



NVIDIA External Multi-Host Adapter Kit for OCP 3.0 User Manual

Table of Contents

Introduction.....	6
Externally Connected Multi-Host Solution (eMH)	6
System Requirements	7
Package Contents	7
Product Overview	8
Features and Benefits	9
Interfaces	11
Product Layout and Interface Information	11
Interfaces Detailed Description.....	12
Multi-Host Board	12
OCP 3.0 Card	12
PCI Express Interface.....	12
Mini-SAS Connectors	12
Networking Interface.....	13
NC-SI Management Interface	13
Mini USB Type B Interface	13
1GbE OOB Management Interface	13
External PCIe Power Supply Connector	14
Pinouts Description	15
NC-SI Management Interface	15
External Power Supply Connector	16
Hardware Installation	17
Safety Warnings	17
Installation Procedure Overview.....	17
System Requirements	17
Hardware Requirements	17
Airflow Requirements	18
Software Requirements	18
Unpacking.....	18
Setup Preparation.....	18
Installation Instructions	19
Cables and Modules.....	23

Kit Power Sequence	24
DIP Switches Control	25
DIP Switches Description.....	25
SW6-1, SW6-2, SW5-1	25
MAIN_POWER_ENABLE and AUX_POWER_ENABLE DIP-Switches.....	25
DIP-Switches Bifurcation.....	26
Driver Installation.....	28
Linux Driver Installation	28
Prerequisites.....	28
Downloading MLNX_OFED	28
Installing MLNX_OFED	29
Driver Load Upon System Boot.....	34
Additional Installation Procedures	36
Performance Tuning	42
Specifications.....	43
MEMH-OCP3-AK70 Specifications	43
Board Mechanical Drawings and Dimensions.....	43
Multi-Host Board	44
MiniSAS Auxiliary Card.....	44
OCP 3.0 SFF Card	45
Finding the GUID/MAC and Serial Number on the Product	46
Document Revision History	47

About This Manual

This User Manual describes NVIDIA® External Multi-Host Adapter Kit for OCP 3.0 with a connectivity of up to four standard servers. It provides details as to the interfaces of the board, specifications, required software and firmware for operating the board, and relevant documentation.

EOL'ed (End of Life) Ordering Part Numbers

The table below provides the ordering part numbers (OPN) for the available adapter card kit.

OPN	Marketing Description
MEMH-OCP3-AK70	External Multi-host OCP3.0 Adapter with a connectivity of up to 4 standard servers

Intended Audience

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

Technical Support

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

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

Customers who purchased NVIDIA M-1 Global Support Services, please see your contract for details regarding Technical Support.

Customers who purchased NVIDIA products through an NVIDIA-approved reseller should first seek assistance through their reseller.

Related Documentation

NVIDIA MLNX_OFED for Linux User Manual and Release Notes	User Manual and Release Notes describing MLNx_OFED features, performance, band diagnostic, tools content and configuration. See NVIDIA MLNX_OFED for Linux Documentation .
WinOF-2 for Windows User Manual and Release Notes	User Manual describing WinOF-2 features, performance, Ethernet diagnostic, tools content and configuration. See WinOF-2 for Windows Documentation .
NVIDIA VMware for Ethernet User Manual and Release Notes	User Manual describing the various components of the NVIDIA ConnectX® NATIVE ESXi stack. See VMware® ESXi Documentation .
NVIDIA Firmware Update	NVIDIA firmware update and query utility used to update the firmware. See NVIDIA Firmware Utility (mlxup) Documentation .
NVIDIA Firmware Tools (MFT) User Manual	User Manual describing the set of MFT firmware management tools for a single node. See MFT User Manual .


NVIDIA MLNX_OFED for Linux User Manual and Release Notes	User Manual and Release Notes describing MLNX_OFED features, performance, band diagnostic, tools content and configuration. See NVIDIA MLNX_OFED for Linux Documentation .
IEEE Std 802.3 Specification	IEEE Ethernet Specifications
PCI Express Specifications	Industry Standard PCI Express Base and Card Electromechanical Specifications .
LinkX Interconnect Solutions	LinkX Ethernet cables and transceivers are designed to maximize the performance of High-Performance Computing networks, requiring high-bandwidth, low-latency connections between compute nodes and switch nodes. NVIDIA offers one of the industry's broadest portfolio of 40GbE, 56GbE, 100GbE, 200GbE and 400GbE cables, including Direct Attach Copper cables (DACs), copper splitter cables, Active Optical Cables (AOCs) and transceivers in a wide range of lengths from 0.5m to 10km. In addition to meeting Ethernet standards, NVIDIA tests every product in an end-to-end environment ensuring a Bit Error Rate of less than 1E-15. Read more at LinkX Cables and Transceivers .

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 size in MegaBits. IB is used in this document to mean InfiniBand. In this document, PCIe is used to mean PCI Express.

Introduction

This is the User Guide for NVIDIA External Multi-Host Adapter Kit for OCP 3.0 with a connectivity of up to 4 standard servers.

 Caution: Powering up the MiniSAS auxiliary cards before the OCP 3.0 multi-host board is powered up may cause not power-up the card.


Externally Connected Multi-Host Solution (eMH)


NVIDIA Multi-Host technology enables connecting up to 4 compute / storage hosts to a single OCP 3.0 multi-host adapter. The deployment of multi-host platforms significantly reduces the overall number of data-center network connections, enabling great infrastructure efficiency and simplicity with CAPEX and OPEX cost savings. NVIDIA Multi-Host technology is built into ConnectX SmartNICs and BlueField DPUs and is ideal for high-performance, compute-intensive, data-center environments delivering cloud, web 2.0, and telecom services.

The externally connected multi-host solution leverages the same multi-host technology built into the network adapter ASIC. The external Mini-SAS harnesses make up the connectivity between each host and the network card. The solution is perfectly positioned for highly dense 4-node 2U chassis systems, enabling 25Gb/s connectivity per connected node, as illustrated in the figure below.

The externally connected Multi-host solution offers a superior price/performance ratio for both new and existing scale-out computing fabrics, while dramatically reducing the total cost of ownership on the following:

- Network adapters, operating as a single network adapter, serve up to 4 nodes.
- Network switch ports, operating as a single switch port, serve up to 4 nodes.
- Active cabling, as a single cable now serves up to four nodes.
- Rack space, power and cooling, attributed to the overall reduction in network connections in the data center.

 The below figure is for illustration purposes only. The OCP 3.0 card is not included in the package.



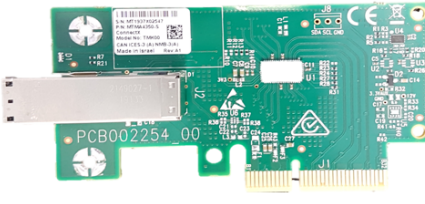
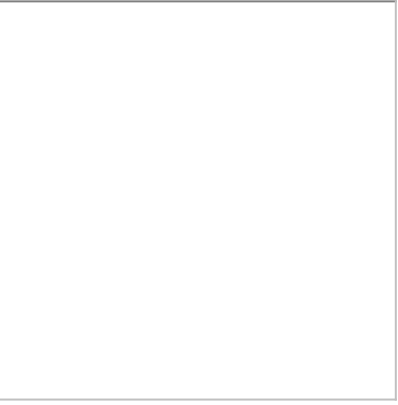


System Requirements

Power Supply	Minimum of 120W external ATX power supply source (not included in the package contents)
External Fan	Critical requirement: at least one external fan (not included in the package contents). Please refer to the airflow requirements listed in the specifications table.
Operating System	OpenFabrics Enterprise Distribution (OFED)
Connectivity	<ul style="list-style-type: none"> • Interoperable with up to 100Gb/s Ethernet switches • Passive copper cable with ESD protection • Powered connectors for optical and active cable support

Package Contents

Prior to unpacking your product, it is important to make sure your system meets all the requirements listed above for a smooth installation. Be sure to inspect each piece of equipment shipped in the packing box. If anything is missing or damaged, contact your reseller.

	Item	Illustration
Cards	1x Multi-Host OCP 3.0 adapter	

	Item	Illustration
	4x MiniSAS Auxilliary boards	
Cables	x4 MiniSAS HD cables	
	USB type-A to Mini-USB type B cable	
Accessories	I/O Panel ^(a)	


(a) The I/O panel is included in the package and can be used when inserting the product into a server. Currently, the option is not supported.


Product Overview

Model	OCP 3.0 Multi-Host Kit
Part Number	MEMH-OCP3-AK70
Network Connector Type	Supports any OCP 3.0 adapter card
Dimensions	Multi-Host Board: 6.69in x 6.69in (170mm x 170mm) OCP 3.0 Adapter Card: Small Form Factor (SFF) OCP 3.0 - 4.52 in. x 2.99 in (115.00mm x 76.00mm) 4x MiniSAS Auxiliary Cards: 1.96 in. x 3.54 in (50.0mm x 90.0 mm) 4x Mini SAS Harnesses: 70cm each

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.

 For visualization scenarios, please contact NVIDIA support.

Feature	Description
PCI Express (PCIe)	Uses the following PCIe interfaces: <ul style="list-style-type: none"> • OCP 3.0 Card: PCIe Gen 3.0 SERDES @ 8.0 GT/s through x8 edge connector • MiniSAS Auxilliary Cards: 4x PCIe Gen 3.0 SERDES @ 8.0GT/s through x4 edge connector
Memory	<ul style="list-style-type: none"> • SPI Quad - includes 128Mbit SPI Quad Flash device (W25Q128JVS1Q device WINBOND-NUVOTON). • FRU EEPROM - Provides the parameters and personality of the card. The EEPROM capacity is 128Kbit. FRU I²C address is (0x50) and is accessible through the PCIe SMBus.
NC-SI over RMII	A Network Controller Sideband Interface (NC-SI) is a combination of logical and physical paths that interconnect the Management Controller and Network Controller(s) for the purpose of transferring management communication traffic among them. NC-SI includes commands and associated responses which the Management Controller uses to control the status and operation of the Network Controller(s). NC-SI also includes a mechanism for transporting management traffic and asynchronous notifications. Please connect an ethernet cable from the management console to J104 of the OCP 3.0 host adapter.
Wake-on-LAN	Wake-on-LAN (WoL) is a feature that allows a network professional to remotely power on a server (4 servers together) or be awakened by a network message. The feature is applicable only when the OCP 3.0 card supports the feature.
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. OCP Multi-Host card 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)	OCP Multi-Host card, 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 OCP Multi-Host card, 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. OCP Multi-Host card 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


Feature	Description
Quality of Service (QoS)	Support for port-based Quality of Service enabling various application requirements for latency and SLA.
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 the target machine • Erasure Coding • T10-DIF Signature Handover
SR-IOV	OCP Multi-Host card SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server.
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

Interfaces

This section describes the product's supported interfaces. Each numbered interface that is referenced in the figures is described in the following table with a link to detailed information.

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

Product Layout and Interface Information



Item	Interface	Description
1	"Multi-Host Board"	External Multi-host board with a connectivity of up to 4 standard servers.
2	"OCP 3.0 card"	The multi-host board populates an OCP 3.0 card.
3	"PCI Express Interface"	PCIe Gen 3.0/4.0 through an edge connector.

Item	Interface	Description
4	"MiniSAS Connectors"	Used to connect the Multi-host board to the MiniSAS Auxiliary cards using the supplied harnesses.
5	"Networking Port"	InfiniBand or Ethernet traffic is transmitted through the OCP 3.0 card networking connectors. The networking connectors allow for the use of modules, optical and passive cable interconnect solutions.
6	"Mini USB Type B Interface"	Mounted on the multi-host board for OS image loading.
7	"NC-SI Management Interface"	BMC connection for remote management
8	"1GbE OOB Management Interface"	1GbE BASE-T OOB management interface.
9	"External PCIe Power Supply Connector"	An external 12V power connection through an 8-pin ATX connector

Interfaces Detailed Description

Multi-Host Board

The Multi-Host OCP 3.0 board enables connecting up to four standard servers to a single multi-host network adapter.

OCP 3.0 Card

The Multi-host board can populate an OCP 3.0 adapter card. The card is not included in the package.

PCI Express Interface

The multi-host board supports PCI Express Gen 3.0/4.0 (1.1 and 2.0 compatible) through edge connectors. The following lists PCIe interface features:

- PCIe Gen 4.0 and 3.0 compliant, 2.0 and 1.1 compatible
- 2.5, 5.0, or 8.0, or 16.0 GT/s link rate
- Auto-negotiates to x16, x8, x4, x2, or x1
- Support for MSI/MSI-X mechanisms

Mini-SAS Connectors

The Mini-SAS connectors are used to connect the Mini-SAS auxiliary cards with the multi-host board using the supplied harnesses.

The MiniSAS auxiliary cards are shipping with short brackets assembled on the card.

Networking Interface

The network ports of the OCP 3.0 card are compliant with the IEEE 802.3 InfiniBand and Ethernet standards. InfiniBand and Ethernet traffic is transmitted through the cards' networking connectors.

NC-SI Management Interface

The multi-host board enables the connection of a Baseboard Management Controller (BMC) to a set of Network Interface Controller (NICs) for the purpose of enabling out-of-band remote manageability. The NC-SI management is supported over RMII and has a connector on the multi-host board. NC-SI over the RMII interface is routed through the RMII connector with a flat cable. The cable Please refer to [NC-SI Management Interface](#) for pins.

Mini USB Type B Interface

The Multi-host board incorporates a Mini USB Type B connector. The purpose of the USB connector is to load operating system images. In order to use this interface, please use the supplied USB cable in the package.

1GbE OOB Management Interface

The multi-host incorporates a 1GbE RJ45 out-of-band port that allows the network operator to establish trust boundaries in accessing the management function to apply it to network resources. It can also be used to ensure management connectivity (including the ability to determine the status of any network component) independent of the status of other in-band network components. Use an RJ45 cable to connect to the 1GbE OOB management interface.

RJ45 Cable for the 1GbE OOB Management Cable



1GbE OOB Management LEDs Interface

The multi-host board incorporates a 1GbE RJ45 out-of-band port that allows the network operator to establish trust boundaries in accessing the management function to apply it to network resources. It can also be used to ensure management connectivity (including the ability to

determine the status of any network component) independent of the status of other in-band network components.

There are two I/O LEDs, LED1 is green and LED2 is Amber, to indicate link activity as described in the below table.

LED Pin	LED Definition	Link Activity
LED2 (Amber)	OFF	Link OFF
	ON	Link ON (Any speed)
LED1 (Green)	OFF	No Activity
	Blinking	Activity (RX, Tx)

External PCIe Power Supply Connector

The Multi-host board requires an external 12V power connection through an 8-pin ATX connector to power up. The external power supply source is not included in the package.

For the external PCIe power supply pins, please refer to [External PCIe Power Supply Connector Pins](#).

Pinouts Description

NC-SI Management Interface

The below table lists the NC-SI pins description. For further details, please refer to [NC-SI Management Interface](#).



Pin#	Signal Name	I/O	Description
1	REF_CLK	I	50M REF CLK for NCSI BUS
2	GND	GND	Ground
3	ARB_IN	I	NCSI hardware arbitration input
4	GND	GND	Ground
5	ARB_OUT	O	NCSI hardware arbitration output
6	GND	GND	Ground
7	RX_D0	O	Receive data
8	GND	GND	Ground
9	RX_D1	O	Receive data
10	GND	GND	Ground
11	CRS_DV	O	Carrier sense/Receive Data Valid
12	GND	GND	Ground
13	TX_D0	I	Transmit data
14	RSRV		Reserved
15	TX_D1	I	Transmit data
16	GND	GND	Ground
17	TX_EN	I	Transmit enable
18	GND	GND	Ground
19	NC		
20	GND	GND	Ground
21	I2C_SDA	B	I2C Serial Data
22	GND	GND	Ground
23	I2C_SCL	I	I2C Serial Clock
24	GND	GND	Ground
25	GND	GND	Ground
26	GND	GND	Ground
27	3.3V	VDD	Power
28	3.3V	VDD	Power
29	3.3V	VDD	Power
30	3.3V	VDD	Power

External Power Supply Connector

The below table provides the External Power Supply pins of the external power supply interfaces on the multi-host motherboard. For further details, please refer to [External PCIe Power Supply Connector](#).

The mechanical pinout of the 8-pin external +12V power connector is shown below. The +12V connector is a GPU power PCIe standard connector. Care should be taken to ensure the power is applied to the correct pins as some 8-pin ATX type connector can have different pinouts.


Pin#	Signal Name	Description
1	GND	Power Return
2	GND	Power Return
3	GND	Power Return
4	GND	Power Return
5	12V	ATX Supplied 12V
6	12V	ATX Supplied 12V
7	12V	ATX Supplied 12V
8	12V	ATX Supplied 12V



Hardware Installation

Installation and initialization of the external multi-host OCP 3.0 adapter kit require attention to the mechanical attributes, power specification, and precautions for electronic equipment.

Safety Warnings

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

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

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


Installation Procedure Overview

The product can be used either in a desktop mode (without installing it in a server), or by inserting it in a designated server. Follow the below procedure for kit installation.

Step	Procedure
1	Check your setup meets the hardware and software requirements.
2	Unpack the product.
3	Prepare your setup.
5	Install the main OCP 3.0 card in the multi-host board.
6	Connect the multi-host board with the external ATX 12V power supply.
7	Install the MiniSAS Auxiliary board in the server.
8	Connect the MiniSAS Auxiliary boards using the supplied MiniSAS HD harnesses.
9	Power up the server.
10	Identify the cards in the system.

System Requirements

Hardware Requirements

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

Airflow Requirements


The user is required to supply at least one external fan for cooling the multi-host board populated with the OCP 3.0 card. Please refer to [Specifications](#) for airflow requirements.

 Failure to supply sufficient airflow may cause damage to the product.

Software Requirements

- See Operating Systems/Distributions section under the Introduction section.
- Software Stacks - NVIDIA OpenFabric software package MLNX_OFED for Linux. See the Driver Installation section.

Unpacking

 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 multi-host OCP 3.0 adapter kit package, place it on an antistatic surface, and verify you have received the following items:
 - a. The multi-host OCP 3.0 board
 - b. 4x MiniSAS Auxiliary boards, with short brackets assembled on each card.
 - c. 4x Mini-SAS HD cables.
 - d. 1x USB type A to mini-USB type B cable.
 - e. I/O panel (to be used when the multi-host OCP 3.0 board is installed in a server. Currently, this option is not supported).
2. Check the parts for visible damage that may have occurred during shipping.
3. Shut down your system if active. Refer to the system documentation for instructions.

Setup Preparation

Follow the below steps prior to installing the product.

1. Verify the DIP-switches position is as stated in the below table and figure.




SW#	Position
SW7-2	ON
SW7-1	ON
SW9-2	ON
SW9-1	OFF
SW8-2	ON
SW8-1	ON
SW6-2	OFF
SW6-1	OFF
SW5-2	ON
SW5-1	ON
SW4-2	OFF
SW4-1	OFF
SW3-2	ON
SW3-2	ON

Installation Instructions




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

➤ Place the multi-host board on an antistatic pad.




➤ Install the OCP 3.0 adapter card in the multi-host board.


1. Prior to installing the OCP 3.0 adapter card in the multi-host board, unscrew the screw as shown in the below figure.




2. Applying even pressure at both corners of the card, insert the OCP 3.0 card into the multi-host board using the guiding rails until a click sound is heard.



- Secure the OCP card with the screw you removed in the previous step.





- When the OCP 3.0 adapter is properly seated, the card's golden fingers are connected to the multi-host OCP 3.0 connector.



➤ Connect the external main power connector with the 12V power source.

External Power Supply Connector



TBD: What kind of power cable? need pin description.

Should we add anything regarding LED operations?

- *Install the MiniSAS auxiliary cards in a server and connect them to the multi-host board.*
 1. Applying even pressure at both corners of the card, insert the MiniSAS auxiliary cards to a server.


Server



2. When the MiniSAS auxiliary cards are properly seated, the harness connector is aligned with the slot opening, and the MiniSAS auxiliary faceplate is visible against the system chassis.

➤ *Connect the MiniSAS auxiliary cards with the OCP multi-host board using the supplied MiniSAS cables.*

➤ **Power up the server. Please refer to the [power sequence instructions](#).**


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

Cables and Modules

Cable Installation

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


 When installing cables make sure that the latches engage.


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

3. After inserting a cable into a port, the Green LED indicator will light when the physical connection is established (that is, when the unit is powered on and a cable is plugged into the port with the other end of the connector plugged into a functioning port). See [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 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.

Kit Power Sequence

- Install the OCP 3.0 card to the multi-host board and power up.
- Install the MiniSAS auxiliary cards and connect them to the multi-host board using the supplied harnesses.
- Power up the server.

 Powering up the auxiliary cards while the OCP 3.0 card is not powered up may cause serious damage to the cards.

 Once powered up, the OCP 3.0 board FRU EEPROM address in multi-host mode is changed to 0x56.

In power down, power down the MiniSAS auxiliary cards and then the OCP 3.0 card.

DIP Switches Control


DIP Switches Description

SW6-1, SW6-2, SW5-1

The below table defines the SMBus connectivity of the OCP 3.0 board to hosts 0-3.

SW#	SW6-1	SW6-2	SW5-1	Notes
I2C_MUX[Y0]: NC	OFF	OFF	OFF	
I2C_MUX[Y1]: HOST#0	OFF	OFF	ON	Default
I2C_MUX[Y2]: HOST#1	OFF	ON	OFF	
I2C_MUX[Y3]: HOST#2	OFF	ON	ON	
I2C_MUX[Y4]: HOST#3	ON	OFF	OFF	
I2C_MUX[Y5]: NCSI	ON	OFF	ON	
I2C_MUX[Y1]: HOST#0	ON	ON	OFF	Unused (Defined as Host#0)
I2C_MUX[Y1]: HOST#0	ON	ON	ON	Unused (Defined as Host#0)

TBD




MAIN_POWER_ENABLE and AUX_POWER_ENABLE DIP-Switches

The below table defines the MAIN_POWER_ENABLE and AUX_POWER_ENABLE DIP-switches status.

⚠ The OCP3.0 board should be connected to MEMH-OCP3 multi-host adapter.


SW#	SW3-2 (bit[2];AUX_POWER_EN ABLE)	SW3-1 (bit[1];MAIN_POWER_EN ABLE)	Notes
MAIN_POWER_ENABLE = constant "0"	N/A	OFF	
MAIN_POWER_ENABLE = constant "1" (Illegal mode, the OCP board is turned off)	OFF	ON	
MAIN_POWER_ENABLE = "1" according to the power sequence flow	ON	ON	Default
AUX_POWER_ENABLE = "0"	OFF	N/A	
AUX_POWER_ENABLE = "1" according to the power sequence flow	ON	N/A	Default



DIP-Switches Bifurcation

The below table defines the bifurcation bits status. The default value of the multi-host board is 0b110.

SW#	SW8-1	SW9-2	SW9-1
BIF [2]	ON	-	-
BIF [1]	-	ON	-
BIF [0]	-	-	OFF



Driver Installation

Please use the relevant driver installation section.

Linux Driver Installation

This section describes how to install and test the MLNX_OFED for Linux package on a single server with a ConnectX-5 adapter card installed.

Prerequisites

Requirements	Description
Platforms	A server platform with a ConnectX-5 Ethernet 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 MLNX_OFED

1. Verify that the system has a network adapter installed by running `lspci` command. The below table provides output examples per ConnectX-5 card configuration.

ConnectX-5 Card Configuration	Output Examples
Single-port Socket Direct Card (2x PCIe x16)	<pre>[root@mftqa-009 ~]# lspci grep mellanox -i a3:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-5] e3:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-5]</pre>
Single-port PCIe x16 Card	<pre>[root@mftqa-009 ~]# lspci grep mellanox -ia 3:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-5]</pre>
Dual-port PCIe x16 Card	<pre>[root@mftqa-009 ~]# lspci grep mellanox -ia 86:00.0 Network controller: Mellanox Technologies MT2910 Family [ConnectX-5] 86:00.1 Network controller: Mellanox Technologies MT2910 Family [ConnectX-5]</pre>

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](#)

- i. Scroll down to the Download wizard, and click the Download tab.
 - ii. Choose your relevant package depending on your host operating system.
 - iii. Click the desired ISO/tgz package.
 - iv. To obtain the download link, accept the End User License Agreement (EULA).
3. Use the Hash utility to confirm the file integrity of your ISO image. Run the following command and compare the result to the value provided on the download page.

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

Installing MLNX_OFED

Installation Script


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

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

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

Example:


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

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

Usage

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

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

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

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

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

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

```
dracut -f
```

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

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

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

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

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

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

Building OFED RPMs. Please wait...
Removing OFED RPMs...
Created /tmp/MLNX_OFED_LINUX-x.x-x-rhel7.3-x86_64-ext.tgz
```

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

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

For the list of installation options, run:

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

Installation Procedure

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

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

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

- Run the installation script.

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

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

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

- ⚠** MLNX_OFED for Ubuntu should be installed with the following flags in chroot environment:

```
./mlnxofedinstall --without-dkms --add-kernel-support --kernel <kernel
version in chroot> --without-fw-update --force
```

For example:

```
./mlnxofedinstall --without-dkms --add-kernel-support --kernel 3.13.0-85-
generic --without-fw-update --force
```

Note that the path to kernel sources (`--kernel-sources`) should be added if the sources are not in their default location.

⚠ In case your machine has the latest firmware, no firmware update will occur and the installation script will print at the end of installation a message similar to the following:

```
Device #1:
-----
Device Type: ConnectX-X
Part Number: MCXXXX-XXX
PSID: MT_<version>
PCI Device Name: 0b:00.0
Base MAC: 0000e41d2d5cf810
Versions: Current Available
FW XX.XX.XXXX
Status: Up to date
```

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

```
Error message #1:
Device #1:
-----
Device Type: ConnectX-X
Part Number: MCXXXX-XXX
PSID: MT_<version>
PCI Device Name: 0b:00.0
Base MAC: 0000e41d2d5cf810
Versions: Current Available
FW XX.XX.XXXX
Status: No matching image found
Error message #2:
The firmware for this device is not distributed inside NVIDIA driver:
0000:01:00.0 (PSID: IBM2150110033)
To obtain firmware for this device, please contact your HW vendor.
```

- d. 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”.

- e. (InfiniBand only) Run the `hca_self_test.ofed` utility to verify whether or not the InfiniBand link is up. The utility also checks for and displays additional information such as:
- HCA firmware version
 - Kernel architecture
 - Driver version
 - Number of active HCA ports along with their states
 - Node GUID
- For more details on `hca_self_test.ofed`, see the file `docs/readme_and_user_manual/hca_self_test.readme`.

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

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

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

Installation Results

Software	<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"> • <code>fca</code> and <code>ibutils</code> • <code>iproute2</code> (rdma tool) - installed under <code>/opt/Mellanox/iproute2/sbin/rdma</code> • The kernel modules are installed under <ul style="list-style-type: none"> • <code>/lib/modules/`uname -r`/updates</code> on SLES and Fedora Distributions • <code>/lib/modules/`uname -r`/extra/mlnx-ofa_kernel</code> on RHEL and other RedHat like Distributions • <code>/lib/modules/`uname -r`/updates/dkms/</code> on Ubuntu
----------	--

Firmware	<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.IBM2150110033.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:

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

```
blacklist mlx5_core
blacklist mlx5_ib
```

- Set "ONBOOT=no" in the "/etc/infiniband/openib.conf" file.
- If the modules exist in the initramfs file, they can automatically be loaded by the kernel. To prevent this behavior, update the initramfs using the operating systems' standard tools.

Note: The process of updating the initramfs will add the blacklists from step 1, and will prevent the kernel from loading the modules automatically.

mlnxofedinstall Return Codes

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

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

Software	<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
Firmware	<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

- a. Log into the installation machine as root.
- b. 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
```

- c. 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]
```

- d. 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]:
```

- e. 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>)
```

- f. 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
```

- g. 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

- Log into the installation machine as root.
- 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
```

- 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
```

- Open the tarball.

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

- 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
```

- 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:

- 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 3.17.4-3
01.fc21.x86_64 (without KMP support)
mlnx-ofed-guest-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED guest installer package for kernel 3.17.4-3
01.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.

- b. 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.rhel7ul 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
sdpnstat.x86_64 0:1.60-26
srptools.x86_64 0:1.0.2-12
Complete!
```

⚠ Installing MLNX_OFED using the “YUM” tool does not automatically update the firmware.

To update the firmware to the version included in MLNX_OFED package, run:

```
# yum install mlnx-fw-updater
```

Installing MLNX_OFED Using apt-get

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

Setting up MLNX_OFED apt-get Repository

- a. Log into the installation machine as root.
- b. 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.
- c. 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 ./
```

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

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

e. 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
```

f. Update the apt-get cache.

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

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

- Log into the installation machine as root.
- 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
```

c. 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
```

d. Open the tarball.

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

e. 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 ./
```

f. 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:

- 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

b. Install the desired group.

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

Example:

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

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

To update the firmware to the version included in MLNX_OFED package, run:

```
# apt-get install mlnx-fw-updater
```

Performance Tuning

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


Specifications

MEMH-OCP3-AK70 Specifications


Physical	Multi-Host Board: 6.67 in. x 6.69 in (170mm x 170mm) OCP 3.0 Adapter Card: Small Form Factor (SFF) OCP 3.0 - 4.52 in. x 2.99 in (115.00mm x 76.00mm) 4x MiniSAS Auxiliary Cards: 1.96 in. x 3.54 in (50.0mm x 90.0 mm) 4x Mini SAS Harnesses: 70cm each		
OCP 3.0 Adapter Card Power	Voltage: 12V		
	Typical Power	Up to 80W	
	Maximum Power	Up to 120W	
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C ^a
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Airflow	Airflow to the Multi-host board populated with the OCP 3.0 card	Minimum of 500LFM @ 45°C	
Regulatory	Safety: CB / cTUVus / CE		
	EMC: CE / FCC / VCCI / ICES / RCM		
	RoHS compliant		

^a The non-operational storage temperature specifications apply to the product without its package.


Board Mechanical Drawings and Dimensions

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


Multi-Host Board



MiniSAS Auxiliary Card



OCP 3.0 SFF Card



Finding the GUID/MAC and Serial Number on the Product

Each adapter kit has a different identifier printed on the label: serial number, the card MAC and P/N.

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

External Multi-Host OCP 3.0 Board Label



Document Revision History

Date	Revision	Description of Changes
Ddc. 2020	1.0	First Release

Notice

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

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice. Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

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

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

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

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

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

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

Trademarks

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



Copyright

© 2023 NVIDIA Corporation & affiliates. All Rights Reserved.

