



ConnectX®-4 Multi-Host Evaluation Kit User Manual

P/N:
MCX4MHEVK-ECAA

Rev 1.0

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (“PRODUCT(S)”) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES “AS-IS” WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER’S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies
350 Oakmead Parkway Suite 100
Sunnyvale, CA 94085
U.S.A.
www.mellanox.com
Tel: (408) 970-3400
Fax: (408) 970-3403

© Copyright 2015. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, ConnectX®, Connect-IB®, CoolBox®, CORE-Direct®, InfiniBridge®, InfiniHost®, InfiniScale®, MetroX®, MLNX-OS®, TestX®, PhyX®, ScalableHPC®, SwitchX®, UFM®, Virtual Protocol Interconnect® and Voltaire® are registered trademarks of Mellanox Technologies, Ltd.

ExtendX™, FabricIT™, HPC-X™, Mellanox CloudX™, Mellanox Open Ethernet™, Mellanox PeerDirect™, Mellanox Virtual Modular Switch™, MetroDX™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Revision History

This document was printed on August 13, 2015.

Table 1 - Revision History Table

Date	Rev	Comments/Changes
August 2015	1.0	First Release

1 Overview

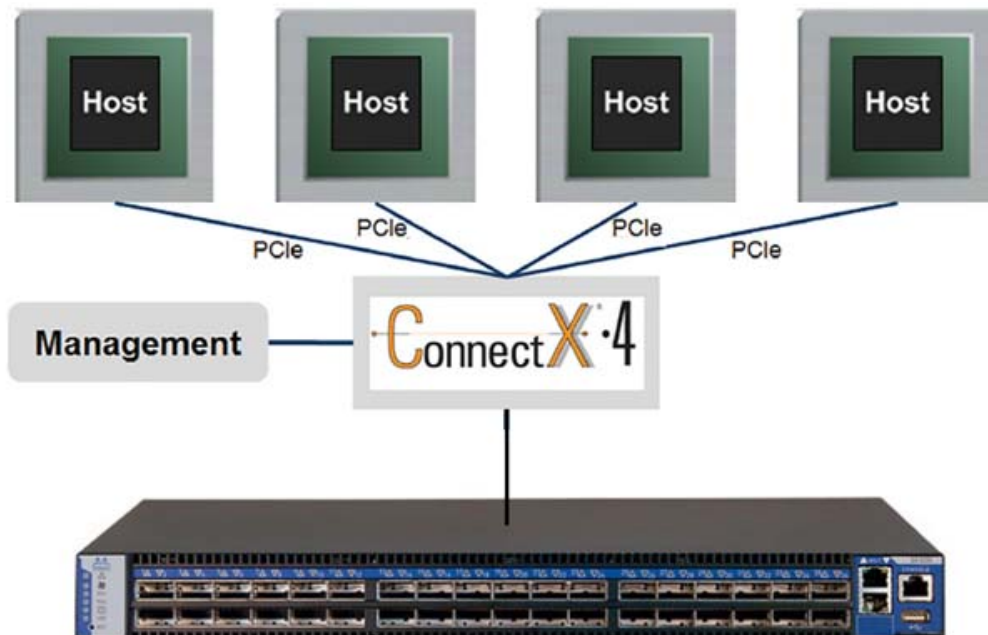
Mellanox's ConnectX®-4 Multi-Host technology enables connecting multiple hosts into a single interconnect adapter by separating the ConnectX-4 PCIe interface into multiple and independent PCIe interfaces. Each interface is connected to a separate host with no performance degradation.

The connection of four fully-independent PCIe buses to four hosts lowers total cost of ownership in the data center. It reduces CAPEX requirements from four cables, four adapters (NICs), and four switch ports to only one of each. Furthermore, it reduces OPEX by cutting down on switch port management and overall power consumption.

Multi-Host technology features uncompromising independent host management, with full independent NC-SI/MCTP support to each host and to the NIC. IT managers can remotely control the configuration and power state of each host individually, such that management of one host does not affect host traffic performance or the management of the other hosts, guaranteeing host security and isolation. To further lower the total cost of ownership, ConnectX-4 supports management of the multiple hosts using a single BMC, with independent NC-SI/MCTP management channels for each of the managed hosts.

The below figure provides an illustration of the a Multi-Host technology that the ConnectX-4 Multi-Host Evaluation Board (EVB) enables. The ConnectX-4 adapter connects to the hosts over the PCI Express interface, with dedicated PCIe lanes per host.

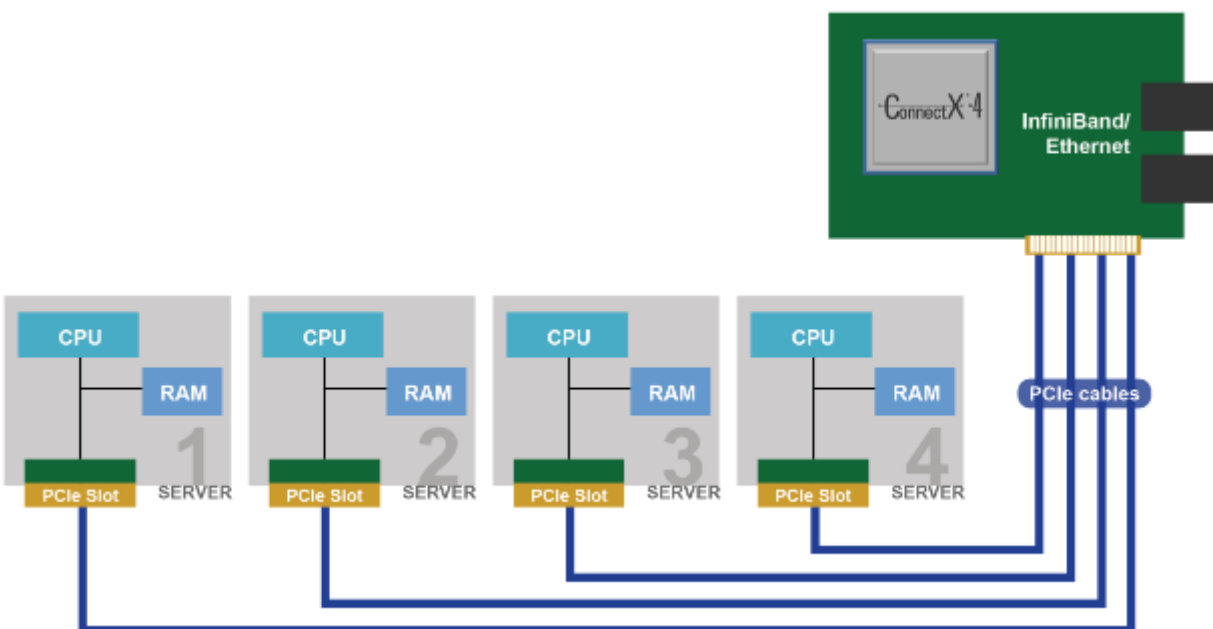
Figure 1: Multi-Host Support Feature on ConnectX-4 EN



2 Unpacking the EVB Kit

Before using your new MCX4MHEVB-ECAA, unpack the system and check to make sure that all the parts have been sent. Check the parts for visible damage that may have occurred during shipping.

Note: If anything is damaged or missing, contact your customer representative immediately.



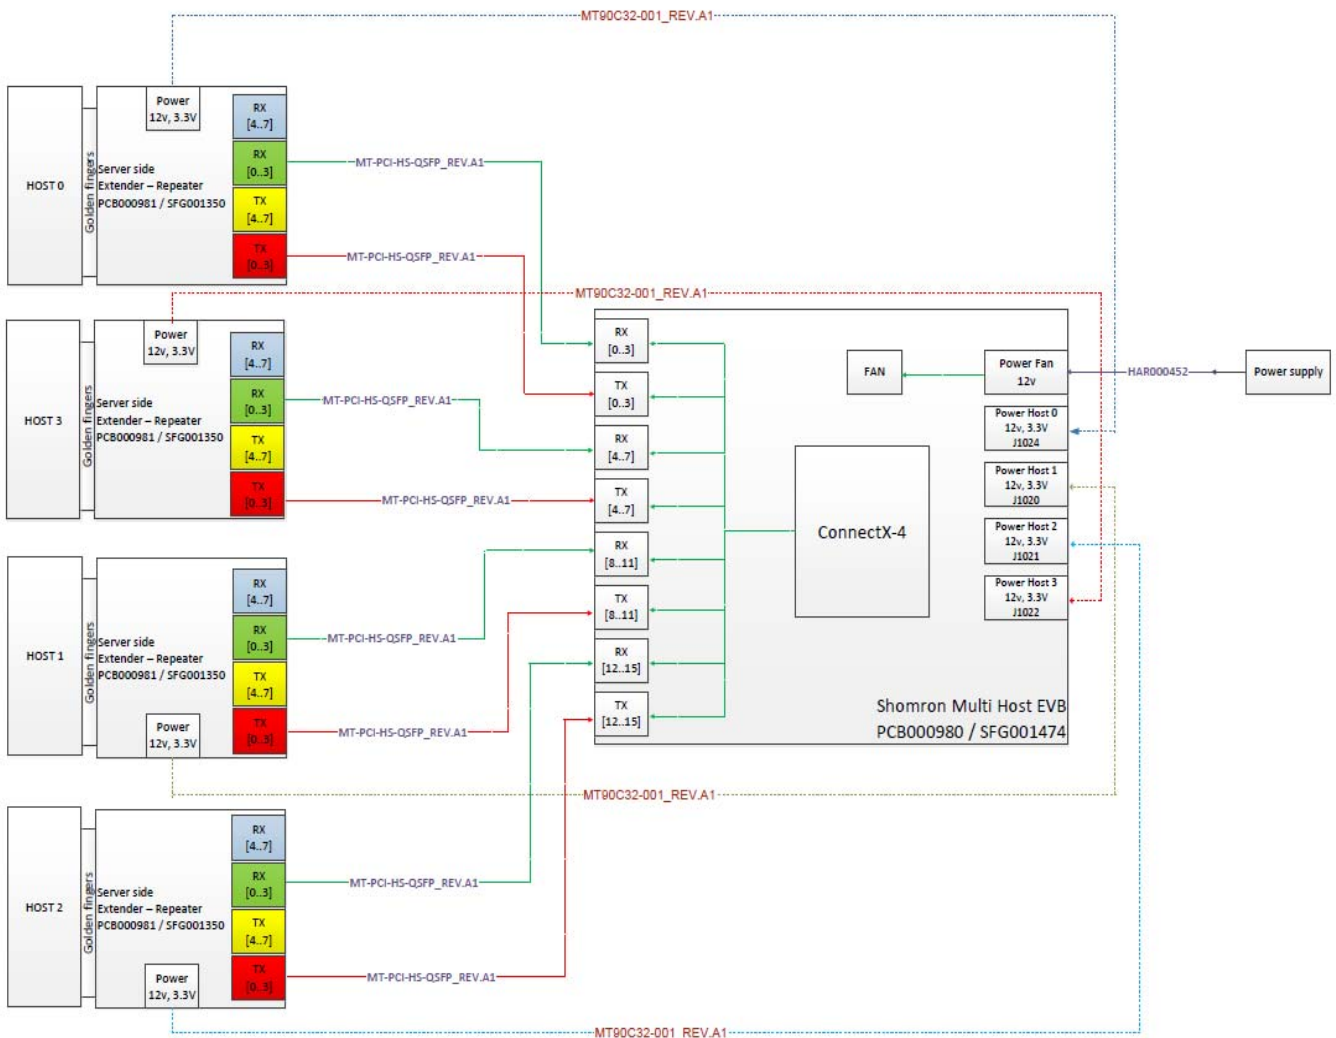
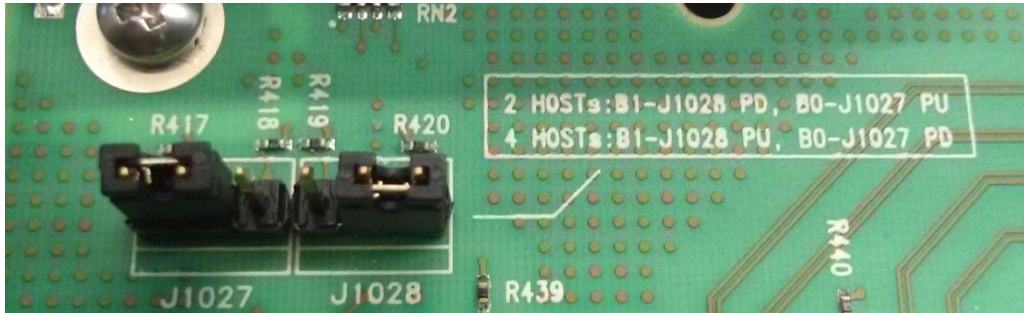
2.1 Kit Components

- 1 Multi-Host EVB (MCX4MHEVB-ECAA)
- 4 x Server-side extender PCI Express (PCIe) repeater cards (PCB00981) - to be connected to independent servers.
- 8 QSFP 2m passive cables (MT-PCI-HS-QSFP_REV.A1) - used to connect to the PCI repeater cards.
- 4x PCIe power cables (MT90C32-001_REV.A1) - to be connected to the power connector of the PCIe repeater cards.
- 1 external 12V power supply (MW155RA1251F01) - to power the board independently (with no need to connect the power from the PCIe repeater cards).
- 1 Harness connecting the power supply and the EVB.

3.2 Four Hosts Configuration

In a four hosts configuration, Multi-Host provides four x4 PCIe lanes per host.

In order to work in the 4 hosts (PCIe 4x per host) configuration, assemble shunt jumpers according to the below.



3.3 Assembly Instructions

1. For each multi-host configuration option, connect PCB000981 (server side extender with repeater) into the slot available on the server.
2. Connect PCB000981 with PCB000980 Multi-Host EVB using QSFP cables MT-PCI-HS-QSFP_REV.A1. Connect the server side extender with repeater to the Multi-Host EVB according to the colored connectivity scheme shown in [Section 3.1, “Two Hosts Configuration”](#) and [Section 3.2, “Four Hosts Configuration”](#).
3. The Multi-Host EVB can also share power rails from four different servers. In order to test this feature, connect the 12V and 3.3V power rails from PCB000981 to PCB000980 using the power cable MT90C32-001_REV.A1.
4. For comfortable work without overload of power cables, it is possible to use the power supply included in the kit. Connect the power supply to the electricity plug from one end and connect it to the adapter harness on the other end. Then connect the harness to J1024 connector on the EVB.
5. When power of the EVB is supplied from only one host, the maximum power, may exceed the 25W per slot (according to the PCIe spec). Instead of connecting to only one host, it is best to connect at least two power sources (i.e. to two hosts) for better power distribution between the hosts.

4 Driver and Firmware

4.1 Linux Driver

For Linux, download and install the latest OpenFabrics Enterprise Distribution (OFED) software package available for ConnectX-4 via the Mellanox web site at: <http://www.mellanox.com> => Products => Software => InfiniBand/VPI Drivers => Linux SW/Driver => Download.

4.2 Windows Driver

For Windows, download and install the latest Mellanox WinOF VPI for Windows software package available for ConnectX-4 via the Mellanox web site at: <http://www.mellanox.com> => Products => Software => InfiniBand/VPI Drivers => Windows SW/Driver => Download.

4.3 Adapter Firmware

In order to receive the latest firmware for this product, please contact Mellanox support.

5 Specifications

5.1 Specifications

Table 2 - MCX4MHEVB-ECAT Specifications Table

Physical	Size: 7.5 in. x10.5 in. (192.58mm x 267mm).
	Connector: QSFP28 InfiniBand/Ethernet (Copper and optical).
Protocol Support	InfiniBand: IBTA v1.2.1 Auto-Negotiation ^a : 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.
	Data Rate: Up to 100Gb/s EDR– InfiniBand Up to 100Gb/s - Ethernet
	PCI Express Gen3: SERDES @ 8.0GT/s, 16 lanes (2.0 and 1.1 compatible).
Power and Environmental	Voltage: 12V, 3.3V
	Typ Power: Passive Cables 23.5W
	Max Power: Passive Cables 25.6W
	Max power available through QSFP+ port: 3.0W
	Temperature: Operational 0°C to 55°C Non-operational 0°C to 70°C
	Humidity: 90% relative humidity ^b
	Air Flow: 600LFM ^c
Regulatory	Safety: IEC/EN 60950-1:2006 ETSI EN 300 019-2-2 IEC 60068-2- 64, 29, 32
	RoHS: RoHS-R6

- a. The ConnectX-4 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another Mellanox InfiniBand product.
- b. For both operational and non-operational states.
- c. Air flow is measured ~1” from the heat sink between the heat sink and the cooling air inlet. The measured air flow is in the case of using a 3.5W QSFP module. In case of lower QSFP module power, a lower air flow can be used. Please advise Mellanox for more information.

5.2 LED Operation

There is a one bicolor LED per port.

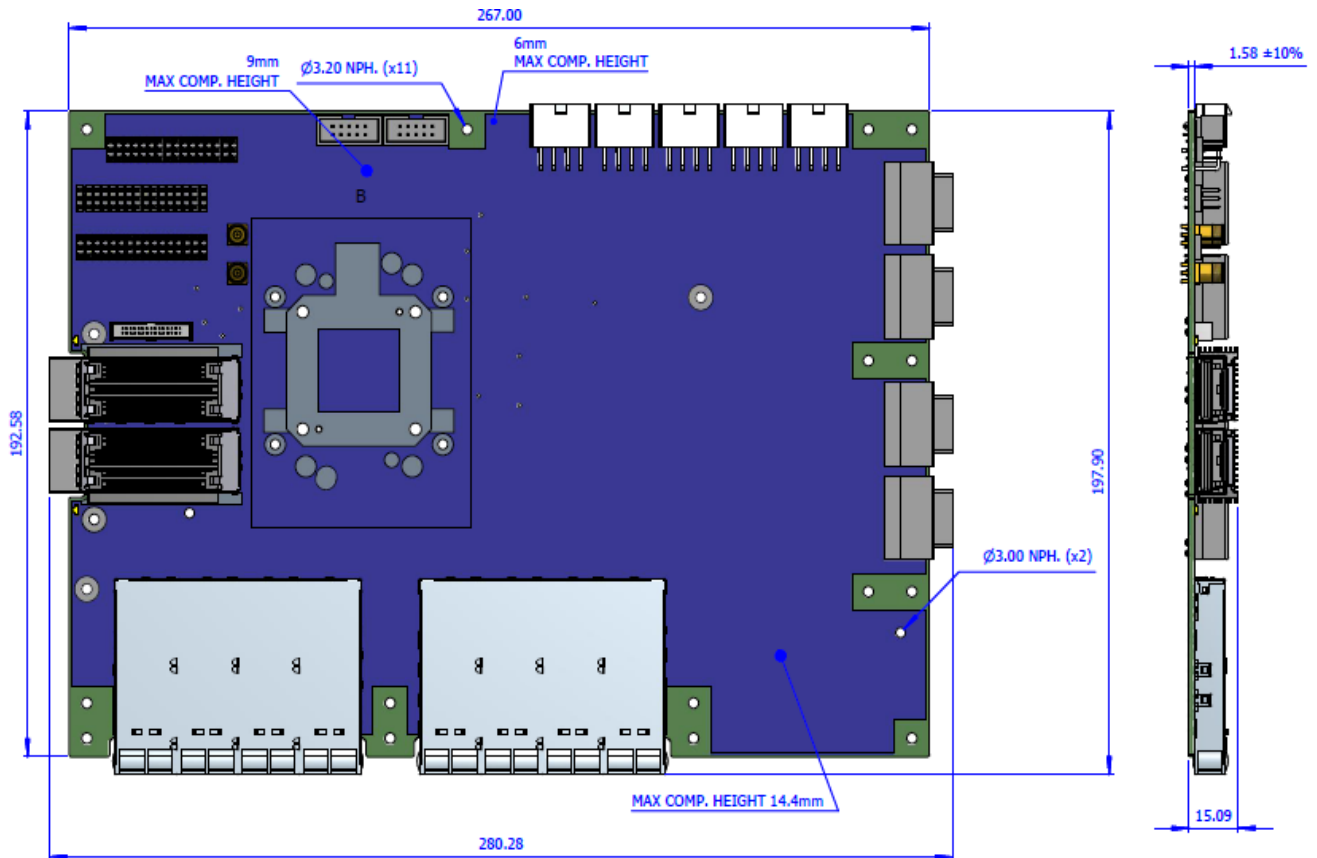
RefDes	LED	Function
D1, D2	Yellow - physical link	<ul style="list-style-type: none"> Constant on indicates a good physical link. Blinking indicates a problem with the physical link. If neither LED is lit, then the physical link has not been established.
	Green - logical (data activity) link	<ul style="list-style-type: none"> A constant yellow indicates a valid logical (data activity) link without data transfer. A blinking yellow indicates a valid logical link with data transfer. If only the green LED is lit and the yellow LED is off, then the logical link has not been established.
D4, D9	Blue	<ul style="list-style-type: none"> Thermal warning in the EQSFP ports.
D3	Yellow	<ul style="list-style-type: none"> Thermal warning - ConnectX-4 reached 105°C.
D6	Red	<ul style="list-style-type: none"> Thermal Shutdown - ConnectX-4 is overheated and reached 120°C.
D8	Green	<ul style="list-style-type: none"> 12V is valid on board.
D7	Green	<ul style="list-style-type: none"> Power good is valid and all power supplies are functioning.
	Red	<ul style="list-style-type: none"> Power good has an error.

5.3 Board Mechanical Drawing and Dimensions



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

5.3.1 Multi-Host Evaluation Board



5.3.2 Server Side Extender with Repeater

