



Mellanox ConnectX-4/ConnectX-5 ESXi 6.7 Inbox Driver Release Notes

Rev 4.17.9.12-vmw

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) 9703400
Fax: (408) 9703403

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

Mellanox®, Mellanox logo, Accelio®, BridgeX®, CloudX logo, CompustorX®, ConnectIB®, 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 MuHost®, 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

Table of Contents	3
List Of Tables	4
Release Update History	5
Chapter 1 Overview	6
1.1 Supported HCAs Firmware Versions	6
Chapter 2 Changes and New Features in Rev 4.17.9.12-vmw	7
Chapter 3 Known Issues	9

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:	Changes and New Features.....	7
Table 5:	Known Issues.....	9

Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 4.17.9.12-vmw	July 26, 2018	Initial version of this release

1 Overview

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

Table 2 - Supported Uplinks to Servers

Version	OS	Uplink Speed
Rev 4.17.9.12-vmw	ESXi 6.7	10/25/40/50/100GbE

1.1 Supported HCAs Firmware Versions

Mellanox ConnectX-4/ConnectX-5 ESXi 6.7 Inbox Driver Rev 4.17.9.12-vmw supports the following Mellanox Ethernet HCA:

Table 3 - Supported HCAs Firmware Versions

HCAs	Recommended Firmware Rev.
ConnectX-5	16.22.1002
ConnectX-5 Ex	16.22.1002
ConnectX-4	12.22.1002
ConnectX-4 Lx	14.22.1002

For the latest firmware versions, visit:

- http://www.mellanox.com/page/products_dyn?product_family=29

2 Changes and New Features in Rev 4.17.9.12-vmw

Table 4 - Changes and New Features

Feature/Change	Description
SR-IOV max_vfs module parameter Type Modification	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.
DCBX Negotiation Support for PFC	PFC port configuration can now be auto-negotiated with switches that support the DCBX protocol.
ESXi CLI	ESXi CLI support for ESXi 6.7
Geneve Stateless Offload	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.
Remote Direct Memory Access (RDMA)	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. Note: 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.
Set Link Speed	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> ”.
Priority Flow Control (PFC)	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> ”.
NetQ RSS	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> ”.
Default Queue RSS (DRSS)	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> ”.
SR-IOV	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
RX/TX Ring Resize	Allows the network administrator to set new RX\TX ring buffer size.

Table 4 - Changes and New Features

Feature/Change	Description
VXLAN Hardware Stateless Offloads	Added support for VXLAN hardware offload. VXLAN hardware offload enables the traditional offloads to be performed on the encapsulated traffic. With ConnectX®-3 Pro, data center operators can decouple the overlay network layer from the physical NIC performance, thus achieving native performance in the new network architecture.
NetDump	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 core dump collection.
NetQueue	NetQueue is a performance technology in VMware ESXi that significantly improves performance in Ethernet virtualized environments.
Wake-on-LAN (WoL)	Allows a network administrator to remotely power on a system or to wake it up from the sleep mode.
Hardware Offload	<ul style="list-style-type: none"> • Large Send Offload (TCP Segmentation Offload) • RSS (Device RSS)
Hardware Capabilities	<ul style="list-style-type: none"> • Multiple Tx/Rx rings • Fixed Pass-Through • Single/Dual port • MSI-X
Ethernet Network	<ul style="list-style-type: none"> • TX/RX checksum • Auto moderation and Coalescing • VLAN stripping offload

3 Known Issues

The following is a list of general limitations and known issues of the various components of this Mellanox ConnectX-4/ConnectX-5 ESXi 6.7 Inbox Driver release.

Table 5 - Known Issues

Internal Ref.	Description
-	<p>Description: Adapter cards that come with a pre-configured link type as InfiniBand cannot be controlled by the driver and cannot be seen by MFT. Thus its link type cannot be changed.</p> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Unload the driver. <code>unload nmlx5_core module</code> 2. Make the device visible to MFT by loading the driver in a recovery mode. <code>vmkload_mod nmlx5_core mst_recovery=1</code> <code>kill the devmgr</code> 3. Check the device available on your machine. <code>/opt/mellanox/bin/mst status</code> 4. Change the link type to Ethernet using MFT. <code>opt/mellanox/bin/mlxconfig -d mt4115_pciconf0 set LINK_TYPE_P1=2</code> <code>LINK_TYPE_P2=2</code> 5. Power Cycle the host. <p>Keywords: Link type, InfiniBand, MFT</p>