



Connect. Accelerate. Outperform.™

Mellanox ConnectX-4 NATIVE ESX Driver for VMware vSphere 6.0 User Manual

Rev 4.15.2.0

www.mellanox.com

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 NON-INFRINGEMENT 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®, GPUDirect®, InfiniBridge®, InfiniHost®, InfiniScale®, Kotura®, Kotura logo, Mellanox Connect. Accelerate. Outperform logo, Mellanox Federal Systems®, Mellanox Open Ethernet®, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, Open Ethernet logo, PhyX®, ScalableHPC®, SwitchX®, TestX®, The Generation of Open Ethernet logo, UFM®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

CyPU™, ExtendX™, FabricIT™, FPGADirect™, HPC-X™, Mellanox Care™, Mellanox CloudX™, Mellanox NEO™, Mellanox Open Ethernet™, Mellanox PeerDirect™, NVMeDirect™, StPU™, Spectrum™, Switch-IB™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Table of Contents	3
List of Tables	4
Document Revision History	5
About this Manual	6
Chapter 1 Introduction	8
1.1 nmlx5 Driver	8
1.2 Mellanox NATIVE ESX ConnectX-4 Package	8
1.2.1 Software Components	8
1.3 Module Parameters	8
1.3.1 nmlx5 Module Parameters	8
Chapter 2 Installation	10
2.1 Hardware and Software Requirements	10
2.2 Installing Mellanox ConnectX-4 NATIVE ESX Driver for VMware vSphere ..	10
2.3 Removing Mellanox OFED Driver	11
2.4 Loading/Unloading Driver Kernel Modules	11
2.5 Firmware Programming	12
Chapter 3 Troubleshooting	13
3.1 General Related Issues	13
3.2 Ethernet Related Issues	13
3.3 Installation Related Issues	14

List of Tables

Table 1:	Document Revision History	5
Table 2:	Abbreviations and Acronyms	6
Table 3:	Reference Documents	7
Table 4:	nmlx5_core Parameters	9
Table 5:	Software and Hardware Requirements	10
Table 6:	General Related Issues	13
Table 7:	Ethernet Related Issues	13
Table 8:	Installation Related Issues	14

Document Revision History

Table 1 - Document Revision History

Release	Date	Description
Rev 4.15.2.0	September, 2015	Initial release of this MLNX-NATIVE-ESX-ConnectX-4 version

About this Manual

This preface provides general information concerning the scope and organization of this User's Manual.

Intended Audience

This manual is intended for system administrators responsible for the installation, configuration, management and maintenance of the software and hardware of VPI (in Ethernet mode), and Ethernet adapter cards. It is also intended for application developers.

Common Abbreviations and Acronyms

Table 2 - Abbreviations and Acronyms (Sheet 1 of 2)

Abbreviation / Acronym	Whole Word / Description
B	(Capital) 'B' is used to indicate size in bytes or multiples of bytes (e.g., 1KB = 1024 bytes, and 1MB = 1048576 bytes)
b	(Small) 'b' is used to indicate size in bits or multiples of bits (e.g., 1Kb = 1024 bits)
FW	Firmware
HCA	Host Channel Adapter
HW	Hardware
LSB	Least significant <i>byte</i>
lsb	Least significant <i>bit</i>
MSB	Most significant <i>byte</i>
msb	Most significant <i>bit</i>
NIC	Network Interface Card
SW	Software
VPI	Virtual Protocol Interconnect
PR	Path Record
RDS	Reliable Datagram Sockets
SDP	Sockets Direct Protocol
SL	Service Level
MPI	Message Passing Interface
QoS	Quality of Service
ULP	Upper Level Protocol

Table 2 - Abbreviations and Acronyms (Sheet 2 of 2)

Abbreviation / Acronym	Whole Word / Description
vHBA	Virtual SCSI Host Bus adapter
uDAPL	User Direct Access Programming Library

Related Documentation

Table 3 - Reference Documents

Document Name	Description
IEEE Std 802.3ae™-2002 (Amendment to IEEE Std 802.3-2002) Document # PDF: SS94996	Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment: Media Access Control (MAC) Parameters, Physical Layers, and Management Parameters for 10 Gb/s Operation
Firmware Release Notes for Mellanox adapter devices	See the Release Notes PDF file relevant to your adapter device. For further information please refer to the Mellanox website. www.mellanox.com -> Support -> Firmware Download
MFT User Manual	Mellanox Firmware Tools User's Manual. For further information please refer to the Mellanox website. www.mellanox.com -> Products -> Ethernet Drivers -> Firmware Tools
MFT Release Notes	Release Notes for the Mellanox Firmware Tools. For further information please refer to the Mellanox website. www.mellanox.com -> Products -> Ethernet Drivers -> Firmware Tools
VMware vSphere 6.0 Documentation Center	VMware website

1 Introduction

Mellanox ConnectX®-4 NATIVE ESX is a software stack which operates across all Mellanox network adapter solutions supporting up to 100Gb/s Ethernet (ETH) and 2.5 or 5.0 GT/s PCI Express 2.0 and 3.0 uplinks to servers.

The following sub-sections briefly describe the various components of the Mellanox ConnectX®-4 NATIVE ESX stack.

1.1 nmlx5 Driver

`nm1x5` is the low level driver implementation for the ConnectX®-4 adapter cards designed by Mellanox Technologies. ConnectX®-4 adapter cards can operate as an InfiniBand adapter, or as an Ethernet NIC. The ConnectX®-4 NATIVE ESX driver supports Ethernet NIC configurations. To accommodate the supported configurations, the driver consist of `nm1x5_core` module.

`nm1x5_core`

A 10/25/40/50/100GigE driver that handles Ethernet specific functions and plugs into the ESX uplink layer

1.2 Mellanox NATIVE ESX ConnectX-4 Package

1.2.1 Software Components

MLNX-NATIVE-ESX-ConnectX-4 contains the following software components:

- Mellanox Host Channel Adapter Drivers
 - `nm1x5_core` (Ethernet)

1.3 Module Parameters

1.3.1 `nm1x5` Module Parameters

To set `nm1x5` parameters:

```
esxcli system module parameters set -m nm1x5_core -p <parameter>=<value>
```

To show all parameters which were set until now:

```
esxcli system module parameters list -m <module name>
```

Parameters which are not set by the user, remain on default value.

The following sections list the available `nm1x5` parameters.

1.3.1.1 nmlx5_core Parameters

Table 1 - nmlx5_core Parameters

Name	Description	Values
enable_nmlx_debug	Enables debug prints for nmlx	<ul style="list-style-type: none">• 1 - enabled• 0 - disabled Default: 0

2 Installation

This chapter describes how to install and test the Mellanox ConnectX-4 NATIVE ESX package on a single host machine with Mellanox Ethernet adapter hardware installed.

2.1 Hardware and Software Requirements

Table 2 - Software and Hardware Requirements

Requirements	Description
Platforms	A server platform with an adapter card based on one of the following Mellanox Technologies' HCA devices: <ul style="list-style-type: none"> • MT27700 ConnectX®-4 (VPI, EN) (firmware: fw-ConnectX4) • MT27630 ConnectX®-4 LX (EN) (firmware: fw-ConnectX4-LX)
Device ID	For the latest list of device IDs, please visit Mellanox website.
Operating System	ESXi 6.0 operating system.
Installer Privileges	The installation requires administrator privileges on the target machine.

2.2 Installing Mellanox ConnectX-4 NATIVE ESX Driver for VMware vSphere



Please uninstall any previous Mellanox driver packages prior to installing the new version.

➤ **To install the driver:**

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

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

Example:

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

3. Reboot the machine.
4. Verify the driver was installed successfully.

```
# esxcli software vib list | grep mlx
nmlx5-core          4.15.2.0-10EM.600.0.0.2768847    MEL    PartnerSupported    2015-09-17
```



After the installation process, all kernel modules are loaded automatically upon boot.

2.3 Removing Mellanox OFED Driver



Please unload the driver before removing it.

➤ **To remove all the drivers:**

1. Log into the ESXi server with root permissions.
2. List the existing ConnectX-4 NATIVE ESX driver modules. (see [Step 4 in Section 2.2, on page 10](#))
3. Remove each module.

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



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

4. Reboot the server.

2.4 Loading/Unloading Driver Kernel Modules

➤ **To unload the driver:**

```
esxcfg-module -u nmlx5_core
```

➤ **To load the driver:**

```
/etc/init.d/sfcbd-watchdog stop
esxcfg-module nmlx5_core
/etc/init.d/sfcbd-watchdog start
kill -POLL $(cat /var/run/vmware/vmkdevmgr.pid)
```

➤ **To restart the driver:**

```
/etc/init.d/sfcbd-watchdog stop
esxcfg-module -u nmlx5_core
esxcfg-module nmlx5_core
/etc/init.d/sfcbd-watchdog start
kill -POLL $(cat /var/run/vmware/vmkdevmgr.pid)
```

2.5 Firmware Programming

1. Download the VMware bootable binary images v4.1.0 from the [Mellanox Firmware Tools \(MFT\)](#) site.
 - **File:** mft-4.1.0.34-10EM-610.0.0.2770837.x86_64.vib
 - **MD5SUM:** 7b8552c1a22b554c4a280d679dfd528f
2. Install the image according to the steps described in the [MFT User Manual](#).



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

3 Troubleshooting

You may be able to easily resolve the issues described in this section. If a problem persists and you are unable to resolve it yourself please contact your Mellanox representative or Mellanox Support at support@mellanox.com.

3.1 General Related Issues

Table 3 - General Related Issues

Issue	Cause	Solution
The system panics when it is booted with a failed adapter installed.	Malfunction hardware component	<ol style="list-style-type: none"> 1. Remove the failed adapter. 2. Reboot the system.
Mellanox adapter is not identified as a PCI device.	PCI slot or adapter PCI connector dysfunctionality	<ol style="list-style-type: none"> 1. Run <code>lspci</code>. 2. Reseat the adapter in its PCI slot or insert the adapter to a different PCI slot. If the PCI slot confirmed to be functional, the adapter should be replaced.
Mellanox adapters are not installed in the system.	Misidentification of the Mellanox adapter installed	Run the command below to identify the Mellanox adapter installed. <code>lspci grep Mellanox'</code>

3.2 Ethernet Related Issues

Table 4 - Ethernet Related Issues

Issue	Cause	Solution
No link.	Mis-configuration of the switch port or using a cable not supporting link rate.	<ul style="list-style-type: none"> • Ensure the switch port is not down • Ensure the switch port rate is configured to the same rate as the adapter's port
No link with break-out cable.	Misuse of the break-out cable or misconfiguration of the switch's split ports	<ul style="list-style-type: none"> • Use supported ports on the switch with proper configuration. For further information, please refer to the MLNX_OS User Manual. • Make sure the QSFP break-out cable side is connected to the SwitchX.
Physical link fails to negotiate to maximum supported rate.	The adapter is running an outdated firmware.	Install the latest firmware on the adapter.

Table 4 - Ethernet Related Issues

Issue	Cause	Solution
Physical link fails to come up.	The cable is not connected to the port or the port on the other end of the cable is disabled.	Ensure that the cable is connected on both ends or use a known working cable

3.3 Installation Related Issues

Table 5 - Installation Related Issues

Issue	Cause	Solution
Driver installation fails.	The install script may fail for the following reasons: <ul style="list-style-type: none">Failed to uninstall the previous installation due to dependencies being usedThe operating system is not supported	<ul style="list-style-type: none">Uninstall the previous driver before installing the new oneUse a supported operating system and kernel