



NVIDIA ConnectX-5 InfiniBand/Ethernet Socket Direct Adapter Cards User Manual

Table of contents

Introduction	9
Interfaces	15
Hardware Installation	17
Driver Installation	41
Linux Driver Installation	41
Troubleshooting	67
Updating Adapter Firmware	70
Specifications	72
Monitoring	78
Finding the GUID/MAC and Serial Number on the Adapter Card	79

List of Figures

Figure 0. Image2019 1 10 11 58 5

Figure 1. Image2021 8 11 8 44 35

Figure 2. Image2021 8 11 8 56 28

Figure 3. Image2021 8 11 8 44 35

Figure 4. Image2021 8 11 8 56 28

Figure 5. Image2021 8 11 8 44 35

Figure 6. Image2021 8 11 8 56 28

Figure 7. Image2021 8 11 8 44 35

Figure 8. Image2021 8 11 8 56 28

Figure 9. Image2021 8 11 8 44 3

Figure 10. Image2018 11 8 16 16 16

Figure 11. Image2018 11 15 11 34 36

Figure 12. Image2019 1 10 17 16 1

Figure 13. Image2019 1 10 17 17 57

Figure 14. Image2019 1 10 17 18 49

Figure 15. Image2019 1 10 17 20 6

Figure 16. Image2019 1 10 17 20 53

Figure 17. Image2018 11 14 15 55 33

Figure 18. Image2019 1 10 17 21 39

Figure 19. Image2019 1 10 17 22 31

Figure 20. Image2019 1 10 17 22 58

Figure 21. Image2019 1 10 17 23 19

Figure 22. Image2019 1 10 18 22 19

Figure 23. Image2019 1 10 18 22 32

Figure 24. Image2019 1 10 18 22 55

Figure 25. Image2019 1 13 10 58 45

Figure 26. Image2019 1 10 18 27 38

Figure 27. Image2019 1 10 18 28 17

List of Tables

Table 0.

Table 1.

Table 2.

Table 3.

Table 4.

Table 5.

Table 6.

Table 7.

Table 8.

Table 9.

Table 10.

Table 11.

Table 12.

Table 13.

Table 14.

Table 15.

Table 16.

Table 17.

Note

The ConnectX-5 product line has moved to end-of-life.

About This Manual

This User Manual describes NVIDIA® ConnectX®-5 InfiniBand/Ethernet Socket Direct adapter card supporting Dual-Socket Servers. The kit includes an adapter card with dual QSFP28 ports with PCI Express x8 edge connector, an auxiliary PCIe connection card with PCI Express x8 edge connector and a Slim-Line SAS cable which connects both cards. The User Manual provides details as to the interfaces of the adapter card, specifications, required software and firmware for operating the adapter card, and relevant documentation

EOL'ed (End of Life) Ordering Part Numbers

The table below provides the ordering part numbers (OPN) for the available ConnectX-5 InfiniBand/Ethernet Socket Direct adapter cards.

NVIDIA SKU	Legacy OPN	Marketing Description
900-9X5AD-0056-DT1	MCX55 6M- ECAT- S25	ConnectX®-5 InfiniBand/Ethernet adapter card with Socket Direct supporting dual-socket server, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, 2x PCIe3.0 x8, 25cm harness, tall bracket
900-9X5AD-0056-DT2	MCX55 6M- ECAT- S35A	ConnectX®-5 InfiniBand/Ethernet adapter card with Socket Direct supporting dual-socket server, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, 2x PCIe3.0 x8, 35cm harness, active auxiliary PCIe connection card, tall bracket

Intended Audience

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

Technical Support

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

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

MLNX_OFED for Linux User Manual and Release Notes	User Manual describing OFED features, performance, band diagnostic, tools content, and configuration. See MLNX_OFED for Linux Documentation .
WinOF-2 for Windows User Manual and	User Manual describing WinOF-2 features, performance, Ethernet diagnostic, tools content, and configuration. See WinOF-2 for Windows Documentation .

Release Notes	
NVIDIA VMware for Ethernet User Manual	User Manual and release notes describing the various components of the NVIDIA ConnectX® NATIVE ESXi stack. See VMware® ESXi Drivers Documentation .
NVIDIA Firmware Utility (mlxup) User Manual and Release Notes	NVIDIA firmware update and query utility used to update the firmware. Refer to Firmware Utility (mlxup) Documentation .
NVIDIA Firmware Tools (MFT) User Manual	User Manual describing the set of MFT firmware management tools for a single node. See MFT User Manual .
InfiniBand Architecture Specification Release 1.2.1, Vol 2 - Release 1.3	InfiniBand Specifications
IEEE Std 802.3 Specification	IEEE Ethernet Specifications

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

Document Conventions

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

Introduction

Product Overview

This is the User Guide for InfiniBand/Ethernet adapter cards based on the ConnectX®-5 integrated circuit device. NVIDIA offers an alternate ConnectX-5 Socket Direct™ card for servers without x16 PCIe slots. The adapter's 16-lane PCIe bus is split into two 8-lane buses, with one bus accessible through a PCIe x8 edge connector and the other bus through an x8 parallel connector to an Auxiliary PCIe Connection Card. The two cards should be installed into two adjacent PCIe x8 slots and connected using a dedicated harness.



Since the two PCIe cards are installed in two PCIe slots, each card gets a unique PCI Bus and Device number. Each of the PCIe x8 busses sees two network ports; in effect, the two 100Gb/s physical ports of the ConnectX-5 Socket Direct Adapter are viewed as four netdevices by the system, with each netdevice providing a maximum of 50Gb/s data transfer rate. In other words, 100Gb/s on a (physical) adapter port (Port 1 or Port 2) will be split on the PCIe side to 50Gb/s on each of the two PCIe cards (the Adapter Card and the Auxiliary PCIe Connection Cards, respectively, in the above picture).

ConnectX-5 InfiniBand/Ethernet Socket Direct Adapter Cards

Model	ConnectX-5 InfiniBand/Ethernet Adapter Cards	
Part Number	MCX556M-ECAT-S25	MCX556M-ECAT-S35A
Dimensions	Low Profile Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm) Auxiliary PCIe Connection Card Size: 4.44 in. x 1.57 in. (113.0mm x 40.00mm)	
	Slim-Line SAS Cable: 9.8 in. x 1.02 in. (250cm x 25.95mm)	Slim-Line SAS Cable: 9.8 in. x 1.02 in. (350cm x 25.95mm)
Network Connector Type	2x QSFP	
Data Rate	InfiniBand: SDR/DDR/QDR/FDR/EDR Ethernet: 10/25/40/50/100 Gb/s	
PCI Express Connectors	PCIe Gen 3.0 x8; SerDes @ 8.0GT/s	
RoHS	RoHS Compliant	
Adapter IC Part Number	MT28808A0-FCCF-EVM	
Device ID (decimal)	4119 for Physical Function (PF) and 4120 for Virtual Function (VF)	

Warning

Full performance of the adapter card (100Gb/s) is possible only when installing both the adapter card and the auxiliary PCIe connection card in the dual PCIe x8 connectors. Please refer to the performance optimization document for further guidance:

<https://community.nvidia.com/s/article/understanding-pcie-configuration-for-maximum-performance>

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

Features and Benefits

Warning

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

Feature	Description
PCI Express (PCIe)	Uses PCIe Gen 3.0 (8GT/s) through two x8 edge connectors. Gen 1.1 and 2.0 compatible.
100Gb/s InfiniBand/Ethernet Adapter	ConnectX-5 offers the highest throughput Ethernet adapter, supporting EDR 100Gb/s InfiniBand and 100Gb/s Ethernet and enabling any standard networking, clustering, or storage to operate seamlessly over any converged network leveraging a consolidated software stack.
InfiniBand Architecture Specification v1.3 compliant	ConnectX-5 delivers low latency, high bandwidth, and computing efficiency for performance-driven server and storage clustering applications. ConnectX-5 is InfiniBand Architecture Specification v1.3 compliant.
EDR InfiniBand	A standard InfiniBand data rate, where each lane of a 4X port runs a bit rate of 25.78125Gb/s with a 64b/66b encoding, resulting in an effective bandwidth of 100Gb/s.
Up to 100 Gigabit Ethernet	NVIDIA adapters comply with the following IEEE 802.3 standards: <ul style="list-style-type: none">• 100GbE / 50GbE / 40GbE / 25GbE / 10GbE / 1GbE• IEEE 802.3bj, 802.3bm 100 Gigabit Ethernet

Feature	Description
	<ul style="list-style-type: none"> • IEEE 802.3by, Ethernet Consortium25, 50 Gigabit Ethernet, supporting all FEC modes • IEEE 802.3ba 40 Gigabit Ethernet • IEEE 802.3ae 10 Gigabit Ethernet • IEEE 802.3ap based auto-negotiation and KR startup • Proprietary Ethernet protocols (20/40GBASE-R2, 50GBASE-R4) • IEEE 802.3ad, 802.1AX Link Aggregation • IEEE 802.1Q, 802.1P VLAN tags and priority • IEEE 802.1Qau (QCN) • Congestion Notification • IEEE 802.1Qaz (ETS) • IEEE 802.1Qbb (PFC) • IEEE 802.1Qbg • IEEE 1588v2 • Jumbo frame support (9.6KB)
Memory	<ul style="list-style-type: none"> • PCI Express - stores and accesses InfiniBand and/or Ethernet fabric connection information and packet data. • SPI Quad - includes 128Mbit SPI Quad Flash device (W25Q128FVSIQ device by ST Microelectronics) • FRU EEPROM - Stores the parameters and personality of the card. The EEPROM capacity is 128Kbit. FRU I²C address is (0x50) and is accessible through the PCIe SMBus. (Note: A ddress 0x58 is reserved.)
Overlay Networks	<p>In order to better scale their networks, data center operators often create overlay networks that carry traffic from individual virtual machines over logical tunnels in encapsulated formats such as NVGRE and VXLAN. While this solves network scalability issues, it hides the TCP packet from the hardware offloading engines, placing higher loads on the host CPU. ConnectX-5 effectively addresses this by providing advanced NVGRE and VXLAN hardware offloading engines that encapsulate and de-capsulate the overlay protocol.</p>
RDMA and RDMA over Converged	<p>ConnectX-5, utilizing IBTA RDMA (Remote Data Memory Access) and RoCE (RDMA over Converged Ethernet) technology, delivers low-latency and high performance over Band and Ethernet networks. Leveraging data center bridging (DCB) capabilities as well as ConnectX-5 advanced congestion control hardware mechanisms, RoCE provides efficient low-latency RDMA services over Layer 2 and Layer 3 networks.</p>

Feature	Description
Ethernet (RoCE)	
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-5 advanced acceleration technology enables higher cluster efficiency and scalability to tens of thousands of nodes.
CPU Offload	Adapter functionality enabling reduced CPU overhead allowing more available CPU for computation tasks.
Open VSwitch (OVS) offload using ASAP2	<ul style="list-style-type: none"> • 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-5 provides dedicated adapter resources and guaranteed isolation and protection for virtual machines within the server.
Storage Acceleration	<p>A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols can leverage InfiniBand RDMA for high-performance storage access.</p> <ul style="list-style-type: none"> • NVMe over Fabric offloads for target machine • Erasure Coding • T10-DIF Signature Handover
SR-IOV	ConnectX-5 SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server.

Feature	Description
High-Performance Accelerations	<ul style="list-style-type: none"> • Tag Matching and Rendezvous Offloads • Adaptive Routing on Reliable Transport • Burst Buffer Offloads for Background Checkpointing

Operating Systems/Distributions

Supported in OpenFabrics Enterprise Distribution (OFED).

Connectivity

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

Manageability

Socket Direct technology maintains support for manageability through a BMC. The Socket Direct PCIe stand-up adapter can be connected to a BMC using MCTP over SMBus or MCTP over PCIe protocols as if it is a standard NVIDIA PCIe stand-up adapter. For configuring the adapter for the specific manageability solution in use by the server, please contact NVIDIA Support.

Interfaces

InfiniBand Interface

The network ports of the ConnectX®-5 adapter cards are compliant with the InfiniBand Architecture Specification, Release 1.3. InfiniBand traffic is transmitted through the cards' QSFP28 connectors.

Ethernet QSFP28 Interface

The network ports of the ConnectX®-5 adapter card are compliant with the IEEE 802.3 Ethernet standards listed in [Features and Benefits](#). Ethernet traffic is transmitted through the cards' QSFP28 connectors.

PCI Express Interface

The ConnectX®-5 adapter card supports PCI Express Gen 3.0 (1.1 and 2.0 compatible) two PCIe x8 edge connectors; x8 edge connector on the adapter card and x8 edge connector on the auxiliary PCIe connection card accessible through the Slim Line-SAS cable. The device can be either a master initiating the PCI Express bus operations, or a slave responding to PCI bus operations.

- PCIe Gen3.0 compliant, 2.0 and 1.1 compatible
- 2.5, 5.0, or 8.0, or 16.0 GT/s link rate x8
- Auto-negotiates to x8, x4, x2, or x1
- Support for MSI/MSI-X mechanisms

LED Interface

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

Physical and Logical Link Indications - Ethernet Mode:

State	LED Behavior
Beacon command for locating the card	A link has not been established
Error	Blinks yellow at 4Hz until error is resolved. The reason for the error may be: <ul style="list-style-type: none"> • Failure of I²C access to the port • Over-current conditions of the port
Link Activity	Blinks green reflecting the link speed
Link Up	Constant green light

Physical and Logical Link Indications - InfiniBand Mode:

State	LED Behavior
Beacon command for locating the card	A link has not been established
Error	Blinks yellow at 4Hz until error is resolved. The reason for the error may be: <ul style="list-style-type: none"> • Failure of I²C access to the port • Over-current conditions of the port
Link Activity	Blinks green reflecting the link speed
Link Up	Constant green light

Hardware Installation

Installation and initialization of ConnectX-5 adapter cards require attention to the mechanical attributes, power specifications, and precautions for electronic equipment.

Safety Warnings

Note

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

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

Note that not all warnings are relevant to all models.



General Installation Instructions

Read all installation instructions before connecting the equipment to the power source.

Jewelry Removal Warning

Before you install or remove equipment that



is connected to power lines, remove jewelry such as bracelets, necklaces, rings, watches, and so on. Metal objects heat up when connected to power and ground and can meltdown, causing serious burns and/or welding the metal object to the terminals.



Over-temperature

This equipment should not be operated in an area with an ambient temperature exceeding the maximum recommended: 55°C (131°F). An airflow of 200LFM at this maximum ambient temperature is required for HCA cards and NICs. To guarantee proper airflow, allow at least 8cm (3 inches) of clearance around the ventilation openings.



During Lightning - Electrical Hazard

During periods of lightning activity, do not work on the equipment or connect or 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-5 Socket Direct adapter cards involve the following steps:

Step	Procedure	Direct Link
1	Check the system's hardware and software requirements	Refer to System Requirements
2	Pay attention to the airflow consideration within the host system	Refer to Airflow Requirements

Step	Procedure	Direct Link
3	Unpack the product	Refer to Unpacking the Package
4	(Optional) Replace the full-height mounting bracket with the supplied short bracket	Refer to Bracket Replacement
5	Connect the slim-line cable to the auxiliary connection card	Refer to Adapter Card Installation Instructions
6	Install the PCIe Auxiliary connection card in the system	
7	Install the adapter card in the system	
5	Connect the slim line-SAS cable to the ConnectX-5 adapter card	
6	Connect cables or modules to the card	Refer to Cables and Modules
7	Identify ConnectX-5 adapter card in the system	Refer to Identifying the Card in Your System

System Requirements

Hardware Requirements

Important

Unless otherwise specified, NVIDIA products are designed to work in an environmentally controlled data center with low levels of gaseous and dust (particulate) contamination.

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

A system with adjacent dual PCIe x8 slots is required for installing the card. For dual-socket servers, this card brings additional benefits of lower latency and lower CPU utilization when one slot is connected to one socket and the other slot is connected to the other socket. The below list specifies the servers tested with the adapter. For specific server compatibility, please contact NVIDIA Support.

- Dell PowerEdge_R720
- SuperMicro SYS-1027R-72RFTP

Important


For proper operation and performance, please make sure to use a PCIe slot with a corresponding bus width and that can supply sufficient power to your card. Refer to the Specifications section of the manual for more power requirements.

Airflow Requirements


ConnectX-5 adapter cards are offered with two airflow patterns: from the heatsink to the network ports, and vice versa, as shown below.

Please refer to [Airflow Specifications](#) for airflow numbers for each specific card model.

Airflow *from the heatsink* to the network ports:



Airflow from the network ports to the heatsink:



i Important

All cards in the system should be planned with the same airflow direction.

Software Requirements

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

Unpacking the Product

Important

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

- Remove any metallic objects from your hands and wrists.
- Make sure to use only insulated tools.
- Verify that the system is powered off and is unplugged.
- It is strongly recommended to use an ESD strap or other antistatic devices.

1. Unpack the adapter package and place them on an antistatic surface, and verify you have received the following items:

Category	Item
Cards	Low profile ConnectX-5 adapter card
	Auxiliary PCIe Connection Card

Category	Item
Harness	Slim-Line SAS Cable 25cm (for MCX556M-ECAT-S25) Slim-Line SAS Cable 35cm (for MCX556M-ECAT-S35A)
Accessories	Low profile card short bracket
	Auxiliary Card PCIe Connection card short bracket
	ConnectX-5 adapter card high-profile bracket (shipped assembled on the card)
	Auxiliary Card PCIe Connection card high-profile bracket (shipped assembled on the card)

Important

Use the Slim-Line SAS harness and Auxiliary PCIe Connection card that are included in the ConnectX-5 Socket Direct adapter card package contents. For MCX556M-ECAT-S25, use the 25cm Slim-line SAS cable and its Auxiliary PCIe connection card.

2. Check the parts for visible damage that may have occurred during shipping.
3. Shut down your system if active:
Turn off the power to the system, and disconnect the power cord. Refer to the system documentation for instructions. Before you install the ConnectX-5 card, make sure that the system is disconnected from power.

Bracket Replacement Instructions

The card is usually shipped with an assembled high-profile bracket. If this form factor is suitable for your requirements, you can skip the remainder of this section and move to [Installation Instructions](#). If you need to replace the high-profile bracket with the short bracket that is included in the shipping box, please follow the instructions in this section.

ⓘ Important

Due to risk of damaging the EMI gasket, it is not recommended to replace the bracket more than three times.

To replace the bracket you will need the following parts:

Removing the Existing Bracket

- The new brackets of the proper height
 - The 2 screws saved from the removal of the bracket
1. Using a torque driver, remove the two screws holding the bracket in place.
 2. Separate the bracket from the ConnectX-5 card.

ⓘ Important

Be careful not to put stress on the LEDs on the adapter card.

3. Save the two screws.

Installing the New Bracket

1. Place the bracket onto the card until the screw holes line up.

ⓘ Important

Do not force the bracket onto the adapter card.


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

 **Important**

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


Installation Instructions

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

 **Warning**

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

1. Before installing the card, make sure that the system is off and the power cord is not connected to the server. Please follow proper electrical grounding procedures.
2. Open the system case.
3. Connect the slim-line SAS connector (male) on the slim lineSAScable to the female connector on the auxiliary PCIe connection card.



(i) Important


Use the Slim-Line SAS harness and auxiliary PCIe connection card that are supplied in the ConnectX-5 Socket Direct adapter card package contents.

4. Locate two available PCI Express slots on the server, one for the adapter card and one for the auxiliary PCIe connection card.

ⓘ Important

For optimal thermal performance, it is preferable to place the auxiliary PCIe connection card component-side facing the adapter card's print-side.

5. Applying even pressure at both corners of the card, insert the auxiliary PCIe connection card into the PCI Express slot until firmly seated.




(i) Important

Do not use excessive force when seating the card, as this may damage the system or the auxiliary PCIe connection card.

6. Secure the auxiliary PCIe connection card using the server's retention mechanisms.


7. Applying even pressure at both corners of the card, insert the adapter card into the PCI Express slot adjacent to the auxiliary PCIe connection card until firmly seated.




(i) Important

Do not use excessive force when seating the card, as this may damage the system or the auxiliary PCIe connection card.

8. Secure the adapter card using the server's retention mechanisms.
9. Connect the slim line-SAS cable to the connector on the adapter card.



10. Close the system case.



i Note

To uninstall the adapter card, see [Uninstalling the Card](#).

Cables and Modules

Cable Installation

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

Warning

When installing cables make sure that the latches engage.

Important

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

3. After inserting a cable into a port, the Green LED indicator will light when the physical connection is established (that is, when the unit is powered on and a cable is plugged into the port with the other end of the connector plugged into a functioning port). See [Adapter Card LED Operations](#).
4. After plugging in a cable, lock the connector using the latching mechanism particular to the cable vendor. When data is being transferred the Green LED will blink. See Adapter Card LED Operations under the Interfaces section.
5. Care should be taken as not to impede the air exhaust flow through the ventilation holes. Use cable lengths which allow for routing horizontally around to the side of the chassis before bending upward or downward in the rack.
6. To remove a cable, disengage the locks and slowly pull the connector away from the port receptacle. 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-5]
```

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-5 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; and DEV is equal to 1018 (for ConnectX-5) – this is a valid NVIDIA PCI Device ID.

Warning

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

Warning

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

Uninstalling the Card

Safety Precautions

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

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

Card Removal


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

Note

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


1. Verify that the system is powered off and unplugged.
2. Wait 30 seconds.

3. Disconnect the slim-line SAS cable from the connector on the adapter card.




4. Disengage the retention mechanisms (clips or screws) on the adapter card.


5. Holding the card from its center, gently pull the adapter card from the PCI Express slot.



6. Disengage the retention mechanisms (clips or screws) on the auxiliary PCIe connection card.
7. Holding the card from its center, gently pull the auxiliary PCIe connection card from the PCI Express slot.d.



8. Disconnect the slim-line SAS cable from the connector on the auxiliary PCIe connection card.



Driver Installation

Please use the relevant driver installation section.

Linux Driver Installation

This section describes how to install and test the NVIDIA OFED for Linux package on a single host machine with Mellanox ConnectX-5 adapter hardware installed.

Prerequisites

Requirements	Description
Platforms	A server platform with a ConnectX-5 adapter card installed.
Required Disk Space for Installation	1GB
Operating System	Linux operating system. For the list of supported operating system distributions and kernels, please refer to the <i>MLNX_OFED Release Notes</i> file.
Installer Privileges	The installation requires administrator (root) privileges on the target machine.

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:

```
[root@mftqa-009 ~]# lspci |grep mellanox -i
05:00.0 Infiniband controller: Mellanox Technologies MT27800 Family
[ConnectX-5]
```

```
05:00.1 Infiniband controller: Mellanox Technologies MT27800 Family
[ConnectX-5]
82:00.0 Infiniband controller: Mellanox Technologies MT27800 Family
[ConnectX-5]
82:00.1 Infiniband controller: Mellanox Technologies MT27800 Family
[ConnectX-5]
```

In the output example above, the first two rows indicate that one card is installed in a PCI slot with PCI Bus address 05 (hexadecimal), PCI Device number 00 and PCI Function number 0 and 1. The other card is installed in a PCI slot with PCI Bus address 82 (hexadecimal), PCI Device number 00 and PCI Function number 0 and 1.

Since the two PCIe cards are installed in two PCIe slots, each card gets a unique PCI Bus and Device number. Each of the PCIe x8 busses sees two network ports; in effect, the two 100Gb/s physical ports of the ConnectX-5 Socket Direct Adapter are viewed as four netdevices by the system.

2. Download the ISO image to your host.

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

You can download and install the latest OpenFabrics Enterprise Distribution (OFED) software package available via the NVIDIA web site at nvidia.com/en-us/networking_Products_Software_InfiniBand_Drivers_NVIDIA_MLNX_OFED

1.

1. Scroll down to the Download wizard, and click the Download tab.
2. Choose your relevant package depending on your host operating system.
3. Click the desired ISO/tgz package.
4. To obtain the download link, accept the End User License Agreement (EULA).

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

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

Installing MLNX_OFED

Installation Script

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

- Discovers the currently installed kernel
- Uninstalls any software stacks that are part of the standard operating system distribution or another vendor's commercial stack
- Installs the `MLNX_OFED_LINUX` binary RPMs (if they are available for the current kernel)
- Identifies the currently installed InfiniBand and Ethernet network adapters and automatically upgrades the firmware

Note: To perform a firmware upgrade using customized firmware binaries, a path can be provided to the folder that contains the firmware binary files, by running `--fw-image-dir`. Using this option, the firmware version embedded in the `MLNX_OFED` package will be ignored.

Example:

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

Warning

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

Usage

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

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

Warning

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

- If you need to install OFED on an entire (homogeneous) cluster, a common strategy is to mount the ISO image on one of the cluster nodes and then copy it to a shared file system such as NFS. To install on all the cluster nodes, use cluster-aware tools (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.

Warning

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

Warning

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

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

```
dracut -f
```

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

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

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

Warning

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

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

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

```
Removing OFED RPMs...
```

```
Created /tmp/MLNX_OFED_LINUX-x.x-x-rhel7.3-x86_64-ext.tgz
```

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

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

For the list of installation options, run:

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

Installation Procedure

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

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

```
host1# mount -o ro,loop MLNX_OFED_LINUX-<ver>-<OS label>-<CPU arch>.iso  
/mnt
```

3. Run the installation script.

```
/mnt/mlnxofedinstall
```

Logs dir: /tmp/MLNX_OFED_LINUX-x.x-x.logs

This program will install the MLNX_OFED_LINUX [package](#) on your machine.

Note that all other Mellanox, OEM, OFED, RDMA or Distribution IB packages will be removed.

Those packages are removed due to conflicts with MLNX_OFED_LINUX, [do](#) not reinstall them.

Starting MLNX_OFED_LINUX-x.x.x installation ...

.....

.....

Installation finished successfully.

Attempting to perform Firmware update...

Querying Mellanox devices firmware ...

Warning

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

```
script:/mnt/mlnxofedinstall --force
```

Warning

MLNX_OFED for Ubuntu should be installed with the following flags in chroot environment: `./mlnxofedinstall --without-dkms --`

```
add-kernel-support --kernel <kernel version in chroot> --
```

```
without-fw-update --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.

Warning

In case your machine has the latest firmware, no firmware update will occur and the installation script will print at the end of installation a message similar to the following: Device #1:-----
---Device Type: ConnectX-XPart Number: MCXXXX-XXXPSID:
MT_2190110032PCI Device Name: 0b:00.0Base MAC:
0000e41d2d5cf810Versions: Current AvailableFW 12.14.0114
12.14.0114Status: Up to date

Warning

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

Error message #1:Device #1:-----Device Type: ConnectX-XPart
Number: MCXXXX-XXXPSID: MT_2190110032PCI Device Name:
0b:00.0Base MAC: 0000e41d2d5cf810Versions: Current
AvailableFW 12.14.0114 N/AStatus: No matching image found

Error message #2:The firmware for this device is not distributed
inside NVIDIA driver: 0000:01:00.0 (PSID: IBM2150110033)To
obtain firmware for this device, please contact your HW vendor.

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

Action \ Adapter	Driver Restart	Standard Reboot (Soft Reset)	Cold Reboot (Hard Reset)
Standard ConnectX-4/ConnectX-4 Lx or higher	-	+	-
Adapters with Multi-Host Support	-	-	+
Socket Direct Cards	-	-	+

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

5. (InfiniBand only) Run the `hca_self_test.ofed` utility to verify whether or not the InfiniBand link is up. The utility also checks for and displays additional information such as:

- HCA firmware version
- Kernel architecture
- Driver version
- Number of active HCA ports along with their states
- Node GUID

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

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

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

Warning

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

Installation Results

Soft ware	<ul style="list-style-type: none">• Most of MLNX_OFED packages are installed under the “/usr” directory except for the following packages which are installed under the “/opt” directory:<ul style="list-style-type: none">◦ fca and ibutils◦ iproute2 (rdma tool) - installed under /opt/Mellanox/iproute2/sbin/rdma• The kernel modules are installed under<ul style="list-style-type: none">◦ /lib/modules/`uname -r`/updates on SLES and Fedora Distributions◦ /lib/modules/`uname -r`/extra/mlnx-ofa_kernel on RHEL and other RedHat like Distributions◦ /lib/modules/`uname -r`/updates/dkms/ on Ubuntu
Firm ware	<ul style="list-style-type: none">• The firmware of existing network adapter devices will be updated if the following two conditions are fulfilled:<ul style="list-style-type: none">◦ The installation script is run in default mode; that is, without the option ‘--without- fw-update’◦ The firmware version of the adapter device is older than the firmware version included with the OFED ISO imageNote: If an adapter’s Flash was originally programmed with an Expansion ROM image, the automatic firmware update will also burn an Expansion ROM image.• In case your machine has an unsupported network adapter device, no firmware update will occur and the error message below will be printed. "The firmware for this device is not distributed inside NVIDIA driver: 0000:01:00.0 (PSID: IBM2150110033) To obtain firmware for this device, please contact your HW vendor."

Installation Logging

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

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

Example:

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

Driver Load Upon System Boot

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

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

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

```
blacklist mlx5_core  
blacklist mlx5_ib
```

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

mlnxofedinstall Return Codes

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

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

Soft ware	<ul style="list-style-type: none"> • Most of MLNX_OFED packages are installed under the “/usr” directory except for the following packages which are installed under the “/opt” directory: <ul style="list-style-type: none"> ◦ fca and ibutils ◦ iproute2 (rdma tool) - installed under /opt/Mellanox/iproute2/sbin/rdma • The kernel modules are installed under <ul style="list-style-type: none"> ◦ /lib/modules/`uname -r`/updates on SLES and Fedora Distributions ◦ /lib/modules/`uname -r`/extra/mlnx-ofa_kernel on RHEL and other RedHat like Distributions ◦ /lib/modules/`uname -r`/updates/dkms/ on Ubuntu
Firm ware	<ul style="list-style-type: none"> • The firmware of existing network adapter devices will be updated if the following two conditions are fulfilled: <ul style="list-style-type: none"> ◦ The installation script is run in default mode; that is, without the option ‘--without- fw-update’ ◦ The firmware version of the adapter device is older than the firmware version included with the OFED ISO image <p>Note: If an adapter’s Flash was originally programmed with an Expansion ROM image, the automatic firmware update will also burn an Expansion ROM image.</p>

- In case your machine has an unsupported network adapter device, no firmware update will occur and the error message below will be printed.
"The firmware for this device is not distributed inside NVIDIA driver:
0000:01:00.0 (PSID: IBM2150110033)
To obtain firmware for this device, please contact your HW vendor."

Installation Logging

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

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

Example:

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

Uninstalling MLNX_OFED

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

Additional Installation Procedures

Installing MLNX_OFED Using YUM

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

Setting up MLNX_OFED YUM Repository

1. Log into the installation machine as root.

2. Mount the ISO image on your machine and copy its content to a shared location in your network.

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

3. Download and install NVIDIA's GPG-KEY:

The key can be downloaded via the following link:

<http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-Mellanox>

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

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

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

4. Install the key.

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

5. Check that the key was successfully imported.

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

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

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

7. Check that the repository was successfully added.

```
# yum repolist  
Loaded plugins: product-id, security, subscription-manager  
This system is not registered to Red Hat Subscription Management. You can  
use subscription-manager to register.  
repo id repo name status  
mlnx_ofed MLNX_OFED Repository 108  
rpmforge RHEL 6Server - RPMforge.net - dag 4,597  
  
repolist: 8,351
```

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

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

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


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

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

Checking if all needed packages are installed...
Building MLNX_OFED_LINUX RPMS . Please wait...
Creating metadata-rpms for 3.10.0-957.21.3.el7.x86_64 ...
WARNING: If you are going to configure this package as a repository, then please
note
WARNING: that it contains unsigned rpms, therefore, you need to disable the
gpgcheck
WARNING: by setting 'gpgcheck=0' in the repository conf file.
Created /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-ext.tgz
```

4. Open the tarball.

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

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

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

6. Check that the repository was successfully added.

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

repolist: 8,351
```

Installing MLNX_OFED Using the YUM Tool

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

1. View the available package groups by invoking:

```
# yum search mlnx-ofed-
mlnx-ofed-all.noarch : MLNX_OFED all installer package (with KMP support)
mlnx-ofed-all-user-only.noarch : MLNX_OFED all-user-only installer package
(User Space packages only)
mlnx-ofed-basic.noarch : MLNX_OFED basic installer package (with KMP
support)
mlnx-ofed-basic-user-only.noarch : MLNX_OFED basic-user-only installer
package (User Space packages only)
mlnx-ofed-bluefield.noarch : MLNX_OFED bluefield installer package (with KMP
support)
mlnx-ofed-bluefield-user-only.noarch : MLNX_OFED bluefield-user-only
installer package (User Space packages only)
mlnx-ofed-dpdk.noarch : MLNX_OFED dpdk installer package (with KMP
support)
mlnx-ofed-dpdk-upstream-libs.noarch : MLNX_OFED dpdk-upstream-libs
installer package (with KMP support)
```

mlnx-ofed-dpdk-upstream-libs-user-only.noarch : MLNX_OFED dpdk-upstream-libs-user-only installer [package](#) (User Space packages only)
 mlnx-ofed-dpdk-user-only.noarch : MLNX_OFED dpdk-user-only installer [package](#) (User Space packages only)
 mlnx-ofed-eth-only-user-only.noarch : MLNX_OFED eth-only-user-only installer [package](#) (User Space packages only)
 mlnx-ofed-guest.noarch : MLNX_OFED guest installer [package](#) (with KMP support)
 mlnx-ofed-guest-user-only.noarch : MLNX_OFED guest-user-only installer [package](#) (User Space packages only)
 mlnx-ofed-hpc.noarch : MLNX_OFED hpc installer [package](#) (with KMP support)
 mlnx-ofed-hpc-user-only.noarch : MLNX_OFED hpc-user-only installer [package](#) (User Space packages only)
 mlnx-ofed-hypervisor.noarch : MLNX_OFED hypervisor installer [package](#) (with KMP support)
 mlnx-ofed-hypervisor-user-only.noarch : MLNX_OFED hypervisor-user-only installer [package](#) (User Space packages only)
 mlnx-ofed-kernel-only.noarch : MLNX_OFED kernel-only installer [package](#) (with KMP support)
 mlnx-ofed-vma.noarch : MLNX_OFED vma installer [package](#) (with KMP support)
 mlnx-ofed-vma-eth.noarch : MLNX_OFED vma-eth installer [package](#) (with KMP support)
 mlnx-ofed-vma-eth-user-only.noarch : MLNX_OFED vma-eth-user-only installer [package](#) (User Space packages only)
 mlnx-ofed-vma-user-only.noarch : MLNX_OFED vma-user-only installer [package](#) (User Space packages only)
 mlnx-ofed-vma-vpi.noarch : MLNX_OFED vma-vpi installer [package](#) (with KMP support)
 mlnx-ofed-vma-vpi-user-only.noarch : MLNX_OFED vma-vpi-user-only installer [package](#) (User Space packages only)

where:

mlnx-ofed-all	Installs all available packages in MLNX_OFED
mlnx-ofed-basic	Installs basic packages required for running NVIDIA cards

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

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

Example:

```

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

```

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

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

2. Install the desired group.

```
# yum install mlnx-ofed-all
Loaded plugins: langpacks, product-id, subscription-manager
Resolving Dependencies
--> Running transaction check
---> Package mlnx-ofed-all.noarch 0:3.1-0.1.2 will be installed
--> Processing Dependency: kmod-iser = 1.0-OFED.3.1.0.1.2.1.g832a737.rhel7u1
for package: mlnx-ofed-all-3.1-0.1.2.noarch
.....
.....
qperf.x86_64 0:0.4.9-9
rds-devel.x86_64 0:2.0.7-1.12
rds-tools.x86_64 0:2.0.7-1.12
sdpNetstat.x86_64 0:1.60-26
srptools.x86_64 0:1.0.2-12
```

Complete!

Warning

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

Installing MLNX_OFED Using apt-get

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

Setting up MLNX_OFED apt-get Repository

1. Log into the installation machine as root.
2. Extract the MLNX_OFED package on a shared location in your network. It can be downloaded from <https://www.nvidia.com/en-us/networking/Products Software InfiniBand Drivers>.
3. Create an apt-get repository configuration file called `"/etc/apt/sources.list.d/mlnx_ofed.list"` with the following content:

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

4. Download and install NVIDIA's Technologies GPG-KEY.

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

5. Verify that the key was successfully imported.

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

6. Update the apt-get cache.

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

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

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

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

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

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

Checking if all needed packages are installed...
Building MLNX_OFED_LINUX RPMS . Please wait...
Creating metadata-rpms for 3.10.0-957.21.3.el7.x86_64 ...
WARNING: If you are going to configure this package as a repository, then please
note
WARNING: that it contains unsigned rpms, therefore, you need to disable the
gpgcheck
WARNING: by setting 'gpgcheck=0' in the repository conf file.
```

```
Created /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-ext.tgz
```

4. Open the tarball.

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

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

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

6. Update the apt-get cache.

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

Installing MLNX_OFED Using the apt-get Tool

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

1. View the available package groups by invoking:

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


mlnx-ofed-hpc-user-only - MLNX_OFED hpc-user-only installer [package](#) (User Space packages only)

mlnx-ofed-dpdk-user-only - MLNX_OFED dpdk-user-only installer [package](#) (User Space packages only)

mlnx-ofed-all-exact - MLNX_OFED all installer [package](#) (with DKMS support) (exact)

mlnx-ofed-all - MLNX_OFED all installer [package](#) (with DKMS support)

mlnx-ofed-vma-vpi-user-only - MLNX_OFED vma-vpi-user-only installer [package](#) (User Space packages only)

mlnx-ofed-eth-only-user-only - MLNX_OFED eth-only-user-only installer [package](#) (User Space packages only)

mlnx-ofed-vma-user-only - MLNX_OFED vma-user-only installer [package](#) (User Space packages only)

mlnx-ofed-hpc - MLNX_OFED hpc installer [package](#) (with DKMS support)

mlnx-ofed-bluefield-user-only - MLNX_OFED bluefield-user-only installer [package](#) (User Space packages only)

mlnx-ofed-dpdk - MLNX_OFED dpdk installer [package](#) (with DKMS support)

mlnx-ofed-vma-eth-user-only - MLNX_OFED vma-eth-user-only installer [package](#) (User Space packages only)

mlnx-ofed-all-user-only - MLNX_OFED all-user-only installer [package](#) (User Space packages only)

mlnx-ofed-vma-eth - MLNX_OFED vma-eth installer [package](#) (with DKMS support)

mlnx-ofed-vma - MLNX_OFED vma installer [package](#) (with DKMS support)

mlnx-ofed-dpdk-upstream-libs-user-only - MLNX_OFED dpdk-upstream-libs-user-only installer [package](#) (User Space packages only)

mlnx-ofed-basic-user-only - MLNX_OFED basic-user-only installer [package](#) (User Space packages only)

mlnx-ofed-basic-exact - MLNX_OFED basic installer [package](#) (with DKMS support) (exact)

mlnx-ofed-basic - MLNX_OFED basic installer [package](#) (with DKMS support)

mlnx-ofed-dpdk-upstream-libs - MLNX_OFED dpdk-upstream-libs installer [package](#) (with DKMS support)

where:

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

2. Install the desired group.

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

Example:

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

Warning

Installing MLNX_OFED using the “apt-get” tool does not automatically update the firmware. To update the firmware to

```
the version included in MLNX_OFED package, run:# apt-get install  
mlnx-fw-updater
```

Performance Tuning

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

Troubleshooting

General Troubleshooting

Server unable to find the adapter	<ul style="list-style-type: none">• Ensure that the adapter is placed correctly• Make sure the adapter slot and the adapter are compatible Install the adapter in a different PCI Express slot• Use the drivers that came with the adapter or download the latest• Make sure your motherboard has the latest BIOS• Try to reboot the server
The adapter no longer works	<ul style="list-style-type: none">• Reseat the adapter in its slot or a different slot, if necessary• Try using another cable• Reinstall the drivers for the network driver files may be damaged or deleted• Reboot the server
Adapters stopped working after installing another adapter	<ul style="list-style-type: none">• Try removing and re-installing all adapters• Check that cables are connected properly• Make sure your motherboard has the latest BIOS
Link indicator light is off	<ul style="list-style-type: none">• Try another port on the switch• Make sure the cable is securely attached• Check you are using the proper cables that do not exceed the recommended lengths• Verify that your switch and adapter port are compatible
Link light is on, but with no communication established	<ul style="list-style-type: none">• Check that the latest driver is loaded• Check that both the adapter and its link are set to the same speed and duplex settings

<p>Event message received of insufficient power</p>	<ul style="list-style-type: none"> • When [adapter's current power consumption] > [PCIe slot advertised power limit] – a warning message appears in the server's system even logs (Eg. dmesg: "Detected insufficient power on the PCIe slot") • It's recommended to use a PCIe slot that can supply enough power. • If a message of the following format appears – "mlx5_core 0003:01:00.0: port_module:254:(pid 0): Port module event[error]: module 0, Cable error, One or more network ports have been powered down due to insufficient/unadvertised power on the PCIe slot" please upgrade your Adapter's firmware. • If the message remains – please consider switching from Active Optical Cable (AOC) or transceiver to Direct Attached Copper (DAC) connectivity.
---	--

Linux Troubleshooting

<p>Environment Information</p>	<pre>cat /etc/issue uname -a cat /proc/cupinfo grep 'model name' uniq ofed_info -s ifconfig -a ip link show ethtool <interface> ethtool -i <interface_of_Mellanox_port_num> ibdev2netdev</pre>
<p>Card Detection</p>	<pre>lspci grep -i Mellanox</pre>
<p>Mellanox Firmware Tool (MFT)</p>	<p>Download and install MFT: MFT Documentation Refer to the User Manual for installation instructions. Once installed, run:</p> <pre>mst start mst status flint -d <mst_device> q</pre>
<p>Ports Information</p>	<pre>ibstat ibv_devinfo</pre>

Firmware Version Upgrade	To download the latest firmware version, refer to the NVIDIA Update and Query Utility .
Collect Log File	<pre>cat /var/log/messages dmesg >> system.log journalctl (Applicable on new operating systems) cat /var/log/syslog</pre>

Windows Troubleshooting

Environment Information	<p>From the Windows desktop choose the Start menu and run: msinfo32</p> <p>To export system information to a text file, choose the Export option from the File menu.</p> <p>Assign a file name and save.</p>
Mellanox Firmware Tool (MFT)	<p>Download and install MFT: MFT Documentation</p> <p>Refer to the User Manual for installation instructions.</p> <p>Once installed, open a CMD window and run:</p> <pre>WinMFT mst start mst status flint -d <mst_device> q</pre>
Ports Information	<pre>vstat</pre>
Firmware Version Upgrade	<p>Download the latest firmware version using the PSID/board ID from here.</p> <pre>flint -d <mst_device> -i <firmware_bin_file> b</pre>
Collect Log File	<ul style="list-style-type: none"> • Event log viewer • MST device logs: <ul style="list-style-type: none"> ◦ mst start ◦ mst status • <code>flint -d <mst_device> dc > dump_configuration.log</code> • <code>mstdump <mst_device> dc > mstdump.log</code>

Updating Adapter Firmware

Each card is shipped with the latest version of qualified firmware at the time of manufacturing. However, NVIDIA issues firmware updates occasionally that provide new features and bug fixes. To check that your card is programmed with the latest available firmware version, download the mlxup firmware update and query utility. The utility can query for available NVIDIA adapters and indicate which adapters require a firmware update. If the user confirms, mlxup upgrades the firmware using embedded images. The latest mlxup executable and documentation are available in [inmlxup - Update and Query Utility](#).

```
[server1]# ./mlxup
Querying Mellanox devices firmware ...
Device Type: ConnectX-5
Part Number: MCX556M-ECAT-S25
Description: ConnectX®-5 VPI adapter card with Socket Direct supporting dual-socket server, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, 2x PCIe3.0 x8, 25cm harness, tall bracket
PSID: MT_2190110032
PCI Device Name: 0000:06:00.0
Base GUID: e41d2d0300fd8b8a
Versions: Current Available
FW 16.00.0000 16.00.0000
Status: Up to date
Device Type: ConnectX-5
Part Number: MCX556M-ECAT-S25
Description: ConnectX®-5 VPI adapter card with Socket Direct supporting dual-socket server, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, 2x PCIe3.0 x8, 25cm harness, tall bracket
PSID: MT_2170110021
PCI Device Name: 0000:07:00.0
```

Base MAC: 0000e41d2da206d4

Versions: Current Available

FW 16.00.0000 16.00.0000

Status: Update required

Perform FW update? [y/N]: y

Device #1: Up to date


Device #2: Updating FW ... Done

Restart needed for updates to take effect.

Log File: /var/log/mlxup/mlxup-yyyymmdd.log

Specifications

MCX556M-ECAT-S25 Specifications

Physical	Low Profile Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm) Auxiliary PCIe Connection Card Size: 4.44 in. x 1.57 in. (113.0mm x 40.00mm) Slim-Line SAS Cable: 9.8 in. x 1.02 in. (250cm x 25.95mm)					
	Connector: Dual QSFP28 InfiniBand and Ethernet (copper and optical)					
Protocol Support	InfiniBand: IBTA v1.3 ^a Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port					
	Ethernet: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR					
	Data Rate	<table border="1"> <tr> <td>InfiniBand</td> <td>SDR/DDR/QDR/FDR/EDR</td> </tr> <tr> <td>Ethernet</td> <td>1/10/25/40/50/100 Gb/s</td> </tr> </table>	InfiniBand	SDR/DDR/QDR/FDR/EDR	Ethernet	1/10/25/40/50/100 Gb/s
	InfiniBand	SDR/DDR/QDR/FDR/EDR				
Ethernet	1/10/25/40/50/100 Gb/s					
PCI Express Gen3: SERDES @ 8.0GT/s, 2x x8 lanes (2.0 and 1.1 compatible)						
<p> Warning PCIe 3.0 x16 bus can supply a maximum bandwidth of 128Gb/s only (=16 * 8GT/s, including overhead), and therefore cannot support 200Gb/s when both network ports of MCX556A-ECAT run at 100Gb/s.</p>						
Power and	Voltage: 12V					
	Power	Cable				

Airflow	Typical Power ^b	Passive Cables	17.1W
	Maximum Power	Passive Cables	20.0W
		1.5W Active Cables	23.3W
		2.5W Active Cables	25.5W
	Maximum power available through QSFP28 port: 5W		
	Airflow: see Airflow Specifications		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)		3050m
Regulatory	Safety: CB / cTUVus / CE		
	EMC: CE / FCC / VCCI / ICES / RCM		
	RoHS: RoHS compliant		


^a The ConnectX-5 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.

^b Typical power for ATIS traffic load.

^c For both operational and non-operational states.

MCX556M-ECAT-S35A Specifications

Physical	Low Profile Adapter Card Size: 6.6 in. x 2.71 in. (167.65mm x 68.90mm) Auxiliary PCIe Connection Card Size: 4.44 in. x 1.57 in. (113.0mm x
-----------------	---

	40.00mm) Slim-Line SAS Cable: 9.8 in. x 1.02 in. (350cm x 25.95mm)		
	Connector: Dual QSFP28 InfiniBand and Ethernet (copper and optical)		
Protocol Support	InfiniBand: IBTA v1.3 ^a Auto-Negotiation: 1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port		
	Ethernet: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR, 10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR		
	Data Rate	InfiniBand	SDR/DDR/QDR/FDR/EDR
		Ethernet	1/10/25/40/50/100 Gb/s
	PCI Express Gen3: SERDES @ 8.0GT/s, 2x x8 lanes (2.0 and 1.1 compatible)		
<div style="background-color: #f8d7da; padding: 10px;"> <p> Warning PCIe 3.0 x16 bus can supply a maximum bandwidth of 128Gb/s only (=16 * 8GT/s, including overhead), and therefore cannot support 200Gb/s when both network ports of MCX556A-ECAT run at 100Gb/s</p> </div>			
Adapter Card Power	Voltage: 12V		
	Power	Cable	
	Typical Power ^b	Passive Cables	17.1W
		Passive Cables	20.0W
		1.5W Active Cables	23.3W
	Maximum Power	2.5W Active Cables	25.5W
Maximum power available through QSFP28 port: 5W			

Active Auxiliary PCIe Connecti on Card Power	Typical Power	1.46W	
	Maximum Power	1.85W	
	Airflow: see Airflow Specifications		
Environm ental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
Altitude (Operational)	3050m		
Regulator y	Safety: CB / cTUVus / CE		
	EMC: CE / FCC / VCCI / ICES / RCM		
	RoHS: RoHS compliant		

^a The ConnectX-5 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.

^b Typical power for ATIS traffic load.

^c For both operational and non-operational states.

Airflow Specifications

Airflow (LFM)					
Airflow Direction - Heat Sink to Port					
Cable Type	Passive	Active 1.5W	Active 2.5W	Active 3.5W	Active 5W
MCX556M-ECAT-S25	400	400	500	600	900
MCX556M-ECAT-S35A	400	400	500	600	900

LED Specifications

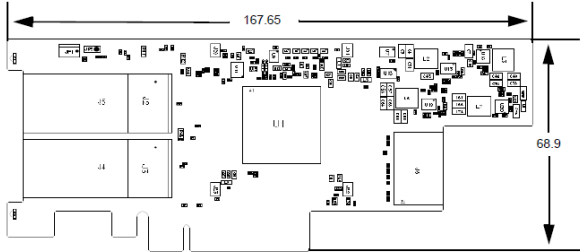
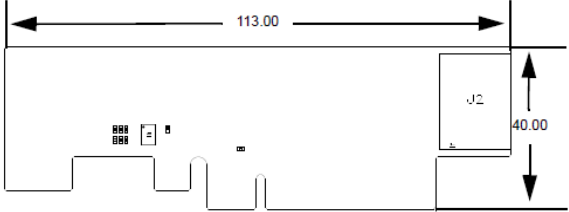
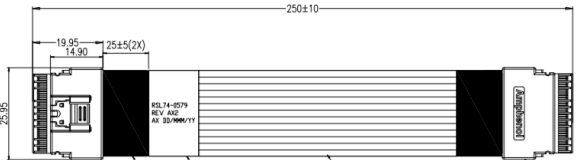
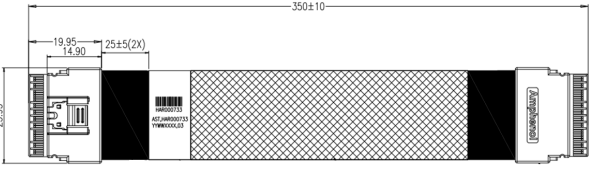
See [LED Interface](#).

Board Mechanical Drawings and Dimensions



Warning

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

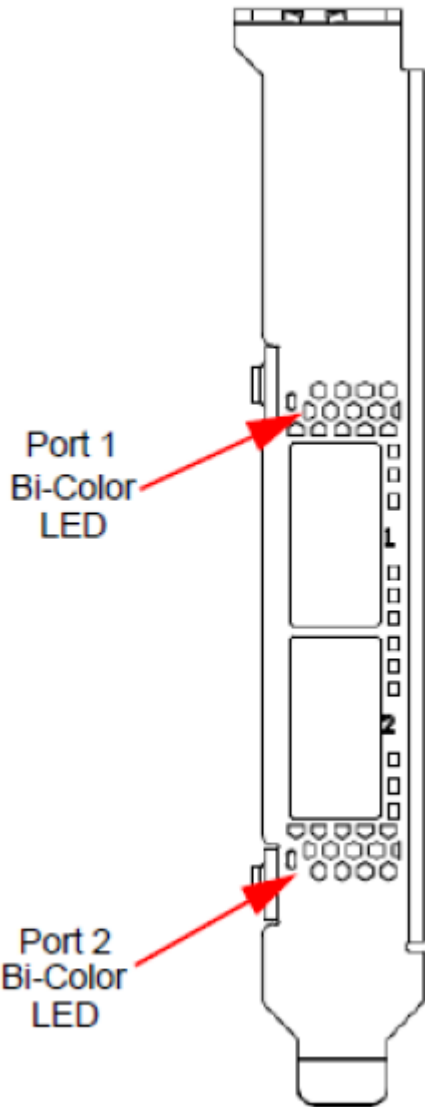
For the 3D Model of the card, please refer to http://www.nvidia.com/page/3d_models.

<p>Dual-port Adapter Card</p> 	<p>Auxiliary PCIe Connection Card</p> 
<p>Slim-Line SAS Cable (25cm)</p> 	<p>Slim-Line SAS Cable (35cm)</p> 


Bracket Mechanical Drawing

<p>Adapter Card Tall Bracket</p> 	<p>Auxiliary PCIe Connection Card Tall Bracket</p> 
--	---

Adapter Card Tall Bracket



Auxiliary PCIe Connection Card Tall Bracket



Monitoring

Thermal Sensors

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

Three thermal threshold definitions impact the overall system operation state:

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

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

Adapter Card Heatsink

The heatsink is attached to the ConnectX-5 IC to dissipate the heat from the ConnectX-5 IC. It is attached by spring-loaded screws.

ConnectX-5 IC has a thermal shutdown safety mechanism that automatically shuts down the ConnectX-5 card in cases of high-temperature events, improper thermal coupling, or heatsink removal.

For the required airflow (LFM) per OPN, please refer to the [Specifications](#) chapter.

Finding the GUID/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 and the card GUID for the InfiniBand protocol. VPI cards have both a GUID and a MAC (derived from the GUID).

i Note

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


MCX556M-ECAT-S25 Board Label (Example)



Auxiliary PCIe Connection Card Board Label



Slim-Line SAS Harness Board Label



HAR000629

AST,HAR000629

YYWWXXX,08

© Copyright 2023, NVIDIA. PDF Generated on 06/07/2024