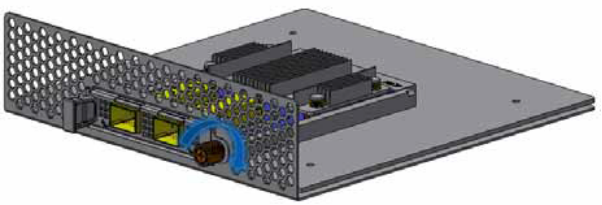

Cards with Internal Lock Bracket

1. Before installing the card, ensure the system is off, and the power cord is not connected to the server. Please follow proper electrical grounding procedures.
2. Open the system case.
3. Align the card with the system rails.

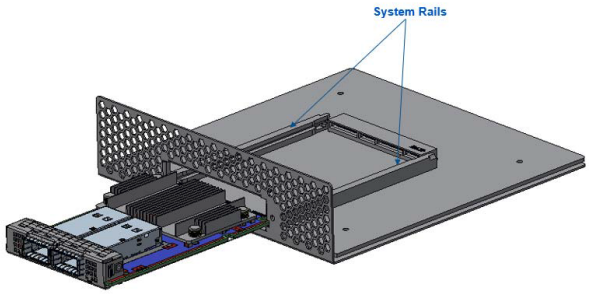

4. Push the card until connectors are in full mate and a clicking sound is heard.

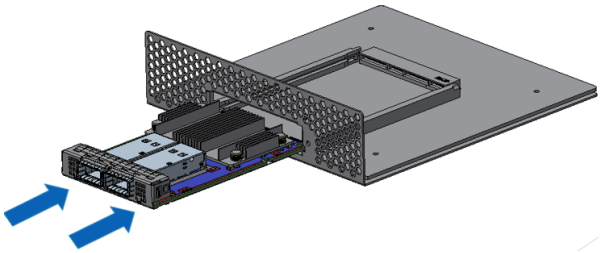

Cards with Ejector-Latch Bracket

1. Before installing the card, ensure the system is off, and the power cord is not connected to the server. Please follow proper electrical grounding procedures.
2. Open the system case.
3. Align the card with the system rails while ensuring the ejector latch is open.

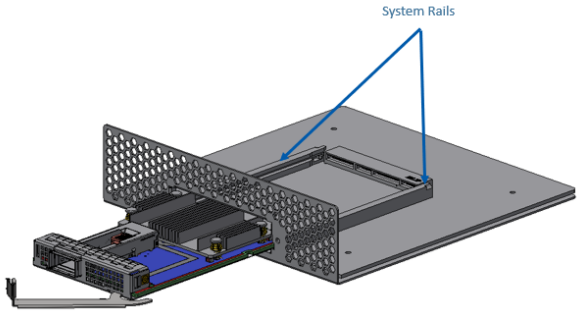

4. Push the card until connectors are in full mate.

Make sure the ejector latch is open before inserting the card.

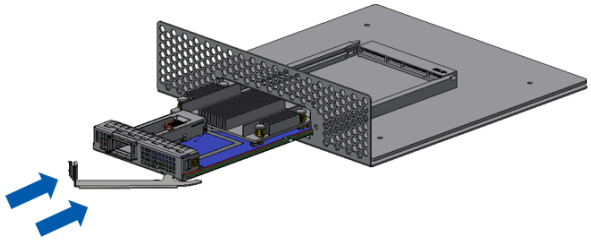

5. To secure the card, close the ejector.

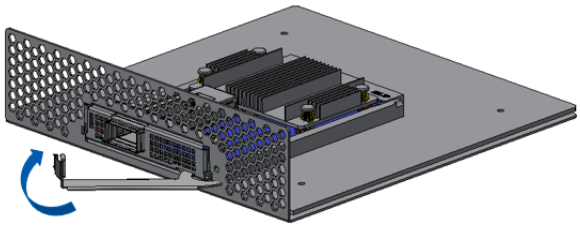

To uninstall the adapter card, see Uninstalling the Card.

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 so as not to impede the air exhaust flow through the ventilation holes. Use cable lengths that allow for routing horizontally around to the side of the chassis before bending upward or downward in the rack.
6. To remove a cable, disengage the locks and slowly pull the connector away from the port receptacle. The LED indicator will turn off when the cable is unseated.

Identifying the Card in Your System

On Linux

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

```
lspci |grep -i Mellanox
Network controller: Mellanox Technologies MT28800 Family [ConnectX-6 Dx]
```

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 Dx 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)

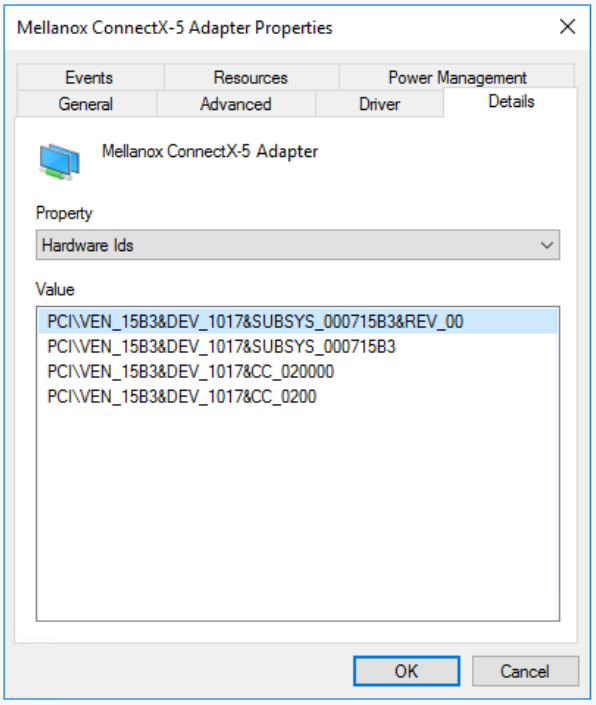

5. In the Value display box, check the fields VEN and DEV (fields are separated by '&'). In the display example above, notice the sub-string "PCI\VEN_15B3&DEV_1003": VEN is equal to 0x15B3 – this is the Vendor ID of NVIDIA; DEV is equal to 1018 (for ConnectX-6 Dx) – this is a valid NVIDIA PCI Device ID.

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

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

Card Extraction Instructions

Safety Precautions

The adapter is installed in a system that operates with voltages that can be lethal. Before uninstalling the adapter card, please observe the following precautions to avoid injury and prevent damage to system components.

1. Remove any metallic objects from your hands and wrists.
2. It is strongly recommended to use an ESD strap or other antistatic devices.
3. Turn off the system and disconnect the power cord from the server.

Cards Extraction Instructions

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

Cards with Thumbscrew (Pull-tab) Bracket

1. Verify that the system is powered off and unplugged.
2. Wait 30 seconds.
3. Rotate the captive screw counterclockwise.

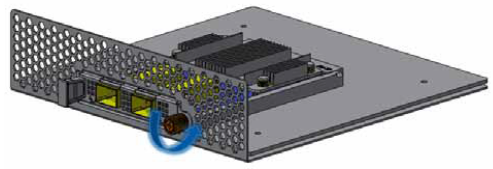

4. While holding the tab, carefully pull out the adapter card.

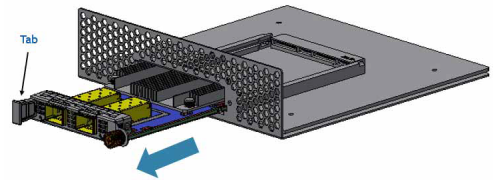

Cards with Internal-Lock Bracket Extraction

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

1. Verify that the system is powered off and unplugged.
2. Wait 30 seconds.
3. Gently extract the adapter card from the server.

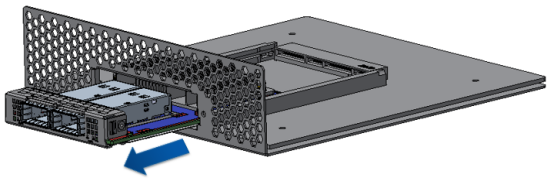

Cards with Ejector Latch Extraction

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

1. Verify that the system is powered off and unplugged.
2. Wait 30 seconds.
3. Open the ejector latch.

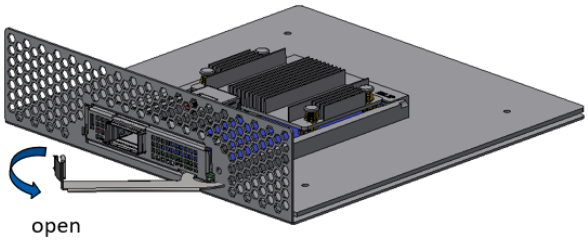

4. Gently extract the adapter card from the server.

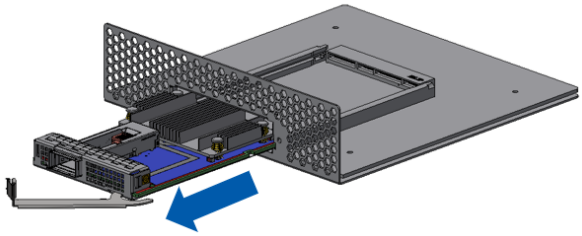

Driver Installation

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

> ⚠ 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.

Hardware and Software Requirements

Description	Package
Windows Server 2012 R2	MLNX_WinOF2-1_10_All_x64.exe
Windows Server 2012	
Windows Server 2016	
Windows 8.1 Client (64 bit only)	
Windows 10 Client (64 bit only)	

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

Downloading NVIDIA WinOF-2 Driver

To download the executable file according to your operating system, please follow the steps below:

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

```
echo %PROCESSOR_ARCHITECTURE%
```

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

2. Go to the NVIDIA WinOF-2 web page at http://www.nvidia.com > Products >InfiniBand/VPI Drivers > Windows SW/Drivers.
3. Download the exe image according to the architecture of your machine (see Step 1). The name of the .exe is in the following format: MLNX_WinOF2-<version>_x<arch>.exe.

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

Installing NVIDIA WinOF-2 Driver

This section provides instructions for two types of installation procedures:

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

⚠ Both Attended and Unattended installations require administrator privileges.

⚠ WinOF-2 supports adapter cards based on the NVIDIA ConnectX®-4 and above family of adapter IC devices only. If you have ConnectX-3 and ConnectX-3 Pro on your server, you will need to install WinOF driver. For details on how to install WinOF driver, please refer to *WinOF User Manual*.

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_WinOF2-1_10_<revision_version>_All_x64.exe /v"/l*vx [LogFile]"
```

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

```
MLNX_WinOF2-1_10_<revision_version>_All_x64.exe /v" MT_SKIPFWUPGRD=1"
```

4. Click Next in the Welcome screen.

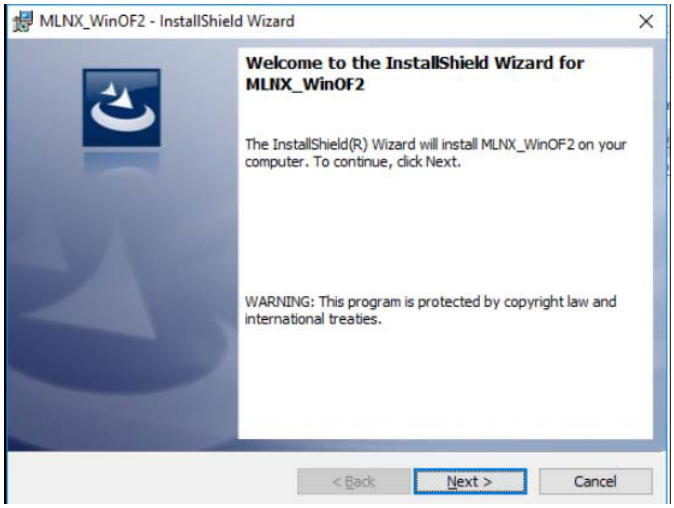

5. Read then accept the license agreement and click Next.

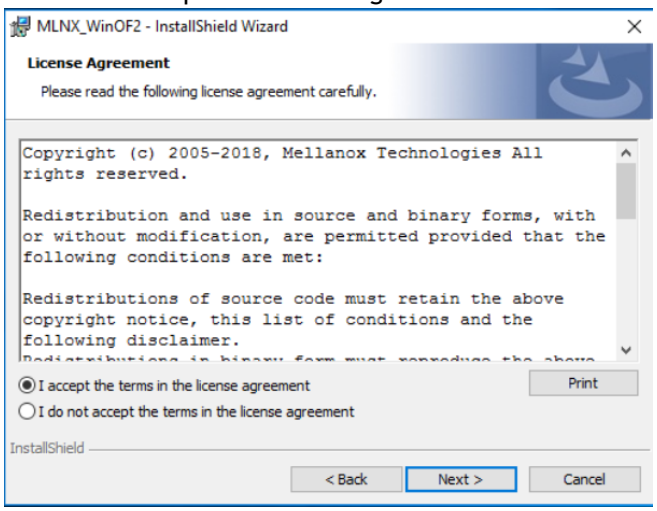

6. Select the target folder for the installation.

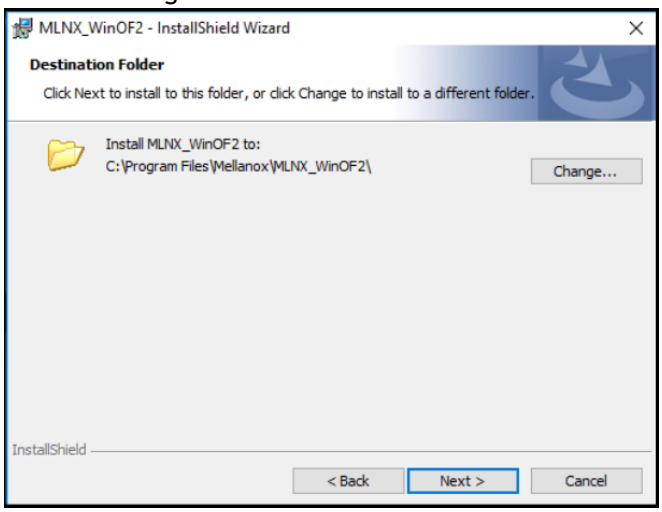

7. The firmware upgrade screen will be displayed in the following cases:
    a. If the user has an OEM card. In this case, the firmware will not be displayed.

b. If the user has a standard NVIDIA card with an older firmware version, the firmware will be updated accordingly. However, if the user has both an OEM card and a NVIDIA card, only the NVIDIA card will be updated.

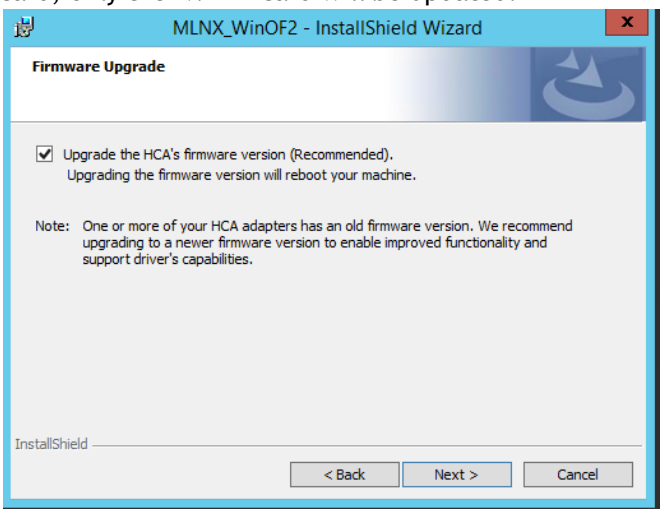

8. Select a Complete or Custom installation, follow Step a and on.

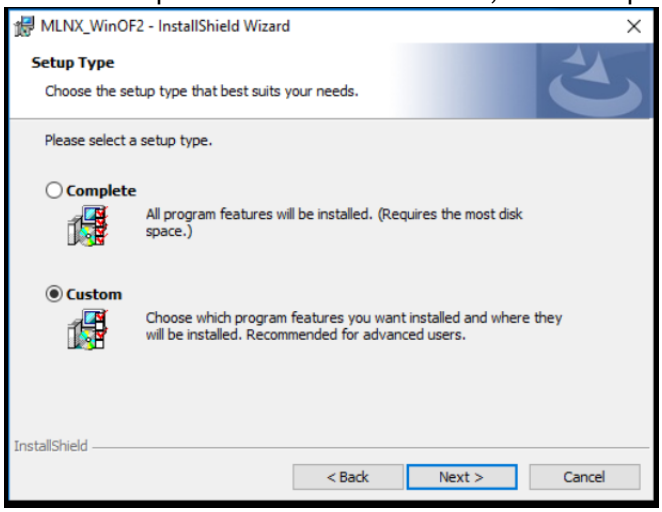

a. Select the desired feature to install:
   i. Performance tools - install the performance tools that are used to measure performance in user environment
   ii. Documentation - contains the User Manual and Release Notes
   iii. Management tools - installation tools used for management, such as mlxstat

iv. Diagnostic Tools - installation tools used for diagnostics, such as mlx5cmd

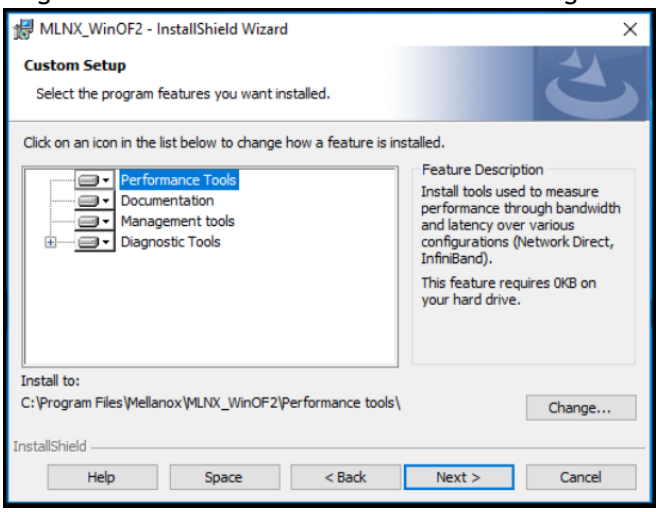

b. Click Next to install the desired tools.

9. Click Install to start the installation.

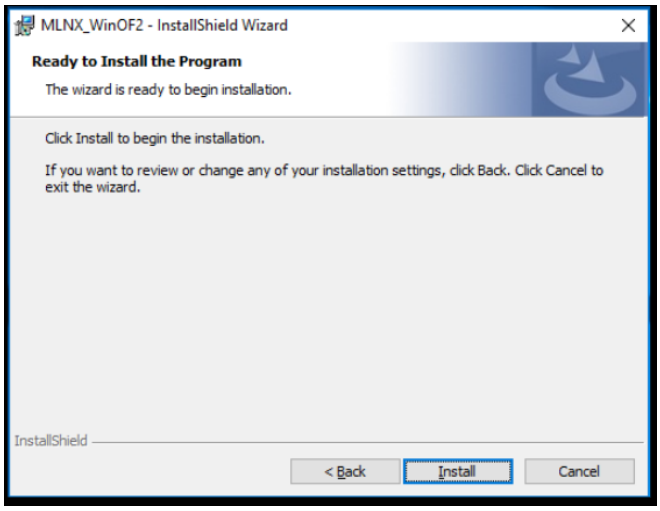

10. In case that firmware upgrade option was checked in Step 7, you will be notified if a firmware upgrade is required. See TBD.

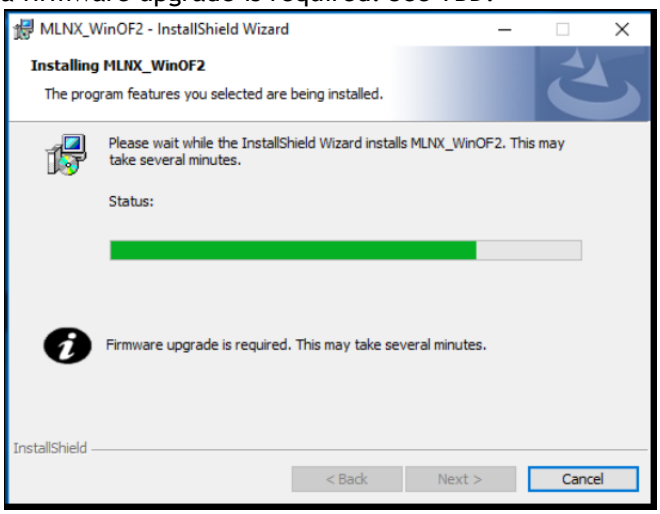

11. Click Finish to complete the installation.

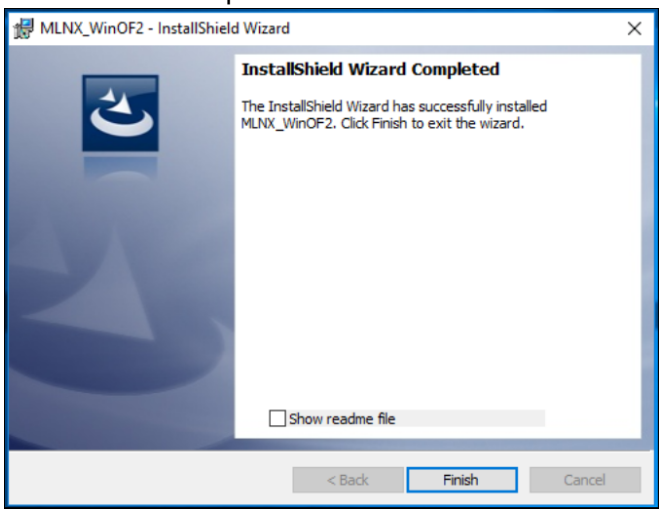

Unattended Installation

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

The following is an example of an unattended installation session.

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

```
MLNX_WinOF2-1_10_<revision_version>_All_x64.exe /S /v/qn
```

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

```
MLNX_WinOF2-1_10_All_x64.exe /S /v/qn /v"/l*vx [LogFile]"
```

4. [Optional] if you want to control whether to install ND provider or not.
(Note: MT_NDPROPERTY default value is True.)

```
MLNX_WinOF2-1_10_All_x64.exe /vMT_NDPROPERTY=1
```

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

```
MLNX_WinOF2-1_10_All_x64.exe /vMT_SKIPFWUPGRD=1
```

> Applications that hold the driver files (such as ND applications) will be closed during the unattended installation.

Installation Results

Upon installation completion, you can verify the successful addition of the network card(s) through the Device Manager.
Upon installation completion, the inf files can be located at: %ProgramFiles% \Mellanox\MLNX_WinOF2\Drivers\<OS>
To see the NVIDIA network adapters, display the Device Manager and pull down the Network adapters menu.

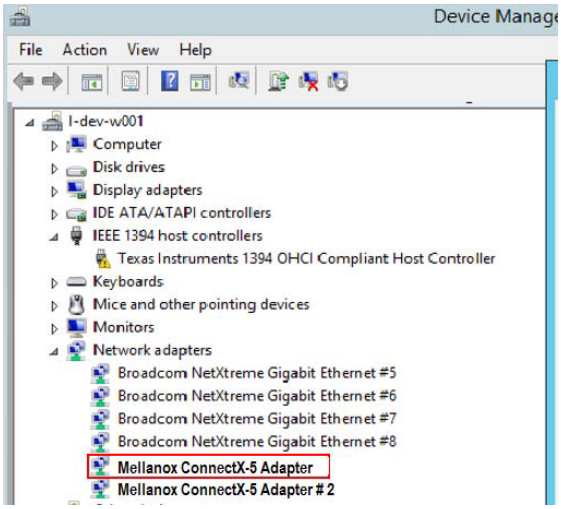

Extracting Files Without Running Installation

To extract the files without running installation, perform the following steps.

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

```
MLNX_WinOF2-1_10_All_x64 /a
```

To extract only the driver files.

```
MLNX_WinOF2-1_10_All_x64 /a /vMT_DRIVERS_ONLY=1
```

3. Click Next to create a server image.

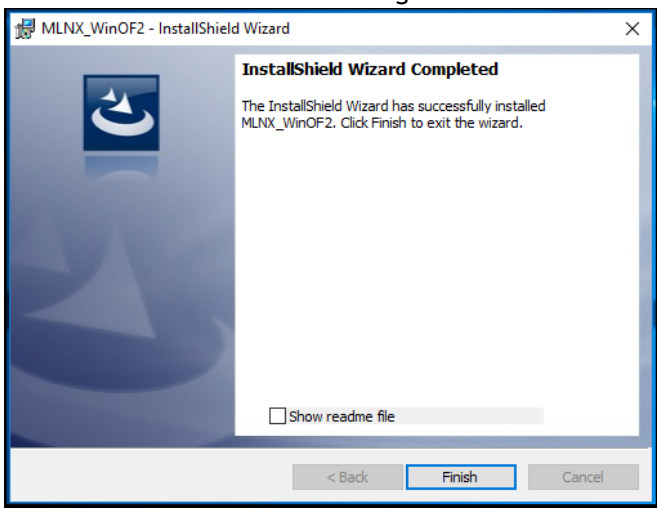

4. Click Change and specify the location in which the files are extracted to.

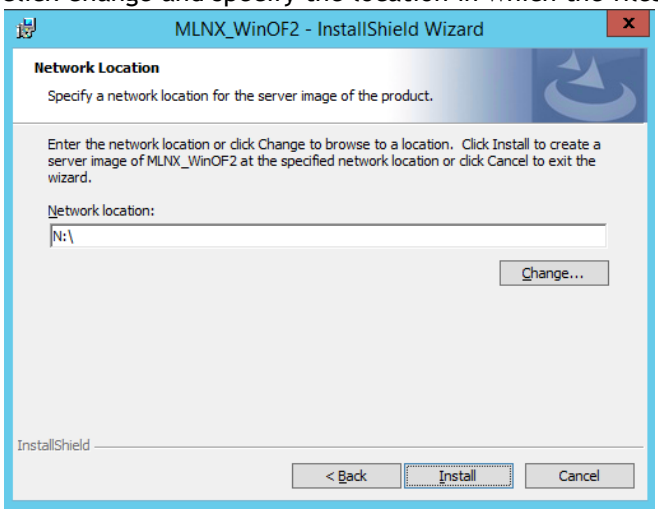

5. Click Install to extract this folder, or click Change to install to a different folder.

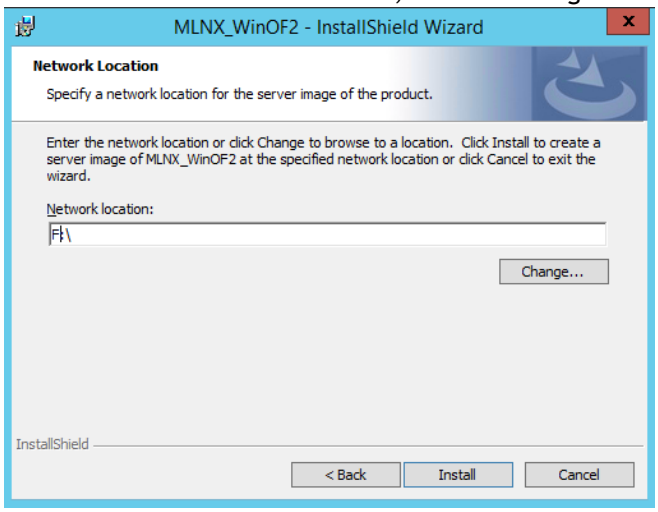

6. To complete the extraction, click Finish.

Uninstalling NVIDIA WinOF-2 Driver

Attended Uninstallation

To uninstall MLNX_WinOF2 on a single node:

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

Unattended Uninstallation

To uninstall MLNX_WinOF2 in unattended mode:

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

```
MLNX_WinOF2-1_10_All_x64.exe /S /x /v"/qn"
```

Firmware Upgrade

If the machine has a standard NVIDIA card with an older firmware version, the firmware will be automatically updated as part of the WinOF-2 package installation.
For information on how to upgrade firmware manually, please refer to the MFT User Manual at www.nvidia.com >Products > Ethernet Drivers > Firmware Tools.

Deploying the Driver on a Nano Server

Offline Installation

To deploy the Driver on a Nano Server:

1. Go to the NVIDIA WinOF web page at http://www.nvidia.com > Products > Ethernet Drivers > Windows SW/Drivers.
2. Download the driver (MLNX_WinOF2_MLNX_WinOF2-1_64_mlx5_All_win2016_x64_fre_1_64_15407.exe).
3. Extract the driver to a local directory (see Extracting Files Without Running Installation).
4. Copy the contents of this directory to C:\WS2016TP5_Drivers.
   The directory contents should appear as follows:

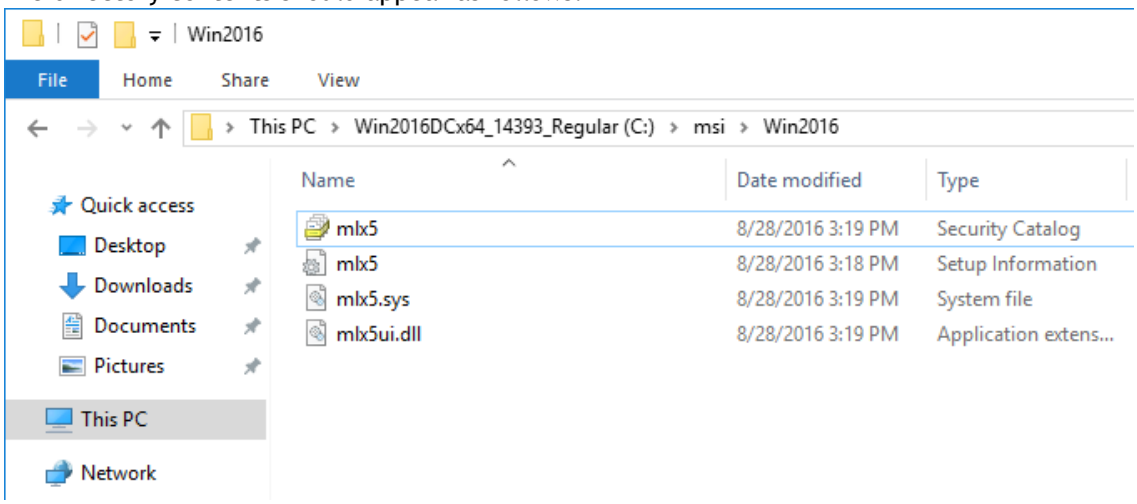

   This location should be specified for DriversPath property when injecting driver into the Nano server image:

```
New-NanoServerImage -MediaPath \\Path\To\Media\en_us -BasePath .\Base -TargetPath .\InjectingDrivers.vhdx
-DriversPath C:\WS2016TP5_Drivers
```

5. Create the Nano server image. Follow the instructions in the TechNet article for "Getting Started with Nano Server".

Online Update

1. Go to the NVIDIA WinOF web page at http://www.nvidia.com > Products >Ethernet Drivers > Windows SW/Drivers.
2. Download the driver package.
3. Extract the NVIDIA drivers to a local directory (see Extracting Files Without Running Installation).
4. Open a remote connection to the Nano server.
5. Copy all the driver files to the Nano server
6. Install the driver:

```
pnputil -I -a <Driver_Folder>\mlx5.inf
```

When upgrading the driver on a server where the remote connection was done over NVIDIA device, there will be a loss of connectivity. To upgrade, it is recommended to run a script that will execute all the required upgrade commands.

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 Dx adapter card installed.

Prerequisites

Requirements	Description
Platforms	A server platform with ConnectX®-6 Dx (EN) (firmware: fw-ConnectX6 Dx)
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.

Downloading NVIDIA OFED

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

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

2. Download the ISO image to your host.
   The image's name has the format `MLNX_OFED_LINUX-<ver>-<OS label><CPU arch>.iso`.
   You can download and install the latest OpenFabrics Enterprise Distribution (OFED) software package available via the NVIDIA website at http://www.nvidia.com > Products > Software > Ethernet Drivers > Linux SW/Drivers > Download..
   a. Scroll down to the Download wizard, and click the Download tab.

b. Choose your relevant package depending on your host operating system.
c. Click the desired ISO/tgz package.
d. To obtain the download link, accept the End User License Agreement (EULA).

3. Use the md5sum utility to confirm the file integrity of your ISO image. Run the following command and compare the result to the value provided on the download page.

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

Installing NVIDIA OFED

Installation Script

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

- Discovers the currently installed kernel
- Uninstalls any software stacks that are part of the standard operating system distribution or another vendor's commercial stack
- Installs the MLNX_OFED_LINUX binary RPMs (if they are available for the current kernel)
- Identifies the currently installed NVIDIA Ethernet network adapters and automatically upgrades the firmware.
  Note: The firmware will not be updated if you run the install script with the '--without-fw-update' option.
  Note: If you wish to perform a firmware upgrade using customized FW binaries, you can provide a path to the folder that contains the FW binary files, by running `--fw-image-dir`. Using this option, the FW version embedded in the MLNX_OFED package will be ignored. Example:

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

Usage

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

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

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

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

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

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

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

Example
The following command will create a MLNX_OFED_LINUX ISO image for RedHat 6.3 under the /tmp directory.

```
# ./MLNX_OFED_LINUX-x.x-x-rhel6.3-x86_64/mlnx_add_kernel_support.sh -m /tmp/MLNX_OFED_LINUX-x.x-x-rhel6.3-
x86_64/ --make-tgz
Note: This program will create MLNX_OFED_LINUX TGZ for rhel6.3 under /tmp directory.
All Mellanox, OEM, OFED, or Distribution IB packages will be removed.
Do you want to continue?[y/N]:y
See log file /tmp/mlnx_ofed_iso.21642.log

Building OFED RPMs. Please wait...
Removing OFED RPMs...
Created /tmp/MLNX_OFED_LINUX-x.x-x-rhel6.3-x86_64-ext.tgz
```

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

For your machine to be part of the Ethernet fabric, a Subnet Manager must be running on one of the fabric nodes. At this point, NVIDIA OFED for Linux has already installed the OpenSM Subnet Manager on your machine.
For the list of installation options, run:

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

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

Installation Procedure

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

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

3. Run the installation script.

```
/mnt/mlnxofedinstall
Logs dir: /tmp/MLNX_OFED_LINUX-x.x-x.logs
This program will install the MLNX_OFED_LINUX package on your machine.
Note that all other Mellanox, OEM, OFED, RDMA or Distribution IB packages will be removed.
Those packages are removed due to conflicts with MLNX_OFED_LINUX, do not reinstall them.
Starting MLNX_OFED_LINUX-x.x.x installation ...
........
........
Installation finished successfully.

Attempting to perform Firmware update...
Querying Mellanox devices firmware ...
```

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

```
/mnt/mlnxofedinstall --force
```

MLNX_OFED for Ubuntu should be installed with the following flags in the root 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 the installation a message similar to the following:

```
Device #1:
----------

Device Type:      ConnectX-6 Dx
Part Number:      MCX623436AC-CDAB
Description:      ConnectX®-6 Dx EN adapter card, 100GbE, OCP3.0, With Host Management,
Dual-port QSFP56, PCIe 4.0 x16, Crypto and Secure Boot, Thumbscrew (Pull Tab) Bracket


PSID: MT_2190110032
PCI Device Name: 0b:00.0
Base MAC:         0000e41d2d5cf810
Versions:         Current        Available
  FW              16.22.0228     16.22.0228
  Status:            Up to date
```

In case your machine has an unsupported network adapter device, no firmware update will occur and one of the following error messages below will be printed. Please contact your hardware vendor for help on firmware updates.
Error message 1:

```
Device #1:
----------

Device Type:      ConnectX-6 Dx
Part Number:      MCX623436AC-CDAB
Description:      ConnectX®-6 Dx EN adapter card, 100GbE, OCP3.0, With Host Management,
Dual-port QSFP56, PCIe 4.0 x16, Crypto and Secure Boot, Thumbscrew (Pull Tab) Bracket

PSID:                  MT_2190110032
PCI Device Name:  0b:00.0
Base MAC:         0000e41d2d5cf810
Versions:         Current           Available
                  FW 16.22.0228     N/A
Status:           No matching image found
```

Error message 2:

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

4. If the installation script has performed a firmware update on your network adapter, complete the step relevant to your adapter card type to load the firmware:
   - ConnectX-6 Dx/ConnectX-6 Dx Ex - perform a standard reboot

   Otherwise, restart the driver by running: `/etc/init.d/openibd restart`

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

Most of the NVIDIA OFED components can be configured or reconfigured after the installation, by modifying the relevant configuration files. See the relevant chapters in this manual for details.

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

Installation Results

Software	• 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
	• 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 Red Hat like Distributions

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 NVIDIA OFED ISO image Note: If an adapter's flash was originally programmed with an Expansion ROM image, the automatic firmware update will also burn an Expansion ROM image.
	• In case that your machine has an unsupported network adapter device, no firmware update will occur and the error message below will be printed.
	`The firmware for this device is not distributed inside Mellanox driver: 0000:01:00.0 (PSID: IBM2150110033)`
	`To obtain firmware for this device, please contact your HW vendor.`

Installation Logs

While installing MLNX_OFED, the install log for each selected package will be saved in a separate log file. The path to the directory containing the log files will be displayed after running the installation script in the following format: "`Logs dir: /tmp/MLNX_OFED_LINUX-<version>.<PD>.logs`".

Example:

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

openibd Script

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

```
OPENIBD_PRE_START
OPENIBD_POST_START
OPENIBD_PRE_STOP
OPENIBD_POST_STOP
```

Example:

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

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

Driver Load Upon System Boot

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

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

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

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

2. Set “ONBOOT=no” in the "/etc/infiniband/openib.conf" file.
3. If the modules exist in the initramfs file, they can automatically be loaded by the kernel. To prevent this behavior, update the initramfs using the operating systems’ standard tools. Note: The process of updating the initramfs will add the blacklists from step 1, and will prevent the kernel from loading the modules automatically.

mlnxofedinstall Return Codes

The table below lists the mlnxofedinstall script return codes and their meanings.

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

Uninstalling MLNX_OFED

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

Installing MLNX_OFED Using YUM

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

Setting up MLNX_OFED YUM Repository

1. Log into the installation machine as root.
2. Mount the ISO image on your machine and copy its content to a shared location in your network.

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

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

```
# wget http://www.nvidia.com/downloads/ofed/RPM-GPG-KEY-Mellanox
--2014-04-20 13:52:30-- http://www.nvidia.com/downloads/ofed/RPM-GPG-KEY-Mellanox
Resolving www.nvidia.com... 72.3.194.0
Connecting to www.nvidia.com|72.3.194.0|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 1354 (1.3K) [text/plain]
Saving to: ?RPM-GPG-KEY-Mellanox?

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

2014-04-20 13:52:30 (247 MB/s) - ?RPM-GPG-KEY-Mellanox? saved [1354/1354]
```

4. Install the key.

```
# sudo rpm --import RPM-GPG-KEY-Mellanox
warning: rpmts_HdrFromFdno: Header V3 DSA/SHA1 Signature, key ID 6224c050: NOKEY
Retrieving key from file:///repos/MLNX_OFED/<MLNX_OFED file>/RPM-GPG-KEY-Mellanox
Importing GPG key 0x6224C050:
Userid: "Mellanox Technologies (Mellanox Technologies - Signing Key v2) <support@mellanox.com>"
From : /repos/MLNX_OFED/<MLNX_OFED file>/RPM-GPG-KEY-Mellanox
Is this ok [y/N]:
```

5. Check that the key was successfully imported.

```
# rpm -q gpg-pubkey --qf '%{NAME}-%{VERSION}-%{RELEASE}\t%{SUMMARY}\n' | grep Mellanox
gpg-pubkey-a9e4b643-520791ba gpg(Mellanox Technologies <support@mellanox.com>)
```

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

```
[mlnx_ofed]
name=MLNX_OFED Repository
baseurl=file:///<path to extracted MLNX_OFED package>/RPMS
enabled=1
gpgkey=file:///<path to the downloaded key RPM-GPG-KEY-Mellanox>
gpgcheck=1
```

7. Check that the repository was successfully added.

```
# yum repolist
Loaded plugins: product-id, security, subscription-manager
This system is not registered to Red Hat Subscription Management. You can use subscription-manager to
register.
repo id      repo name                                  status
mlnx_ofed    MLNX_OFED Repository                        108
rpmforge     RHEL 6Server - RPMforge.net - dag          4,597

repolist: 8,351
```

Installing MLNX_OFED Using the YUM Tool

After setting up the YUM repository for MLNX_OFED package, perform the following:

1. View the available package groups by invoking:

```
# yum search mlnx-ofed-
mlnx-ofed-all.noarch : MLNX_OFED all installer package (with KMP support)
mlnx-ofed-basic.noarch : MLNX_OFED basic installer package (with KMP support)
mlnx-ofed-guest.noarch : MLNX_OFED guest installer package (with KMP support)
mlnx-ofed-hpc.noarch : MLNX_OFED hpc installer package (with KMP support)
mlnx-ofed-hypervisor.noarch : MLNX_OFED hypervisor installer package (with KMP support)
mlnx-ofed-vma.noarch : MLNX_OFED vma installer package (with KMP support)
mlnx-ofed-vma-eth.noarch : MLNX_OFED vma-eth installer package (with KMP support)
```

```
mlnx-ofed-vma-vpi.noarch : MLNX_OFED vma-vpi installer package (with KMP support)
```

Where:

```
mlnx-ofed-all               Installs all available packages in MLNX_OFED.
mlnx-ofed-basic             Installs basic packages required for running Mellanox cards.
mlnx-ofed-guest             Installs packages required by guest OS.
mlnx-ofed-hpc               Installs packages required for HPC.
mlnx-ofed-hypervisor        Installs packages required by hypervisor OS.
mlnx-ofed-vma               Installs packages required by VMA.
mlnx-ofed-vma-eth           Installs packages required by VMA to work over Ethernet.
mlnx-ofed-vma-vpi           Installs packages required by VMA to support VPI.
```

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

Example:

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

If you have an operating system different than RHEL or SLES, or you have installed a kernel that is not supported by default in MLNX_OFED, you can use the mlnx_add_kernel_support.sh script to build MLNX_OFED for your kernel.
The script will automatically build the matching group RPM packages for your kernel so that you can still install MLNX_OFED via yum.
Please note that the resulting MLNX_OFED repository will contain unsigned RPMs, therefore, you should set 'gpgcheck=0' in the repository configuration file.

2. Install the desired group.

```
# yum install mlnx-ofed-all
Loaded plugins: langpacks, product-id, subscription-manager
Resolving Dependencies
--> Running transaction check
---> Package mlnx-ofed-all.noarch 0:3.1-0.1.2 will be installed
--> Processing Dependency: kmod-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!
```

Uninstalling MLNX_OFED Using the YUM Tool

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

Installing MLNX_OFED Using apt-get Tool

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

Setting up MLNX_OFED apt-get Repository

1. Log into the installation machine as root.
2. Extract the MLNX_OFED package on a shared location in your network.
   You can download it from http://www.nvidia.com > Products > Software> Ethernet Drivers.
3. Create an apt-get repository configuration file called "/etc/apt/sources.list.d/mlnx_ofed.list" with the following content:

```
# deb file:/<path to extracted MLNX_OFED package>/DEBS ./
```

4. Download and install NVIDIA Technologies GPG-KEY.

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

5. Check that the key was successfully imported.

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

6. Update the apt-get cache.

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

Installing MLNX_OFED Using the apt-get Tool

After setting up the apt-get repository for MLNX_OFED package, perform the following:

1. View the available package groups by invoking:

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

   Where:

```
mlnx-ofed-all        MLNX_OFED all installer package.
mlnx-ofed-basic      MLNX_OFED basic installer package.
mlnx-ofed-vma        MLNX_OFED vma installer package.
mlnx-ofed-hpc        MLNX_OFED HPC installer package.
mlnx-ofed-vma-eth    MLNX_OFED vma-eth installer package.
mlnx-ofed-vma-vpi    MLNX_OFED vma-vpi installer package.
```

2. Install the desired group.

```
# apt-get install '<group name>'
```

Example:

```
# apt-get install mlnx-ofed-all
```

> Installing MLNX_OFED using the “apt-get” tool does not automatically update the firmware. To update the firmware to the version included in MLNX_OFED package, run: `# apt-get install mlnx-fw-updater`
> Or, update the firmware to the latest version that is described in Updating Adapter Firmware.

Uninstalling MLNX_OFED Using the apt-get Tool

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

Updating Firmware After Installation

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

Updating the Device Online

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

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

Example:

```
mlxfwmanager --online -u -d 0000:09:00.0
Querying Mellanox devices firmware ...
Device #1:
----------
Device Type:      ConnectX-6 Dx
Part Number:      MCX623436AC-CDAB
Description:      ConnectX®-6 Dx EN adapter card, 100GbE, OCP3.0, With Host Management, Dual-port QSFP56, PCIe 4.0
x16, Crypto and Secure Boot, Thumbscrew (Pull Tab) Bracket
PSID:             MT_1020120019
PCI Device Name:  0000:09:00.0
Port1 GUID:       0002c9000100d051
Port2 MAC:        0002c9000002
Versions:         Current              Available
                  FW 2.32.5000         2.33.5000
Status:           Update required
---------
Found 1 device(s) requiring firmware update. Please use -u flag to perform the update.
```

Updating the Device Manually

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

In case that you ran the mlnxofedinstall script with the ‘--without-fw-update’ option or you are using an OEM card and now you wish to (manually) update firmware on your adapter card(s), you

need to perform the steps below. The following steps are also appropriate in case that you wish to burn newer firmware that you have downloaded from NVIDIA's Web site (http://www.nvidia.com > Support > Firmware Download).

1. Get the device's PSID.

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

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

3. Burn the firmware.

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

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

Updating the Device Firmware Automatically upon System Boot

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

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

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

Otherwise, the following message is printed:

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

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

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

```
MLNX_EXCLUDE_DEVICES="00:05.0,00:07.0"
```

UEFI Secure Boot

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

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

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

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

1. Download the x.509 public key.

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

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

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

3. Reboot the system.

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

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

Removing Signature from kernel Modules

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

To remove the signature from the MLNX_OFED kernel modules:

1. Remove the signature.

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

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

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

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

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

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

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

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

Performance Tuning

Depending on the application of the user's system, it may be necessary to modify the default configuration of network adapters based on the ConnectX® adapters. In case that tuning is required, please refer to the Performance Tuning Guide for NVIDIA Network Adapters at https://community.nvidia.com/docs/DOC-2489.

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's devices: • ConnectX®-6 Dx (EN) (firmware: fw-ConnectX6 Dx)
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 NVIDIA NATIVE ESXi Driver for VMware vSphere

⚠ Please uninstall all previous NVIDIA driver packages prior to installing the new version. See Removing Earlier NVIDIA Drivers for further information.

To install the driver:

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

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

Example:

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

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 NVIDIA NATIVE ESXi Driver for VMware vSphere.)
3. Remove each module:

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

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

4. Reboot the server.

Firmware Programming

a. Download the VMware bootable binary images from the NVIDIA Firmware Tools (MFT) site.
b. Install the image according to the steps described in the MFT User Manual.

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

Updating Adapter Firmware

Each adapter card is shipped with the latest version of qualified firmware at the time of manufacturing. However, NVIDIA issues firmware updates occasionally that provide new features and bug fixes. To check that your card is programmed with the latest available firmware version, download the mlxup firmware update and query utility. The utility can query for available NVIDIA adapters and indicate which adapters require a firmware update. If the user confirms, mlxup upgrades the firmware using embedded images. The latest mlxup executable and documentation are available from http://www.nvidia.com > Products > Software > Firmware Tools.

Firmware Update Example

```
[server1]# ./mlxup
Querying Mellanox devices firmware ...
Device Type:      ConnectX-6 Dx
Part Number:      MCX623436AC-CDAB
Description:      ConnectX®-6 Dx EN adapter card, 100GbE, OCP3.0, With Host Management, Dual-port QSFP56, PCIe 4.0
x16, Crypto and Secure Boot, 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 Dx
Part Number:      MCX623436AC-CDAB
Description:      ConnectX®-6 Dx EN adapter card, 100GbE, OCP3.0, With Host Management, Dual-port QSFP56, PCIe 4.0
x16, Crypto and Secure Boot, 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

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

NVIDIA Firmware Tool (NFT)	Download and install MFT: http://www.nvidia.com/content/pages.php?pg=management_tools&menu_section=34 Refer to the User Manual for installation instructions. 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 http://www.nvidia.com/supportdownloader
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.
NVIDIA Firmware Tool (NFT)	Download and install MFT: http://www.nvidia.com/content/pages.php?pg=management_tools&menu_section=34 Refer to the User Manual for installation instructions. 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: http://www.nvidia.com/supportdownloader/ 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

⚠ Install the ConnectX-6 Dx 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.

MCX623432AE-ADAB / MCX623432AN-ADAB / MCX623432AC-ADAB / MCX623432AS-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, 16 lanes (2.0 and 1.1 compatible)</td></tr>
<tr><td rowspan="8">Power and Environmental</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">15.7W</td><td colspan="2">3.8W</td></tr>
<tr><td>Maximum Power</td><td>Passive Cables</td><td colspan="2">19.74W</td><td colspan="2">7.95W</td></tr>
<tr><td colspan="6">Maximum power available through SFP28 port: 1.5W</td></tr>
<tr><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>Humidity</td><td colspan="5">90% relative humidity[c]</td></tr>
<tr><td>Altitude (Operational)</td><td colspan="5">3050m</td></tr>
<tr><td rowspan="2">Environmental</td><td rowspan="2">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>
</table>

		Passive Cable	400LFM	100LFM	TBD	TBD
Regulatory	Safety	CB / cTUVus / CE				
	EMC	CE / FCC / VCCI / ICES / RCM				
	RoHS	RoHS compliant				

a. ConnectX-6 Dx 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.

MCX623432AC-GDAB / MCX623432AN-GDAB / MCX623432AS-GDAI 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: Dual SFP56 Ethernet (copper and optical)			
	Retention Mechanism:	**MCX623432A[N/C]-GDAB:** Thumbscrew (Pull Tab) Bracket		
		MCX623432AS-GDAI: Internal Lock		
Protocol Support	**Ethernet**: 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR			
	Data Rate	Ethernet	1/10/25/40/50 Gb/s	
	PCI Express Gen 4.0: SERDES @ 16.0GT/s, 16 lanes (2.0 and 1.1 compatible)			
Power and Environmental	**Voltage**: 3.3VAUX, 12V			
	Power	**Cable Type**	**Active Mode**	**Standby Mode**
	Typical Power[(b)]	Passive Cables	16.9W	4.54W
	Maximum Power	Passive Cables	21.4W	8.75W
	Maximum power available through SFP56 port: 2W			
	Temperature	Operational	0°C to 55°C	
		Non-operational	-40°C to 70°C	
	Humidity	90% relative humidity[(c)]		

	Altitude (Operational)	3050m				
Environmental	**Airflow**	**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	400LFM	100LFM	TBD	TBD
Regulatory	**Safety**	CB / cTUVus / CE				
	EMC	CE / FCC / VCCI / ICES / RCM				
	RoHS	RoHS compliant				

a. ConnectX-6 Dx 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.

MCX623435MN-CDAB 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 QSFP56 Ethernet (copper and optical)			
	Retention Mechanism: Thumbscrew (Pull Tab) Bracket			
Protocol Support	**Ethernet**: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 100GBASE-CR2, 100GBASE-KR2, 100GBASE-SR2, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR			
	Data Rate:	Ethernet	1/10/25/40/50/100 Gb/s	
	PCI Express Gen 4.0: SERDES @ 16.0GT/s, 16 lanes (2.0 and 1.1 compatible)			
Power and Environmental	**Voltage**: 3.3VAUX, 12V			
	Power	**Cable Type**	**Active Mode**	**Standby Mode**
	Typical Power[b]	Passive Cables	18W	5.7W
	Maximum Power	Passive Cables	23.4W	10W
	Maximum power available through QSFP56 port: 5W			
	Temperature	Operational	0°C to 55°C	
		Non-operational	-40°C to 70°C	
	Humidity	90% relative humidity[c]		

	Altitude (Operational)	3050m				
Environmental	**Airflow**	**Cable Type**	**Hot Aisle HSK to Port**		**Cold Aisle @35C Port to HSK**	
			Active Mode @55C	**Standby Mode @50C**	**Active Mode**	**Standby Mode**
		Passive Cable	500LFM	100LFM	300LFM	100LFM
Regulatory	**Safety**	CB / cTUVus / CE				
	EMC	CE / FCC / VCCI / ICES / RCM				
	RoHS	RoHS compliant				

a. ConnectX-6 Dx 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.

MCX623439MC-CDAB / MCX623439MN-CDAB 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-port QSFP56		
	Retention Mechanism:	Thumbscrew (Pull Tab) Bracket	
Protocol Support	**Ethernet**: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 100GBASE-CR2, 100GBASE-KR2, 100GBASE-SR2, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR		
	Data Rate:	Ethernet	1/10/25/40/50/100 Gb/s
	PCI Express Gen 4.0: SERDES @ 16.0GT/s, 16 lanes (2.0 and 1.1 compatible)		

Power and Environmental	**Voltage**: 3.3VAUX, 12V				
	Power	**Cable Type**	**Active Mode**	**Standby Mode**	
	Typical Power[b]	Passive Cables	17.22W	4.3W	
	Maximum Power	Passive Cables	22.77W	8.1W	
	Maximum power available through QSFP56 port: 5W				
	Temperature	Operational	0°C to 55°C		
		Non-operational	-40°C to 70°C		
	Humidity	90% relative humidity[c]			
	Altitude (Operational)	3050m			
Environmental	**Airflow**	**Cable Type**	**Hot Aisle HSK to Port**		**Cold Aisle @35C Port to HSK**
			Normal Mode @55C	**Standby Mode @45C**	**Normal Mode**	**Standby Mode**
		Passive Cable	600LFM	100LFM	Not Supported	Not Supported
Regulatory	**Safety**	CB / cTUVus / CE			
	EMC	CE / FCC / VCCI / ICES / RCM			
	RoHS	RoHS compliant			

MCX623436AC-CDAB / MCX623436AN-CDAB / MCX623436AS-CDAB / MCX623436MC-CDAB / MCX623436MN-CDAB / MCX623436MS-CDAB / MCX623436AS-CDAI 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: Dual QSFP56 Ethernet (copper and optical)	
	Retention Mechanism:	MCX623436A[N/C]-CDAB: Thumbscrew (Pull Tab) Bracket
		MCX623436AS-CDAI: Internal Lock Bracket
Protocol Support	**Ethernet**: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 100GBASE-CR2, 100GBASE-KR2, 100GBASE-SR2, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR	

	Data Rate:	Ethernet	1/10/25/40/50/100 Gb/s		
	PCI Express Gen 4.0: SERDES @ 16.0GT/s, 16 lanes (2.0 and 1.1 compatible)				
Power and Environmental	**Voltage**: 3.3VAUX, 12V				
	Power	**Cable Type**	**Active Mode**	**Standby Mode**	
	Typical Power[b]	Passive Cables	19W	5.7W	
	Maximum Power	Passive Cables	26.6W	10W	
	Maximum power available through QSFP56 port: 5W				
	Temperature	Operational	0°C to 55°C		
		Non-operational	-40°C to 70°C		
	Humidity	90% relative humidity[c]			
	Altitude (Operational)	3050m			
Environmental	**Airflow**	**Cable Type**	**Hot Aisle HSK to Port**		**Cold Aisle @35C Port to HSK**	
			Normal Mode @55C	**Standby Mode @45C**	**Normal Mode**	**Standby Mode**
		Passive Cable	600LFM	100LFM	Not Supported	Not Supported
Regulatory	**Safety**	CB / cTUVus / CE				
	EMC	CE / FCC / VCCI / ICES / RCM				
	RoHS	RoHS compliant				

a. ConnectX-6 Dx 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.

MCX623430MS-CDAB 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: Dual DSFP Ethernet (copper and optical)	
	Retention Mechanism:	Thumbscrew (Pull Tab) Bracket
Protocol Support	**Ethernet**: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 100GBASE-CR2, 100GBASE-KR2, 100GBASE-SR2, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR	

	Data Rate:	Ethernet	1/10/25/40/50/100 Gb/s		
	PCI Express Gen 4.0: SERDES @ 16.0GT/s, 16 lanes (2.0 and 1.1 compatible)				
Power and Environmental	**Voltage**: 3.3VAUX, 12V				
	Power	**Cable Type**	**Active Mode**	**Standby Mode**	
	Typical Power[b]	Passive Cables	19W	5.7W	
	Maximum Power	Passive Cables	26.6W	10W	
	Maximum power available through DSFP port: 5W				
	Temperature	Operational	0°C to 55°C		
		Non-operational	-40°C to 70°C		
	Humidity	90% relative humidity[c]			
	Altitude (Operational)	3050m			

Environmental	Airflow	Cable Type	Hot Aisle HSK to Port		Cold Aisle @35C Port to HSK	
			Normal Mode @55C	Standby Mode @45C	Normal Mode	Standby Mode
		Passive Cable	600LFM	100LFM	300LFM	100LFM
Regulatory	**Safety**	CB / cTUVus / CE				
	EMC	CE / FCC / VCCI / ICES / RCM				
	RoHS	RoHS compliant				

a. ConnectX-6 Dx 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.

MCX623435AC-CDAB / MCX623435AE-CDAB / MCX623435AN-CDAB / MCX623435AC-VDAB / MCX623435AE-VDAB / MCX623435AN-VDAB / MCX623435AS-VDAI 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 QSFP56 Ethernet (copper and optical)	
	Retention Mechanism:	MCX623435A[N/C]-VDAB: Thumbscrew (Pull Tab) Bracket
		MCX623435AS-VDAI: Internal Lock

<table>
<tr><td rowspan="4">Protocol Support</td><td colspan="5">Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 100GBASE-CR2, 100GBASE-KR2, 100GBASE-SR2, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR</td></tr>
<tr><td>Data Rate:</td><td>Ethernet</td><td colspan="3">1/10/25/40/50/100/200 Gb/s</td></tr>
<tr><td colspan="5">PCI Express Gen 4.0: SERDES @ 16.0GT/s, 16 lanes (2.0 and 1.1 compatible)</td></tr>
<tr><td rowspan="8">Power and Environmental</td><td colspan="5">Voltage: 3.3VAUX, 12V</td></tr>
<tr><td>Power</td><td>Cable Type</td><td colspan="2">Active Mode</td><td>Standby Mode</td></tr>
<tr><td>Typical Power[b]</td><td>Passive Cables</td><td colspan="2">16.9W</td><td>5.6W</td></tr>
<tr><td>Maximum Power</td><td>Passive Cables</td><td colspan="2">24W</td><td>11.4W</td></tr>
<tr><td colspan="5">Maximum power available through QSFP56 port: 5W</td></tr>
<tr><td rowspan="2">Temperature</td><td>Operational</td><td colspan="3">0°C to 55°C</td></tr>
<tr><td>Non-operational</td><td colspan="3">-40°C to 70°C</td></tr>
<tr><td>Humidity</td><td colspan="4">90% relative humidity[c]</td></tr>
<tr><td>Altitude (Operational)</td><td colspan="4">3050m</td></tr>
<tr><td rowspan="3">Environmental</td><td rowspan="3">Airflow</td><td rowspan="2">Cable Type</td><td colspan="2">Hot Aisle
Heatsink to Port</td><td colspan="2">Cold Aisle @35C
Port to Heatsink</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>500LFM</td><td>100LFM</td><td>300LFM</td><td>100LFM</td></tr>
<tr><td rowspan="3">Regulatory</td><td>Safety</td><td colspan="4">CB / cTUVus / CE</td></tr>
<tr><td>EMC</td><td colspan="4">CE / FCC / VCCI / ICES / RCM</td></tr>
<tr><td>RoHS</td><td colspan="4">RoHS compliant</td></tr>
</table>

a. ConnectX-6 Dx 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.

Board Mechanical Drawing and Dimensions

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

ConnectX-6 Dx SFP28/**SFP56** OCP 3.0 Mechanical Drawing and Dimensions

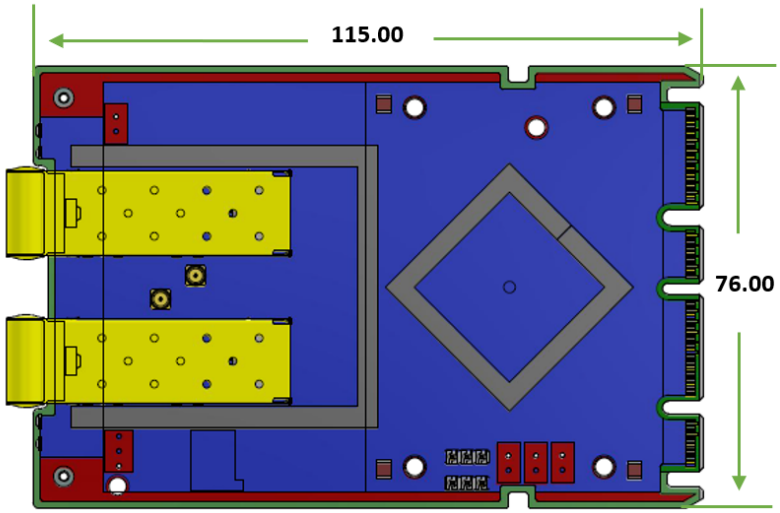

ConnectX-6 Dx Dual-Port QSFP56 OCP 3.0 Mechanical Drawing and Dimensions

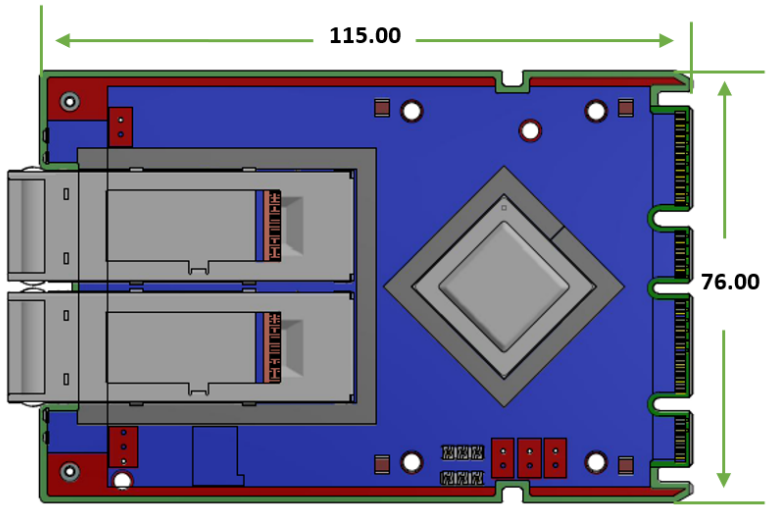

ConnectX-6 Dx Single-Port QSFP56 OCP 3.0 Mechanical Drawing and Dimensions

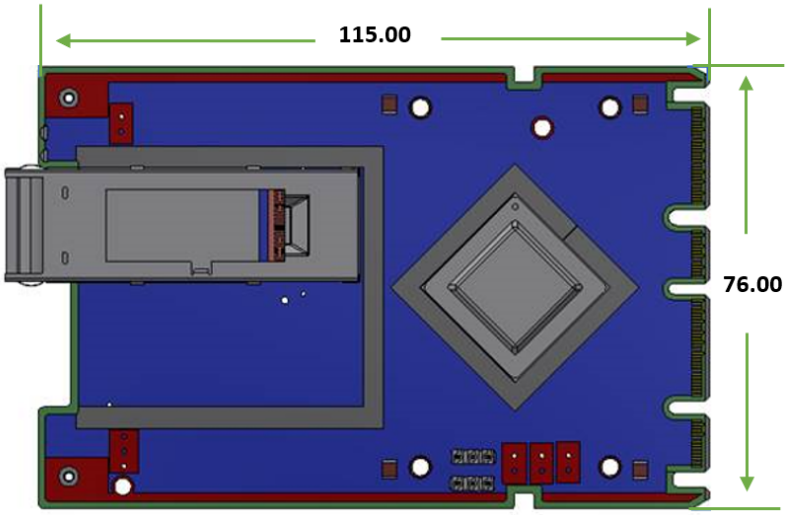

Bracket Mechanical Drawings and Dimensions

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

Dual-port SFP28/SFP56 Thumbscrew Bracket

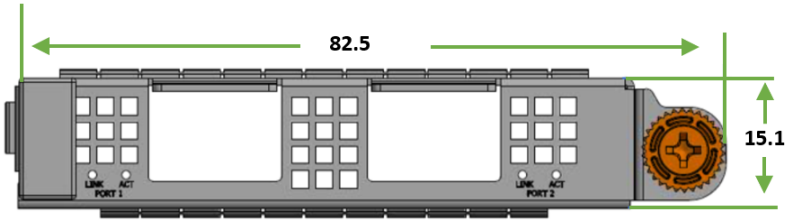

Dual-port QSFP56 Thumbscrew Bracket

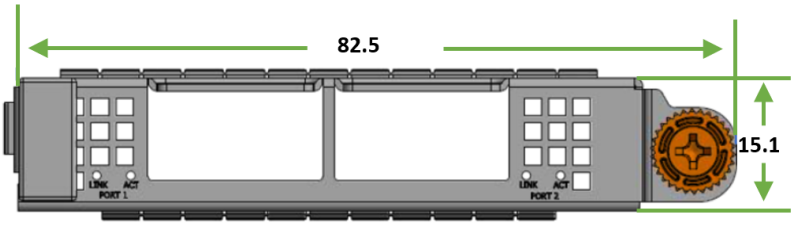

Single-port QSFP56 Thumbscrew Bracket

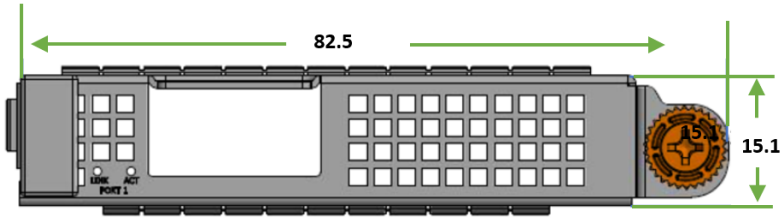

Finding the MAC and Serial Number on the Adapter Card

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.

MCX623432AN-ADAB Board Label Example

MCX623432AC-ADAB Board Label Example

MCX623432AN-GDAB Board Label Example

MCX623432AC-GDAB Board Label Example

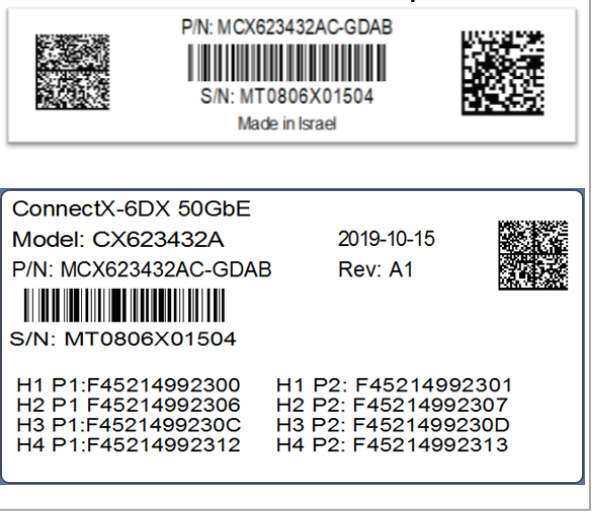

MCX623436AN-CDAB Board Label Example

MCX623436AC-CDAB Board Label Example

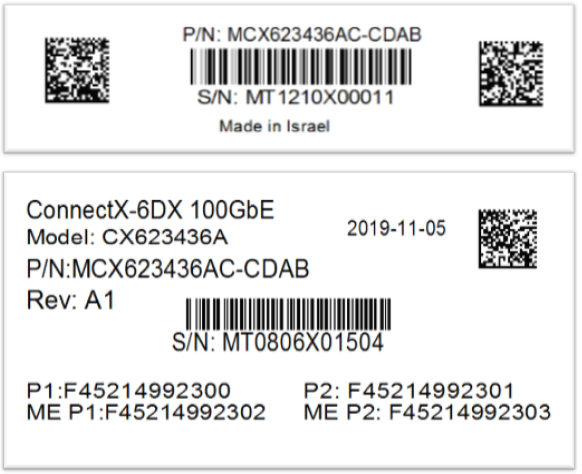

MCX623435AN-VDAB Board Label Example

MCX623435AC-VDAB Board Label Example

Document Revision History

Date	Revision	Description of Changes
Jul. 2022	2.6	Updated brackets' mechanical drawings.
Jun. 2022	2.5	Added Multi-Host description in the Introduction and Supported Interfaces.
Mar. 2022	2.4	Updated MCX623430MS-CDAB Specifications.
Feb. 2022	2.3	Added table Legacy (EOL) Ordering Part Numbers.
Aug. 2021	2.2	Added OPN MCX623430MS-CDAB.
Apr. 2021	2.1	Added OPNs MCX623436MN-CDAB and MCX623436MS-CDAB.
Mar. 2021	1.9	Added OPN MCX623432AS-ADAB.
Feb. 2021	1.8	Added OPN MCX623439MC[/N]-CDAB.
Jan. 2021	1.7	Added OPNs MCX623435AE-CDAB and MCX623435AN-CDAB
Oct. 2020	1.6	• Updated airflow numbers in Specifications. • Updated networking ports LEDs indications.
Sep. 2020	1.5	Added OCP 3.0 bracket replacement instructions.
Jul. 2020	1.4	Updated airflow numbers in Specifications.
Jul. 2020	1.3	Updated all power and airflow numbers in Specifications.
Jul. 2020	1.2	Added support for MCX623435MN-CDAB.
Feb. 2020	1.1	Updated Specifications with power numbers.
Nov. 2019	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.