



# **NVIDIA ConnectX-4 onwards NICs NATIVE ESXi Driver for VMware vSphere Release Notes v4.21.71.101**

# Table of Contents

<b>Overview .....</b>	<b>3</b>
Release Notes Revision History .....	3
<b>General Information .....</b>	<b>4</b>
Content of MLNX-NATIVE-ESX Driver Package.....	4
Supported HCAs Firmware Versions .....	4
Tested Hypervisors in Paravirtualized and SR-IOV Environments .....	5
<b>Changes and New Features.....</b>	<b>6</b>
<b>Known Issues.....</b>	<b>7</b>

---

# 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 10/25/40/50/100 and 200GbE network speeds with ESXi 6.5 onwards, allowing the highest port rate on ESXi today.

## Release Notes Revision History

Date	Revision	Description
May 24, 2021	4.21.71.101	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 7.0 U2. This driver supports the following uplinks to servers.

Version	OS	Uplink Speed
4.21.71.101	ESXi 7.0 U2	10/25/40/50/100/200GbE

## Content of MLNX-NATIVE-ESX Driver Package

ESXi 7.0 U2:

Mellanox-nmlx5\_4.21.71.101-10EM.702.0.0.17630552.zip - Hypervisor bundle for ESXi 7.0 U2 contains the following kernel modules:

- nmlx5\_core
- nmlx5\_rdma

## Supported HCAs Firmware Versions

MLNX-NATIVE-ESX Rev 4.21.71.101 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 <sup>a</sup> (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.

# Changes and New Features

Feature/Change	Description
<b>4.21.71.101</b>	
vSan over RDMA	Added support for vSan over RDMA. For further information, see <a href="https://blogs.nvidia.com/blog/2021/03/24/vsan-over-rdma-hyperconverged-solution/">https://blogs.nvidia.com/blog/2021/03/24/vsan-over-rdma-hyperconverged-solution/</a>
ENS model 0 RSS	Added support for Receive Side Scaling (RSS) for ENS model 0.
ENS model 1 level 1	<b>[Connectx-5 onwards]</b> Added support for ENS FPO Model 1 with Rx path flow lookup offloaded.
ENS model 1 level 2	<b>[Connectx-5 onwards]</b> Added support for ENS FPO Model 1 with Tx path partial action execution offloaded.
ENS model 2A	<b>[Connectx-5 onwards]</b> Added support for ENS FPO Model 2 with SR-IOV as passthrough technology.
Link Speed	Added support for 200GbE link speed.
Adapter Cards	Added support for ConnectX-6 Lx devices.
RDMA 10K QP support	Scaled support for up to 10K connections over RDMA networks.
Hardware DCBX	Added support for Data Center Bridging Capability Exchange (DCBX) protocol with hardware offload.
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. <b>Note:</b> 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.
ens_fallback_model	Added ens_fallback_model to set the default fallback mode when the option to query ENS model from the OS is not supported. Default to Model 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
2204581	<p><b>Description:</b> A mismatch between the uplink and the VF MTU values may result in CQE with error.</p> <p><b>Workaround:</b> Align the uplink and the VF MTU values.</p> <p><b>Keywords:</b> CQE, error, model 2,</p> <p><b>Discovered in Version:</b> 4.21.71.101</p>
2429623	<p><b>Description:</b> Enabling sriov_mc_isolation module parameter may result in vmknic and emulated NICs multicast and IPv6 traffic loss.</p> <p><b>Workaround:</b> Unset or set the module parameter to 0.</p> <p><b>Keywords:</b> Multicast, IPv6, SR-IOV</p> <p><b>Discovered in Version:</b> 4.21.71.101</p>
2372060	<p><b>Description:</b> RDMA is not supported in the Hypervisor with ENS model 2.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> ENS model 2, RDMA</p> <p><b>Discovered in Version:</b> 4.21.71.101</p>
2139469	<p><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.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> MTU, SR-IOV</p> <p><b>Discovered in Version:</b> 4.21.71.101</p>
1340255	<p><b>Description:</b> ECN statistic counters accumulatorsPeriod and ecnMarkedRocePackets display wrong values and cannot be cleared.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> nmlx5 ecn nmlxcli</p> <p><b>Discovered in Version:</b> 4.21.71.101</p>

Internal Ref.	Description
1340275	<p><b>Description:</b> ECN tunable parameter initialAlphaValue for the Reaction Point protocol cannot be modified.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> nmlx5 ecn nmlxcli</p> <p><b>Discovered in Version:</b> 4.21.71.101</p>
2430662	<p><b>Description:</b> Card's speed remains zero after port goes down and reboot is performed.</p> <p><b>Workaround:</b> Turn the port down and then up again.</p> <p><b>Keywords:</b> ConnectX-6 Dx, link speed</p> <p><b>Discovered in Version:</b> 4.21.71.101</p>
1514289	<p><b>Description:</b> RoCE traffic may fail after vMotion when using namespace.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Namespace, RoCE, vMotion</p> <p><b>Discovered in Version:</b> 4.21.71.101</p>
2334405	<p><b>Description:</b> Legacy SR-IOV is not supported with Model 1.</p> <p><b>Workaround:</b> Unset max_vfs or alternatively move to ENS model 0 or Model 2.</p> <p><b>Keywords:</b> SR-IOV, ENS</p> <p><b>Discovered in Version:</b> 4.21.71.101</p>
2449578	<p><b>Description:</b> When in ENS mode, changing the scheduler to HCLK, may cause traffic loss.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> ENS, HCLK scheduler</p> <p><b>Discovered in Version:</b> 4.21.71.101</p>
746100	<p><b>Description:</b> The 'esxcli mellanox uplink link info -u &lt;vmnic_name&gt;' command reports the 'Auto negotiation' capability always as 'true'.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> 'Auto negotiation' capability</p> <p><b>Discovered in Version:</b> 4.21.71.101</p>
1068621	<p><b>Description:</b> SMP MADs (ibnetdiscover, sminfo, iblinkinfo, smpdump, ibqueryerr, ibdiagnet and smpquery) are not supported on the VFs.</p> <p><b>Workaround:</b> N/A</p>

Internal Ref.	Description
	<b>Keywords:</b> SMP MADs
	<b>Discovered in Version:</b> 4.21.71.101
1446060	<p><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:</p> <ul style="list-style-type: none"> <li>• ConnectX-4 / ConnectX-5: up to 127</li> </ul>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> SR-IOV, VFs per port
	<b>Discovered in Version:</b> 4.21.71.101

## 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.