



# Mellanox ConnectX-4/ConnectX-5 NATIVE ESXi Driver for VMware vSphere 6.7 Release Notes

---

Rev 4.17.14.2

## NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT ("PRODUCT(S)") AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES "ASIS" 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](http://www.mellanox.com)  
Tel: (408) 970-3400  
Fax: (408) 970-3403

© Copyright 2018. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, Accelio®, BridgeX®, CloudX logo, CompustorX®, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, EZchip®, EZchip logo, EZappliance®, EZdesign®, EZdriver®, EZsystem®, GPUDirect®, InfiniHost®, InfiniBridge®, InfiniScale®, LinkX®, Kotura®, Kotura logo, Mellanox CloudRack®, Mellanox CloudXMellanox®, Mellanox Federal Systems®, Mellanox HostDirect®, Mellanox Multi-Host®, Mellanox Open Ethernet®, Mellanox OpenCloud®, Mellanox OpenCloud Logo®, Mellanox PeerDirect®, Mellanox ScalableHPC®, Mellanox StorageX®, Mellanox TuneX®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, NPS®, Open Ethernet logo, PhyX®, PlatformX®, PSIPHY®, SiPhy®, StoreX®, SwitchX®, Tiler®, Tiler logo, TestX®, TuneX®, The Generation of Open Ethernet logo, UFM®, Unbreakable Link®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

For the most updated list of Mellanox trademarks, visit <http://www.mellanox.com/page/trademarks>

## Table of Contents

<b>Table of Contents</b> .....	<b>3</b>
<b>List Of Tables</b> .....	<b>4</b>
<b>Release Update History</b> .....	<b>5</b>
<b>Chapter 1 Overview</b> .....	<b>6</b>
1.1 Content of MLNX-NATIVE-ESX Driver Package.....	6
1.2 Supported HCAs Firmware Versions.....	6
1.3 Tested Hypervisors in Paravirtualized and SR-IOV Environments .....	7
<b>Chapter 2 Changes and New Features in Rev 4.17.14.2</b> .....	<b>8</b>
<b>Chapter 3 Known Issues</b> .....	<b>9</b>
<b>Chapter 4 Bug Fixes History</b> .....	<b>11</b>
<b>Chapter 5 Change Log History</b> .....	<b>12</b>

## List Of Tables

Table 1:	Release Update History . . . . .	5
Table 2:	Supported Uplinks to Servers . . . . .	6
Table 3:	Supported HCAs Firmware Versions . . . . .	6
Table 4:	Tested Hypervisors in Paravirtualized and SR-IOV Environments . . . . .	7
Table 5:	Changes and New Features . . . . .	8
Table 6:	Known Issues . . . . .	9
Table 7:	Fixed Bugs List . . . . .	11
Table 8:	Change Log History . . . . .	12

## Release Update History

*Table 1 - Release Update History*

Release	Date	Description
Rev 4.17.14.2	July 29, 2018	Initial release of this MLNX-NATIVE-ESXi-ConnectX-4/ ConnectX-4 Lx/ConnectX-5 driver version

# 1 Overview

These are the release notes of Mellanox ConnectX-4/ConnectX-5 NATIVE ESXi Driver for VMware vSphere 6.7. Mellanox ConnectX-4/ConnectX-5 NATIVE ESXi Driver for VMware vSphere 6.7 supports the following uplinks to servers

**Table 2 - Supported Uplinks to Servers**

Version	OS	Uplink Speed
4.17.14.2	ESXi 6.7	10/25/40/50/100GbE

## 1.1 Content of MLNX-NATIVE-ESX Driver Package

The MLNX-NATIVE-ESX driver package is distributed as an offline bundle (.zip file) and contains:

- **ESXi 6.7:**  
MLNX-NATIVE-ESX-ConnectX-4-5\_4.17.14.2-10EM-670.0.0.8169922.zip - Hypervisor bundle for ESXi 6.7 contains the following kernel modules:
  - nmlx5\_core
  - nmlx5\_rdma

## 1.2 Supported HCAs Firmware Versions

MLNX-NATIVE-ESX Rev 4.17.14.2 supports the following Mellanox Ethernet HCA:

**Table 3 - Supported HCAs Firmware Versions**

HCAs	Recommended Firmware Rev.	Additional Firmware Rev.
ConnectX-4	12.23.1020	12.22.1002
ConnectX-4 Lx	14.23.1020	14.22.1002
ConnectX-5	16.23.1020	16.22.1002
ConnectX-5 Ex	16.23.1020	16.22.1002

For the latest firmware versions, visit:

- [http://www.mellanox.com/page/products\\_dyn?product\\_family=29](http://www.mellanox.com/page/products_dyn?product_family=29)

### 1.3 Tested Hypervisors in Paravirtualized and SR-IOV Environments

*Table 4 - Tested Hypervisors in Paravirtualized and SR-IOV Environments*

Tested Hypervisors	HCAs	Guest Operating System
SR-IOV	ConnectX-4/ConnectX-4 Lx ConnectX-5/ConnectX-5 Ex	Windows Server 2012 R2
		Windows Server 2016 DC
		RedHat 6.6
		RedHat 6.7
		RedHat 6.8
		RedHat 7.0
		RedHat 7.1
		RedHat 7.2
		RedHat 7.3
		RedHat 7.4
		RedHat 7.5
		SLES 11.3
		SLES 12
		SLES 12.1
		SLES 12.2
		SLES 12.3
SLES 12.4		
Paravirtualized <sup>a</sup> (Ethernet Only)	ConnectX-4/ConnectX-4 Lx ConnectX-5/ConnectX-5 Ex	Windows Server 2012 R2
		RedHat 7.2
		RedHat 7.5

a. Paravirtualized RDMA is only tested Linux OSes.

## 2 Changes and New Features in Rev 4.17.14.2

*Table 5 - Changes and New Features*

Feature/Change	Description
<b>Virtualization</b>	Removed a VF driver limitation. Now the driver can support the maximum number of VFs supported by the firmware.
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”</a> , on page 11



### 3 Known Issues

The following is a list of general limitations and known issues of the various components of this MLNX-NATIVE-ESX release.

**Table 6 - Known Issues (Sheet 1 of 2)**

Internal Ref.	Description
1446060	<b>Description:</b> Although the max_vfs module parameter range is "0-128", due to firmware limitations, the following are the supported VFs per single port devices: <ul style="list-style-type: none"> <li>• ConnectX-4: up to 127</li> <li>• ConnectX-5: up to 63</li> </ul>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> SR-IOV, VFs per port
	<b>Discovered in Version:</b> 4.17.14.2
1340275	<b>Description:</b> ECN tunable parameter initialAlphaValue for the Reaction Point protocol cannot be modified.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> nmlx5 ecn nmlxcli
	<b>Discovered in Version:</b> 4.17.13.8
1340255	<b>Description:</b> ECN statistic counters accumulatorsPeriod and ecnMarkedRocePackets display wrong values and cannot be cleared.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> nmlx5 ecn nmlxcli
	<b>Discovered in Version:</b> 4.17.13.8
-	<b>Description:</b> The hardware can offload only up to 256B of headers.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Hardware offload
	<b>Discovered in Version:</b> 4.17.13.8
781277	<b>Description:</b> The "esxcli network sriovnic vf stats" command is not supported. When running this command on a vmknic, a failure message is displayed.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> esxcli SR-IOV
	<b>Discovered in Version:</b> 4.17.13.8
858972	<b>Description:</b> Setting the "Allow Guest MTU Change" option in vSphere Client is currently not functional. Although guest MTU changes in SR-IOV are allowed, they do not affect the port's MTU and the guest's MTU remains the same as the PF MTU.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> MTU, SR-IOV

**Table 6 - Known Issues (Sheet 2 of 2)**

Internal Ref.	Description
746100	<b>Description:</b> The 'esxcli mellanox uplink link info -u <vmnic_name>' command reports the 'Auto negotiation' capability always as 'true'.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> 'Auto negotiation' capability
1068621	<b>Description:</b> SMP MADs (ibnetdiscover, sminfo, iblinkinfo, smpdump, ibqueryerr, ibdiagnet and smpquery) are not supported on the VFs.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> SMP MADs
778371	<b>Description:</b> Wake-on-LAN does not notify when invalid parameters are provided.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> WoL
778572	<b>Description:</b> Nested ESXi might not function properly.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Nested ESXi
765008	<b>Description:</b> Device RSS fails to hash traffic to sufficient RX rings with Broadcast traffic.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> RSS, RX rings
852883	<b>Description:</b> In stress condition 'Watchdog' may appear, leading to uplink going up and down.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> uplink, watchdog
784293	<b>Description:</b> VGT traffic over VXLAN interfaces is currently not supported.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VGT traffic, VXLAN

## 4 Bug Fixes History

Table 7 lists the bugs fixed in this release.

**Table 7 - Fixed Bugs List**

Internal Ref.	Description
1358381	<p><b>Description:</b> Fixed an issue that prevented ESXi from being discovered via the CDP protocol on ConnectX-4 Lx adapter cards.</p> <p><b>Keywords:</b> CDP protocol, ConnectX-4 Lx</p> <p><b>Discovered in Release:</b> 4.17.13.8</p> <p><b>Fixed in Release:</b> 4.17.14.2</p>

## 5 Change Log History

**Table 8 - Change Log History (Sheet 1 of 3)**

Feature/Change	Description
<b>Rev. 4.17.13.8</b>	
<b>Explicit Congestion Notification (ECN)</b>	Explicit Congestion Notification (ECN) is an extension to the Internet Protocol and to the Transmission Control Protocol. ECN allows end-to-end notification of network congestion without dropping packets. To configure ECN behavior, download the nmlxcli tool from the Mellanox site. For further information, refer to the User Manual section <i>Explicit Congestion Notification (ECN)</i> .
<b>Dynamic RSS</b>	Improves network performance by allowing OS Load Balancer better RSS RX queue utilization during heavy traffic of the same type. For further information, refer to the User Manual section <i>Dynamic RSS</i> .
<b>Multiple RSS Engines</b>	Improves network performance by exposing multiple RSS RX queues to the hypervisor network stack. For further information, refer to the User Manual section <i>Multiple RSS Engines</i> .
<b>Packet Capture Utility</b>	Packet Capture utility duplicates all traffic, including RDMA, in its raw Ethernet form (before stripping) to a dedicated "sniffing" QP, and then passes it to an ESX drop capture point. It allows gathering of Ethernet and RoCE bidirectional traffic via pktcap-uw and viewing it using regular Ethernet tools, e.g. Wireshark To enable/disable packet capture, download the nmlxcli tool from the Mellanox site.  For further information, refer to the User Manual section Packet Capture Utility.
<b>SR-IOV max_vfs module parameter Type Modification</b>	Changed the type of the SR-IOV max_vfs module parameter from a single integer value to an array of unsigned integers. For further information, refer to the User Manual.
<b>InfiniBand SR-IOV</b>	Enables the creation of InfiniBand virtual functions, allowing the guests to operate over an InfiniBand fabric.
<b>DCBX Negotiation Support for PFC</b>	PFC port configuration can now be auto-negotiated with switches that support the DCBX protocol.
<b>ESXi CLI</b>	Added ESXi CLI support for ESXi 6.7
<b>Adapter Cards</b>	Added support for ConnectX-5/ConnectX-5 Ex adapter cards. <b>Note:</b> ConnectX-5/ConnectX-5 Ex cards are currently at beta level.
<b>Geneve Stateless Offload</b>	Geneve network protocol is encapsulated into IP frame (L2 tunneling). Encapsulation is suggested as a means to alter the normal IP routing for datagrams, by delivering them to an intermediate destination that would otherwise not be selected based on the (network part of the) IP Destination Address field in the original IP header.

**Table 8 - Change Log History (Sheet 2 of 3)**

Feature/Change	Description
<b>Remote Direct Memory Access (RDMA)</b>	Remote Direct Memory Access (RDMA) is the remote memory management capability that allows server-to-server data movement directly between application memory without any CPU involvement. <b>Note:</b> It is recommended to use RoCE with PFC enabled in driver and network switches. For how to enable PFC in the driver see section <i>Priority Flow Control (PFC)</i> in the User Manual.
<b>Set Link Speed</b>	Enables you to set the link speed to a specific link speed supported by ESXi. For further information, see the User Manual section “ <i>Set Link Speed</i> ”.
<b>Priority Flow Control (PFC)</b>	Applies pause functionality to specific classes of traffic on the Ethernet link. For further information, see the User Manual section “ <i>Priority Flow Control (PFC)</i> ”.
<b>NetQ RSS</b>	Allows the user to configure multiple hardware queues backing up the single RX queue. NetQ RSS improves vMotion performance and multiple streams of IPv4/IPv6 TCP/UDP/IPSEC bandwidth over single interface between the Virtual Machines. For further information, see the User Manual section “ <i>NetQ RSS</i> ”.
<b>Default Queue RSS (DRSS)</b>	Allows the user to configure multiple hardware queues backing up the default RX queue. DRSS improves performance for large scale multicast traffic between hypervisors and Virtual Machines interfaces. For further information, see the User Manual section “ <i>Default Queue Receive Side Scaling (DRSS)</i> ”.
<b>SR-IOV</b>	Single Root IO Virtualization (SR-IOV) is a technology that allows a physical PCIe device to present itself multiple times through the PCIe bus. Support for up to 8 ConnectX-4 ports and up to 16 VFs. For further information, refer to the User Manual
<b>RX/TX Ring Resize</b>	Allows the network administrator to set new RX\TX ring buffer size.
<b>VXLAN Hardware Stateless Offloads for ConnectX®-4</b>	VXLAN hardware offload enables the traditional offloads to be performed on the encapsulated traffic.
<b>NetDump</b>	Enables a host to transmit diagnostic information via the network to a remote netdump service, which stores it on disk. Network-based core-dump collection can be configured in addition to or instead of disk-based coredump collection.
<b>NetQueue</b>	NetQueue is a performance technology in VMware ESXi that significantly improves performance in Ethernet virtualized environments.
<b>Wake-on-LAN</b>	Allows a network administrator to remotely power on a system or to wake it up from sleep mode
<b>Hardware Offload</b>	<ul style="list-style-type: none"> <li>• Large Send Offload (TCP Segmentation Offload)</li> <li>• RSS (Device RSS)</li> </ul>
<b>Hardware Capabilities</b>	<ul style="list-style-type: none"> <li>• Multiple Tx/Rx rings</li> <li>• Fixed Pass-Through</li> <li>• Single/Dual port</li> <li>• MSI-X</li> </ul>

**Table 8 - Change Log History (Sheet 3 of 3)**

Feature/Change	Description
<b>Ethernet Network</b>	<ul style="list-style-type: none"><li>• TX/RX checksum</li><li>• Auto moderation and Coalescing</li><li>• VLAN stripping offload</li></ul>