



NVIDIA ConnectX-4 onwards NICs NATIVE ESXi 6.5.Driver for VMware vSphere Release Notes v4.16.71.1

Table of Contents

Overview.....	3
Release Notes Revision History.....	3
General Information.....	4
Content of MLNX-NATIVE-ESX Driver Package.....	4
Supported HCAs Firmware Versions.....	4
Tested Hypervisors in Paravirtualized and SR-IOV Environments.....	5
Changes and New Features	6
Bug Fixes in the Version.....	7
Known Issues	8
Change Log History	12
Bug Fixes History	15

Overview

NVIDIA® Networking native ESXi drivers enable industry-leading performance and efficiency as non-virtualized environments using hardware offloads such as RDMA over Converged Ethernet (RoCE) on VMware vSphere. NVIDIA® ConnectX®-4 onwards deliver 10/25/40/50 and 100GbE network speeds with ESXi 6.5 onwards, allowing the highest port rate on ESXi today.

Release Notes Revision History

Date	Revision	Description
February 16 2020	4.16.71.1	Initial release of this Release Notes version.

General Information

These are the release notes of NVIDIA® ConnectX®-4 onwards adapter cards for NATIVE ESXi Driver for VMware vSphere 6.5. This driver supports the following uplinks to servers.

Version	OS	Uplink Speed
4.16.71.1	ESXi 6.5	10/25/40/50/100GbE

Content of MLNX-NATIVE-ESX Driver Package

ESXi 6.5:

Mellanox-nmlx5_4.16.71.1-10EM.650.0.0.4598673.zip - Hypervisor bundle for ESXi 6.5 contains the following kernel modules:

- nmlx5_core
- nmlx5_rdma

Supported HCAs Firmware Versions

MLNX-NATIVE-ESX Rev 4.16.71.1 supports the following Mellanox Ethernet HCA and their corresponding firmware version:

HCAs	Minimal Recommended Firmware Rev.
ConnectX-4	12.28.2006
ConnectX-4 Lx	14.29.1016
ConnectX-5 / ConnectX-5 Ex	16.29.1016
ConnectX-6	20.29.1016
ConnectX-6 Dx	22.29.1016
ConnectX-6 Lx	26.29.1016

For the latest firmware versions, visit: <https://www.mellanox.com/support/firmware/firmware-downloads>

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 ConnectX-6 ConnectX-6 Dx ConnectX-6 Lx	Windows Server 2016 DC RedHat 8.0 RedHat 7.5 RedHat 7.3 RedHat 6.10 RedHat 6.3 SLES 12 SP4 SLES 12 SP3
Paravirtualized ^{a,b} (Ethernet Only)	ConnectX-4 ConnectX-4 Lx ConnectX-5/ConnectX-5 Ex ConnectX-6 ConnectX-6 Dx ConnectX-6 Lx	Windows Server 2016 DC RedHat 8.0 RedHat 7.5 RedHat 7.3 RedHat 6.10 RedHat 6.3 SLES 12 SP4 SLES 12 SP3

a. Paravirtualized RDMA is supported only in Linux Operating Systems and in this release it was tested for RedHat 7.5 only.

b. Paravirtualized RDMA is supported in ESXi 6.5U1 build and above.

Changes and New Features

Feature/Change	Description
4.16.71.1	
Adapter Cards	Added support for ConnectX-6 Lx devices.
RDMA 10K QP support	Scaled support for up to 10K connections over RDMA networks.
SR-IOV InfiniBand	SR-IOV InfiniBand is at GA level.
supported_num_ports module parameter	Updated the supported_num_ports default value to 1 to lower memory constraints. Note: The user must set a value corresponding to the amount of ports installed in the system.
sriov_mc_isolation	Added sriov_mc_isolation module parameter to isolate multicast traffic to SR-IOV interfaces. Default value is OFF.

Bug Fixes in the Version

The table below lists the bugs fixed in this release. For older issues, please refer to [Bug Fixes History](#).

Internal Ref.	Description
-	Description: ConnectX-6 currently supports up to 100GbE speed link.
	Keywords: ConnectX-6, 100GbE
	Discovered in Version: 4.16.70.1
	Fixed in Release: 4.16.71.1
1712298	Description: Live unload of the driver is not supported. Doing so may cause a PSOD if the max_vfs parameter is set.
	Keywords: Driver load
	Discovered in Version: 4.16.70.1
	Fixed in Release: 4.16.71.1
2120216	Description: The maximum number of established active RDMA connections (QPs) is currently 5000.
	Keywords: QPs, RDMA
	Discovered in Version: 4.16.70.1
	Fixed in Release: 4.16.71.1
2130911	Description: Setting ETS value to 0 may cause WQE timeout.
	Keywords: ETS, QOS
	Discovered in Version: 4.16.70.1
	Fixed in Release: 4.16.71.1

Known Issues

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

Internal Ref.	Description
2429623	Description: Enabling sriov_mc_isolation module parameter may result in vmknic and emulated NICs multicast and IPv6 traffic loss.
	Workaround: Unset or set the module parameter to 0.
	Keywords: Multicast, IPv6, SR-IOV
	Discovered in Version: 4.16.71.1
2139469	Description: 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.
	Workaround: N/A
	Keywords: MTU, SR-IOV
	Discovered in Version: 4.16.71.1
1340255	Description: ECN statistic counters accumulatorsPeriod and ecnMarkedRocePackets display wrong values and cannot be cleared.
	Workaround: N/A
	Keywords: nmlx5 ecn nmlxcli
	Discovered in Version: 4.16.71.1
1340275	Description: ECN tunable parameter initialAlphaValue for the Reaction Point protocol cannot be modified.
	Workaround: N/A
	Keywords: nmlx5 ecn nmlxcli
	Discovered in Version: 4.16.71.1
2430662	Description: ConnectX-6 Dx speed remains zero after port goes down and reboot is performed.
	Workaround: Turn the down and then up again
	Keywords: ConnectX-6 Dx, link speed
	Discovered in Version: 4.16.71.1

Internal Ref.	Description
1446060	Description: 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"> • ConnectX-4 / ConnectX-5: up to 127
	Workaround: N/A
	Keywords: SR-IOV, VFs per port
	Discovered in Version: 4.16.14.2
1340275	Description: ECN tunable parameter initialAlphaValue for the Reaction Point protocol cannot be modified.
	Workaround: N/A
	Keywords: nmlx5 ecn nmlxcli
	Discovered in Version: 4.16.13.5
1340255	Description: ECN statistic counters accumulatorsPeriod and ecnMarkedRocePackets display wrong values and cannot be cleared.
	Workaround: N/A
	Keywords: nmlx5 ecn nmlxcli
	Discovered in Version: 4.16.13.5
-	Description: The hardware can offload only up to 256B of headers.
	Workaround: N/A
	Keywords: Hardware offload
	Discovered in Version: 4.16.10.3
685558	Description: There is no traffic between PV and SR-IOV VF connected to different ports on the same HCA. This issue is applicable to ESXi 6.5 & ESXi 6.5 UP1. The issue is solved in ESXi 6.5 UP2.
	Workaround: N/A
	Keywords: PV, SR-IOV VF, HCA
858972	Description: 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.
	Workaround: N/A
	Keywords: MTU, SR-IOV

Internal Ref.	Description
-	<p>Description: Geneve options length support is limited to 56B. Received packets with options length bigger than 56B are dropped.</p> <p>WA: N/A</p> <p>Keywords: Geneve</p>
910292	<p>Description: Running with ConnectX-4/ConnectX-4 Lx older firmware versions, might result in the following internal firmware errors:</p> <ul style="list-style-type: none"> • Device health compromised • synd 0x1: firmware internal error • extSync 0x94ee <p>Workaround: Upgrade your firmware to the latest version 12.17.2020/14.17.2020</p> <p>Keywords: Firmware</p>
746100	<p>Description: The 'esxcli mellanox uplink link info -u <vmnic_name>' command reports the 'Auto negotiation' capability always as 'true'.</p> <p>Workaround: N/A</p> <p>Keywords: 'Auto negotiation' capability</p>
1072640	<p>Description: ESXi v4.16.10.3 cannot updated from v4.16.8.8 (GA) or from the Inbox driver using the "esxcli software vib update" command.</p> <p>Workaround: To update it, run the "esxcli software vib install" command.</p> <p>Keywords: Driver update</p>
1068621	<p>Description: SMP MADs (ibnetdiscover, sminfo, iblinkinfo, smpdump, ibqueryerr, ibdiagnet and smpquery) are not supported on the VFs.</p> <p>Workaround: N/A</p> <p>Keywords: SMP MADs</p>
778371	<p>Description: Wake-on-LAN does not notify when invalid parameters are provided.</p> <p>Workaround: N/A</p> <p>Keywords: WoL</p>
778572	<p>Description: Nested ESXi might not function properly.</p> <p>Workaround: N/A</p> <p>Keywords: Nested ESXi</p>
765008	<p>Description: Device RSS fails to hash traffic to sufficient RX rings with Broadcast traffic.</p>

Internal Ref.	Description
	Workaround: N/A
	Keywords: RSS, RX rings
852883	Description: In stress condition 'Watchdog' may appear, leading to uplink going up and down.
	Workaround: N/A
	Keywords: uplink, watchdog

Change Log History

Feature/Change	Description
4.16.70.1	
Power Limitation	An event will be sent to notify the administrator if the power required by the network adapter is higher than that available on the PCIe slot.
Differentiated Services Code Point (DSCP)	Added support for trusting Differentiated Services Code Point (DSCP) and setting default value for RoCE traffic.
SR-IOV VF Counters	Added a new counter that enables the user to query per Virtual Function counters.
RX Counters	Added the RX out-of-buffer counter to indicate any lack of software receive buffers.
SR-IOV	SR-IOV InfiniBand is at beta level.
Rev. 4.16.13.5	
Explicit Congestion Notification (ECN)	<p>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.</p> <p>To configure ECN behavior, download the nmlxcli tool from the Mellanox site.</p> <p>For further information, refer to the User Manual section <i>Explicit Congestion Notification (ECN)</i>.</p>
Bug Fixes	See " Bug Fixes History " section.
Rev. 4.16.12.12	
Packet Capture Utility	<p>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.</p> <p>It allows gathering of Ethernet and RoCE bidirectional traffic via pktcap-uw and viewing it using regular Ethernet tools, e.g. Wireshark</p> <p>To enable/disable packet capture, download the nmlxcli tool from the Mellanox site.</p> <p>For further information, refer to the User Manual section Packet Capture Utility.</p>
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.
Bug Fixes	See " Bug Fixes History " section.
Rev. 4.16.10.3	

InfiniBand SR-IOV	Enables the creation of InfiniBand virtual functions, allowing the guests to operate over an InfiniBand fabric.
ESXi CLI	Added ESXi CLI support for ESXi 6.5
Rev. 4.16.7.8	
Adapter Cards	Added support for ConnectX-5/ConnectX-5 Ex adapter cards. Note: ConnectX-5/ConnectX-5 Ex cards are currently at beta level.
Rev. 4.16.7.8	
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.

	<p>Support for up to 8 ConnectX-4 ports and up to 16 VFs.</p> <p>For further information, refer to the User Manual</p>
RX/TX Ring Resize	Allows the network administrator to set new RX\TX ring buffer size.
VXLAN Hardware Stateless Offloads for ConnectX®-4	VXLAN hardware offload enables the traditional offloads to be performed on the encapsulated traffic.
NetDump	Enables a host to transmit diagnostic information via the network to a remote netdump service, which stores it on disk. Network-based coredump collection can be configured in addition to or instead of disk- based coredump collection.
NetQueue	NetQueue is a performance technology in VMware ESXi that significantly improves performance in Ethernet virtualized environments.
Wake-on-LAN	Allows a network administrator to remotely power on a system or to wake it up from 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

Bug Fixes History

The table below lists the bugs fixed in this release.

Internal Ref.	Description
781277	Description: The "esxcli network sriovnic vf stats" command is not supported. When running this command on a vmknics, a failure message is displayed.
	Keywords: esxcli SR-IOV
	Discovered in Version: 4.6.10.3
	Fixed in Release: 4.16.70.1
1358381	Description: Fixed an issue that prevented ESXi from being discovered via the CDP protocol on ConnectX-4 Lx adapter cards.
	Keywords: CDP protocol, ConnectX-4 Lx
	Discovered in Release: 4.16.12.12
	Fixed in Release: 4.16.14.2
1253564	Description: Disabled multicast loopback to avoid a scenario that prevented MAC learning in some configurations.
	Keywords: MAC, multicast loopback
	Discovered in Release: 4.16.12.12
	Fixed in Release: 4.16.13.5
958154	Description: NetQ RSS for encapsulated traffic is currently not supported. Encapsulated traffic (VXLAN/Geneve) directed to NetQ RSS queue will not be distributed through all queues' channels, thus will not utilize the RSS feature.
	Note: It is highly recommended to avoid requesting RSS for encapsulated interfaces, i.e. refrain from defining the following in the VM configuration file: <iface_name>.pnicsFeatures=4
	Keywords: NetQ RSS, encapsulated traffic
	Discovered in Release: 4.16.8.8
	Fixed in Release: 4.16.12.12
698142/637104	Description: Traffic loss of large packets might occur after MTU change.
	Keywords: MTU, Traffic loss
	Discovered in Release: 4.16.7.8

Internal Ref.	Description
	Fixed in Release: 4.16.10.3
846359	Description: Fixed an issue which caused the adapter card to get stuck in Down state after setting the ring size to 8192.
	Keywords: Ring size
	Discovered in Release: 4.16.7.8
	Fixed in Release: 4.16.8.8

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 (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 Mellanox Technologies Ltd. and/or NVIDIA Corporation 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. For the complete and most updated list of Mellanox trademarks, visit <http://www.mellanox.com/page/trademarks>

Copyright

© 2021 Mellanox Technologies Ltd. All rights reserved.