



MLNX_EN for Linux Release Notes

Rev 4.4-1.0.1.0



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 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) 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

Table of Contents	3
Chapter 1 Overview	5
1.1 Supported Platforms and Operating Systems	5
1.1.1 Tested Hypervisors in Paravirtualized and SR-IOV Environments.....	7
1.2 Supported HCAs Firmware Versions.....	8
Chapter 2 Changes and New Features in Rev 4.4-1.0.1.0	9
2.1 Unsupported Functionalities/Features/HCAs	9
Chapter 3 Known Issues	10
Chapter 4 Bug Fixes History	20
Chapter 5 Change Log History	26



Release Update History

Release	Date	Description
Rev 4.4-1.0.1.0	July 22, 2018	Initial release of this version.

1 Overview

These are the release notes of MLNX_EN for Linux Driver, Rev 4.4-1.0.1.0 which operates across all Mellanox network adapter solutions supporting the following uplinks to servers:

Uplink/HCAs	Driver Name	Uplink Speed
ConnectX®-3/ ConnectX®-3 Pro	mlx4	<ul style="list-style-type: none"> Ethernet: 10GigE, 40GigE and 56GigE^a
ConnectX®-4	mlx5	<ul style="list-style-type: none"> Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE^a, and 100GigE
ConnectX®-4 Lx		<ul style="list-style-type: none"> Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, and 50GigE
ConnectX®-5		<ul style="list-style-type: none"> Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, and 100GigE
ConnectX®-5 Ex		<ul style="list-style-type: none"> Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, and 100GigE
Innova™ IPsec EN		<ul style="list-style-type: none"> Ethernet: 10GigE, 40GigE

- a. 56 GbE is a Mellanox propriety link speed and can be achieved while connecting a Mellanox adapter cards to Mellanox SX10XX switch series or connecting a Mellanox adapter card to another Mellanox adapter card.

1.1 Supported Platforms and Operating Systems

The following are the supported OSs in MLNX_EN Rev 4.4-1.0.1.0 :

Table 1 - Supported Platforms and Operating Systems

Operating System	Platform
RHEL6.3/CentOS6.3	x86_64
RHEL6.5/CentOS6.5	x86_64
RHEL6.6/CentOS6.6	x86_64
RHEL6.7/CentOS6.7	x86_64
RHEL6.8/CentOS6.8	x86_64
RHEL6.9/CentOS6.9	x86_64/PPC64 (Power8)
RHEL6.10/CentOS6.10	x86_64/PPC64 (Power8)
RHEL7.1/CentOS7.1	x86_64
RHEL7.2/CentOS7.2	x86_64/PPC64 (Power8)/PPC64LE (Power8)
RHEL7.3/CentOS7.3	x86_64/PPC64 (Power8)/PPC64LE (Power8)
RHEL7.4/CentOS7.4	x86_64/PPC64 (Power8)/PPC64LE (Power8)
RHEL7.4 ALT (Pegas 1.0)	AArch64 (ARM)
RHEL7.5/CentOS7.5	x86_64/PPC64 (Power8)/PPC64LE (Power8)
RHEL7.5 ALT (Pegas 1.1)	PPC64LE (Power8)/AArch64 (ARM)

Table 1 - Supported Platforms and Operating Systems

Operating System	Platform
Debian 8.5	x86_64
Debian 8.7	x86_64
Debian 8.9	x86_64
Debian 9.0	x86_64
Debian 9.1	x86_64
Fedora 20	x86_64
Fedora 27	x86_64
OL 6.7	x86_64
OL 6.8	x86_64
OL 6.9	x86_64
OL 7.2	x86_64
OL 7.3	x86_64
OL 7.4	x86_64
OL 7.5	x86_64
SLES11 SP3	x86_64/PPC64 (Power8)
SLES11 SP4	x86_64/PPC64 (Power8)
SLES12 SP1	x86_64
SLES12 SP2	x86_64/PPC64LE (Power8)
SLES12 SP3	x86_64/PPC64LE (Power8)/AArch64 (ARM)
SLES15	x86_64/PPC64LE (Power8)
Ubuntu 16.04.3	x86_64/PPC64LE (Power8)
Ubuntu 16.04.4	x86_64/PPC64LE (Power8)
Ubuntu 17.10	x86_64/PPC64LE (Power8)
Ubuntu 18.04	x86_64/PPC64LE (Power8)
Kernel 4.14-4.15	x86_64/PPC64 (Power8)/PPC64LE (Power8)
EulerOS2.0 SP3	x86_64/AArch64 (ARM)
WindRiver 6.0	x86_64



32 bit platforms are no longer supported in MLNX_EN

1.1.1 Tested Hypervisors in Paravirtualized and SR-IOV Environments

Table 2 - Supported Non-Linux VMs

HCA	Windows Virtual Machine Type	WinOF version	Protocol
ConnectX-3	Windows 2012 R2 DC	MLNX_VPI 5.50	ETH
ConnectX-3Pro	Windows 2016 DC	MLNX_VPI 5.50	ETH
ConnectX-4	Windows 2012 R2 DC	MLNX_WinOF2 2.0	ETH
ConnectX-4 Lx	Windows 2016 DC	MLNX_WinOF2 2.0	ETH

Table 3 - Tested Hypervisors in Paravirtualized and SR-IOV Environments

Tested Hypervisors	HCAs	Operating System
SR-IOV	All	RHEL6.10 KVM
		RHEL7.1 KVM
		RHEL7.4 KVM
		RHEL7.5 KVM
		SLES12 SP3 KVM
		SLES15 KVM
		EulerOS2.0 SP3 KVM
		Ubuntu16.04.4 KVM
		WindRiver 6.0
Paravirtualized	All	RHEL7.1 KVM
		RHEL7.4 KVM
		RHEL7.5 KVM
		SLES15 KVM
		Ubuntu 16.04.4 KVM

1.2 Supported HCAs Firmware Versions

MLNX_EN Rev 4.4-1.0.1.0 supports the following Mellanox network adapter cards firmware versions:

Table 4 - Supported HCAs Firmware Versions

HCA	Recommended Firmware Rev.	Additional Firmware Rev. Supported
ConnectX®-3	2.42.5000	2.40.7000
ConnectX®-3 Pro	2.42.5000	2.40.7000
ConnectX®-4	12.23.1000	12.22.1002
ConnectX®-4 Lx	14.23.1000	14.22.1002
ConnectX®-5	16.23.1000	16.22.1002
ConnectX®-5 Ex	16.23.1000	16.22.1002
Innova IPsec EN	14.23.1000	14.22.1002

For the official firmware versions, please see:

http://www.mellanox.com/content/pages.php?pg=firmware_download

2 Changes and New Features in Rev 4.4-1.0.1.0

The following are the changes and/or new features that have been added to this version of MLNX_EN

Table 5 - Changes and New Features in Rev 4.4-1.0.1.0

HCA's	Feature/Change	Description
ConnectX-4/ ConnectX-4 Lx/ ConnectX-5	Adaptive Interrupt Moderation	Added support for adaptive Tx, which optimizes the moderation values of the Tx CQs on runtime for maximum throughput with minimum CPU overhead. This mode is enabled by default.
		Updated Adaptive Rx to ignore ACK packets so that queues that only handle ACK packets remain with the default moderation.
	Docker Containers [Beta]	Added support for Docker containers to run over Virtual RoCE and InfiniBand devices using SR-IOV mode.
	Firmware Tracer	Added a new mechanism for the device's FW/HW to log important events into the event tracing system (/sys/kernel/debug/tracing) without requiring any Mellanox-specific tool. Note: This feature is enabled by default.
	CR-Dump	Accelerated the original cr-dump by optimizing the reading process of the device's CR-Space snapshot.
ConnectX-4/ ConnectX-4 Lx	VST Q-in-Q	Added support for C-tag (0x8100) VLAN insertion to tagged packets in VST mode.
ConnectX-4 Lx/ ConnectX-5	OVS Offload using ASAP ²	Added support for Mellanox Accelerated Switching And Packet Processing (ASAP ²) technology, which allows OVS offloading by handling OVS data-plane, while maintaining OVS control-plane unmodified. OVS Offload using ASAP ² technology provides significantly higher OVS performance without the associated CPU load. For further information, refer to ASAP ² Release Notes under www.mellanox.com -> Products -> Software -> ASAP ²
All	Bug Fixes	See Section 4, "Bug Fixes History" , on page 20.

For additional information on the new features, please refer to MLNX_EN User Manual.

2.1 Unsupported Functionalities/Features/HCA's

The following are the unsupported functionalities/features/HCA's in MLNX_EN:

- ConnectX®-2 Adapter Card

3 Known Issues

The following is a list of general limitations and known issues of the various components of this Mellanox EN for Linux release.

For the list of old known issues, please refer to MLNX_EN Archived Known Issues file at: http://www.mellanox.com/pdf/prod_software/MLNX_EN_Archived_Known_Issues.pdf

Table 6 - Known Issues

Internal Reference Number	Issue
1424233	Description: On RHEL v7.3, 7.4 and 7.5 OSs, setting IPv4-IP-forwarding will turn off LRO on existing interfaces. Turning LRO back on manually using ethtool and adding a VLAN interface may cause a warning call trace.
	Workaround: Make sure IPv4-IP-forwarding and LRO are not turned on at the same time.
	Keywords: IPv4 forwarding, LRO
	Discovered in Release: 4.4-1.0.1.0
1425027	Description: Attempting to establish a RoCE connection on the default GID or on IPv6 link-local address might fail when two or more netdevices that belong to HCA ports are slaves under a bonding master. This might also result in the following error message being printed in the kernel log: “__ib_cache_gid_add: unable to add gid fe80:0000:0000:0000:f652:14ff:fe46:7391 error=-28”.
	Workaround: N/A
	Keywords: RoCE, bonding
	Discovered in Release: 4.4-1.0.1.0
1431282	Description: Software reset may result in an order inversion of interface names.
	Workaround: Restart the driver to re-set the order.
	Keywords: Software reset
	Discovered in Release: 4.4-1.0.1.0
1442507	Description: Retpoline support in GCC causes an increase in CPU utilization, which results in IP forwarding’s 15% performance drop.
	Workaround: N/A
	Keywords: Retpoline, GCC, CPU, IP forwarding, Spectre attack
	Discovered in Release: 4.4-1.0.1.0
1400381	Description: On SLES 11 SP3 PPC64 OSs, a memory allocation issue may prevent the interface from loading after reboot, resulting in a call trace in the message log.
	Workaround: Restart the driver.
	Keywords: SLES11 SP3
	Discovered in Release: 4.4-1.0.1.0

Table 6 - Known Issues

Internal Reference Number	Issue
1425129	Description: MLNX_OFED cannot be installed on SLES 15 OSs using Zypper repository.
	Workaround: Install MLNX_OFED using the standard installation script instead of Zypper repository.
	Keywords: Installation, SLES, Zypper
	Discovered in Release: 4.4-1.0.1.0
1241056	Description: When working with ConnectX-4/ConnectX-5 HCAs on PPC systems with Hardware LRO and Adaptive Rx support, bandwidth drops from full wire speed (FWS) to ~60Gb/s.
	Workaround: Make sure to disable Adaptive Rx when enabling Hardware LRO: <pre>ethtool -C <interface> adaptive-rx off ethtool -C <interface> rx-usecs 8 rx-frames 128</pre>
	Keywords: Hardware LRO, Adaptive Rx, PPC
	Discovered in Release: 4.3-1.0.1.0
1090612	Description: NVMeoF protocol does not support LBA format with non-zero metadata size. Therefore, NVMe namespace configured to LBA format with metadata size bigger than 0 will cause Enhanced Error Handling (EEH) in PowerPC systems.
	Workaround: Configure the NVMe namespace to use LBA format with zero sized metadata.
	Keywords: NVMeoF, PowerPC, EEH
	Discovered in Release: 4.3-1.0.1.0
1309621	Description: In switchdev mode default configuration, stateless offloads/steering based on inner headers is not supported.
	Workaround: To enable stateless offloads/steering based on inner headers, disable encap by running: <pre>devlink dev eswitch show pci/0000:83:00.1 encap disable</pre> Or, in case devlink is not supported by the kernel, run: <pre>echo none > /sys/kernel/debug/mlx5/<BDF>/compat/encap</pre> Note: This is a hardware-related limitation.
	Keywords: switchdev, stateless offload, steering
	Discovered in Release: 4.3-1.0.1.0
1275082	Description: When setting a non-default IPv6 link local address or an address that is not based on the device MAC, connection establishments over RoCEv2 might fail.
	Workaround: N/A
	Keywords: IPV6, RoCE, link local address
	Discovered in Release: 4.3-1.0.1.0

Table 6 - Known Issues

Internal Reference Number	Issue
1307336	Description: In RoCE LAG mode, when running <code>ibdev2netdev -v</code> , the port state of the second port of the <code>mlx4_0</code> IB device will read “NA” since this IB device does not have a second port.
	Workaround: N/A
	Keywords: <code>mlx4</code> , RoCE LAG, <code>ibdev2netdev</code> , bonding
	Discovered in Release: 4.3-1.0.1.0
1296355	Description: Number of MSI-X that can be allocated for VFs and PFs in total is limited to 2300 on Power9 platforms.
	Workaround: N/A
	Keywords: MSI-X, VF, PF, PPC, SR-IOV
	Discovered in Release: 4.3-1.0.1.0
1294934	Description: Firmware reset might cause Enhanced Error Handling (EEH) on Power7 platforms.
	Workaround: N/A
	Keywords: EEH, PPC
	Discovered in Release: 4.3-1.0.1.0
1259293	Description: On Fedora 20 operating systems, driver load fails with an error message such as: “[185.262460] kmem_cache_sanity_check (fs_ft-es_0000:00:06.0): Cache name already exists.” This is caused by SLUB allocators grouping multiple slab <code>kmem_cache_create</code> into one slab cache alias to save memory and increase cache hotness. This results in the slab name to be considered stale.
	Workaround: Upgrade the kernel version to <code>kernel-3.19.8-100.fc20.x86_64</code> . Note that after rebooting to the new kernel, you will need to rebuild <code>MLNX_EN</code> against the new kernel version.
	Keywords: Fedora, driver load
	Discovered in Release: 4.3-1.0.1.0
1264359	Description: When running <code>perftest (ib_send_bw, ib_write_bw, etc.)</code> in <code>rdma-cm</code> mode, the <code>resp_cqe_error</code> counter under <code>/sys/class/infiniband/mlx5_0/ports/1/hw_counters/resp_cqe_error</code> might increase. This behavior is expected and it is a result of receive WQEs that were not consumed.
	Workaround: N/A
	Keywords: <code>perftest</code> , RDMA CM, <code>mlx5</code>
	Discovered in Release: 4.3-1.0.1.0

Table 6 - Known Issues

Internal Reference Number	Issue
1264956	<p>Description: Configuring SR-IOV after disabling RoCE LAG using sysfs (/sys/bus/pci/drivers/mlx5_core/<bdf>/roce_lag_enable) might result in RoCE LAG being enabled again in case SR-IOV configuration fails.</p> <p>Workaround: Make sure to disable RoCE LAG once again.</p> <p>Keywords: RoCE LAG, SR-IOV</p> <p>Discovered in Release: 4.3-1.0.1.0</p>
1263043	<p>Description: On RHEL7.4, due to an OS issue introduced in kmod package version 20-15.el7_4.6, parsing the depmod configuration files will fail, resulting in either of the following issues:</p> <ul style="list-style-type: none"> • Driver restart failure prompting an error message, such as: “ERROR: Module mlx5_core belong to kernel which is not a part of MLNX_OFED, skipping...” • nvmet_rdma kernel module dysfunction, despite installing MLNX_OFED using the "--with-nvme" option. An error message, such as: “nvmet_rdma: unknown parameter 'offload_mem_start' ignored” will be seen in dmesg output <p>Workaround: Go to RedHat webpage to upgrade the kmod package version.</p> <p>Keywords: driver restart, kmod, kmp, nvme, nvmet_rdma</p> <p>Discovered in Release: 4.2-1.2.0.0</p>
-	<p>Description: Packet Size (Actual Packet MTU) limitation for IPsec offload on Innova IPsec adapter cards: The current offload implementation does not support IP fragmentation. The original packet size should be such that it does not exceed the interface's MTU size after the ESP transformation (encryption of the original IP packet which increases its length) and the headers (outer IP header) are added:</p> <ul style="list-style-type: none"> • Inner IP packet size <= I/F MTU - ESP additions (20) - outer_IP (20) - fragmentation issue reserved length (56) • Inner IP packet size <= I/F MTU - 96 <p>This mostly affects forwarded traffic into smaller MTU, as well as UDP traffic. TCP does PMTU discovery by default and clamps the MSS accordingly.</p> <p>Workaround: N/A</p> <p>Keywords: Innova IPsec, MTU</p> <p>Discovered in Release: 4.2-1.0.1.0</p>
-	<p>Description: No LLC/SNAP support on Innova IPsec adapter cards.</p> <p>Workaround: N/A</p> <p>Keywords: Innova IPsec, LLC/SNAP</p> <p>Discovered in Release: 4.2-1.0.1.0</p>

Table 6 - Known Issues

Internal Reference Number	Issue
-	Description: No support for FEC on Innova IPsec adapter cards. When using switches, there may be a need to change its configuration.
	Workaround: N/A
	Keywords: Innova IPsec, FEC
	Discovered in Release: 4.2-1.0.1.0
955929	Description: Heavy traffic may cause SYN flooding when using Innova IPsec adapter cards.
	Workaround: N/A
	Keywords: Innova IPsec, SYN flooding
	Discovered in Release: 4.2-1.0.1.0
-	Description: Priority Based Flow Control is not supported on Innova IPsec adapter cards.
	Workaround: N/A
	Keywords: Innova IPsec, Priority Based Flow Control
	Discovered in Release: 4.2-1.0.1.0
-	Description: Pause configuration is not supported when using Innova IPsec adapter cards. Default pause is global pause (enabled).
	Workaround: N/A
	Keywords: Innova IPsec, Global pause
	Discovered in Release: 4.2-1.0.1.0
1045097	Description: Connecting and disconnecting a cable several times may cause a link up failure when using Innova IPsec adapter cards.
	Workaround: N/A
	Keywords: Innova IPsec, Cable, link up
	Discovered in Release: 4.2-1.0.1.0
-	Description: On Innova IPsec adapter cards, supported MTU is between 512 and 2012 bytes. Setting MTU values outside this range might fail or might cause traffic loss.
	Workaround: Set MTU between 512 and 2012 bytes.
	Keywords: Innova IPsec, MTU
	Discovered in Release: 4.2-1.0.1.0

Table 6 - Known Issues

Internal Reference Number	Issue
1118530	Description: On kernel versions 4.10-4.13, when resetting <code>sriov_numvfs</code> to 0 on PowerPC systems, the following dmesg warning will appear: <code>mlx5_core <BDF>: can't update enabled VF BAR0</code>
	Workaround: Reboot the system to reset <code>sriov_numvfs</code> value.
	Keywords: SR-IOV, numvfs
	Discovered in Release: 4.2-1.0.1.0
1125184	Description: In old kernel versions, such as Ubuntu 14.04 and RedHat 7.1, VXLAN interface does not reply to ARP requests for a MAC address that exists in its own ARP table. This issue was fixed in the following newer kernel versions: Ubuntu 16.04 and RedHat 7.3.
	Workaround: N/A
	Keywords: ARP, VXLAN
	Discovered in Release: 4.2-1.0.1.0
1134323	Description: When using kernel versions older than version 4.7 with IOMMU enabled, performance degradations and logical issues (such as soft lockup) might occur upon high load of traffic. This is caused due to the fact that IOMMU IOVA allocations are centralized, requiring many synchronization operations and high locking overhead amongst CPUs.
	Workaround: Use kernel v4.7 or above, or a backported kernel that includes the following patches: <ul style="list-style-type: none"> • 2aac630429d9 iommu/vt-d: change intel-iommu to use IOVA frame numbers • 9257b4a206fc iommu/iova: introduce per-cpu caching to iova allocation • 22e2f9fa63b0 iommu/vt-d: Use per-cpu IOVA caching
	Keywords: IOMMU, soft lockup
	Discovered in Release: 4.2-1.0.1.0
1135738	Description: On 64k page size setups, DMA memory might run out when trying to increase the ring size/number of channels.
	Workaround: Reduce the ring size/number of channels.
	Keywords: DMA, 64K page
	Discovered in Release: 4.2-1.0.1.0
1159650	Description: When configuring VF VST, VLAN-tagged outgoing packets will be dropped in case of ConnectX-4 HCAs. In case of ConnectX-5 HCAs, VLAN-tagged outgoing packets will have another VLAN tag inserted.
	Workaround: N/A
	Keywords: VST
	Discovered in Release: 4.2-1.0.1.0

Table 6 - Known Issues

Internal Reference Number	Issue
1157770	Description: On Passthrough/VM machines with relatively old QEMU and libvirt, CMD timeout might occur upon driver load. After timeout, no other commands will be completed and all driver operations will be stuck.
	Workaround: Upgrade the QEMU and libvirt on the KVM server. Tested with (Ubuntu 16.10) are the following versions: <ul style="list-style-type: none"> • libvirt 2.1.0 • QEMU 2.6.1
	Keywords: QEMU
	Discovered in Release: 4.2-1.0.1.0
1147703	Description: Using dm-multipath for High Availability on top of NVMeoF block devices must be done with “directio” path checker.
	Workaround: N/A
	Keywords: NVMeoF
	Discovered in Release: 4.2-1.0.1.0
1152408	Description: RedHat v7.3 PPCLE and v7.4 PPCLE operating systems do not support KVM qemu out of the box. The following error message will appear when attempting to run virt-install to create new VMs: Cant find qemu-kvm package to install
	Workaround: Acquire the following rpms from the beta version of 7.4ALT to 7.3/ 7.4 PPCLE (in the same order): <ul style="list-style-type: none"> • qemu-img-.el7a.ppc64le.rpm • qemu-kvm-common-.el7a.ppc64le.rpm • qemu-kvm-.el7a.ppc64le.rpm
	Keywords: Virtualization, PPC, Power8, KVM, RedHat, PPC64LE
	Discovered in Release: 4.2-1.0.1.0
1012719	Description: A soft lockup in the CQ polling flow might occur when running very high stress on the GSI QP (RDMA-CM applications). This is a transient situation from which the driver will later recover.
	Workaround: N/A
	Keywords: RDMA-CM, GSI QP, CQ
	Discovered in Release: 4.2-1.0.1.0
1078630	Description: When working in RoCE LAG over kernel v3.10, a kernel crash might occur when unloading the driver as the Network Manager is running.
	Workaround: Stop the Network Manager before unloading the driver and start it back once the driver unload is complete.
	Keywords: RoCE LAG, network manager
	Discovered in Release: 4.2-1.0.1.0

Table 6 - Known Issues

Internal Reference Number	Issue
1149557	Description: When setting VGT+, the maximal number of allowed VLAN IDs presented in the sysfs is 813 (up to the first 813).
	Workaround: N/A
	Keywords: VGT+
	Discovered in Release: 4.2-1.0.1.0
965591	Description: Lustre support is limited to versions 2.9 and 2.10.
	Workaround: N/A
	Keywords: Lustre
	Discovered in Release: 4.1-1.0.2.0
995665/1165919	Description: In kernels below v4.13, connection between NVMeoF host and target cannot be established in a hyper-threaded system with more than 1 socket.
	Workaround: On the host side, connect to NVMeoF subsystem using <code>--nr-io-queues <num_queues></code> flag. Note that <code>num_queues</code> must be lower or equal to <code>num_sockets</code> multiplied with <code>num_cores_per_socket</code> .
	Keywords: NVMeoF
1039346	Description: Enabling multiple namespaces per subsystem while using NVMeoF target offload is not supported.
	Workaround: To enable more than one namespace, create a subsystem for each one.
	Keywords: NVMeoF Target Offload, namespace
1030301	Description: Creating virtual functions on a device that is in LAG mode will destroy the LAG configuration. The bonding device over the Ethernet NICs will continue to work as expected.
	Workaround: N/A
	Keywords: LAG, SR-IOV
1047616	Description: When node GUID of a device is set to zero (0000:0000:0000:0000), RDMA_CM user space application may crash.
	Workaround: Set node GUID to a nonzero value.
	Keywords: RDMA_CM
1051701	Description: New versions of iproute which support new kernel features may misbehave on old kernels that do not support these new features.
	Workaround: N/A
	Keywords: iproute

Table 6 - Known Issues

Internal Reference Number	Issue
1007830	<p>Description: When working on Xenserver hypervisor with SR-IOV enabled on it, make sure the following instructions are applied:</p> <ol style="list-style-type: none"> 1. Right after enabling SR-IOV, unbind all driver instances of the virtual functions from their PCI slots. 2. It is not allowed to unbind PF driver instance while having active VFs. <p>Workaround: N/A</p> <p>Keywords: SR-IOV</p>
1005786	<p>Description: When using ConnectX-5 adapter cards, the following error might be printed to dmesg, indicating temporary lack of DMA pages:</p> <pre> "mlx5_core ... give_pages:289:(pid x): Y pages alloc time exceeded the max permitted duration mlx5_core ... page_notify_fail:263:(pid x): Page alloca- tion failure notification on func_id(z) sent to fw mlx5_core ... pages_work_handler:471:(pid x): give fail - 12" </pre> <p>Example: This might happen when trying to open more than 64 VFs per port.</p> <p>Workaround: N/A</p> <p>Keywords: mlx5_core, DMA</p>
1008066/ 1009004	<p>Description: Performing some operations on the user end during reboot might cause call trace/panic, due to bugs found in the Linux kernel. For example: Running <code>get_vf_stats</code> (via <code>iptool</code>) during reboot.</p> <p>Workaround: N/A</p> <p>Keywords: mlx5_core, reboot</p>
1009488	<p>Description: Mounting MLNX_EN to a path that contains special characters, such as parenthesis or spaces is not supported. For example, when mounting MLNX_EN to <code>"/media/CDROM(vcd)"/</code>, installation will fail and the following error message will be displayed:</p> <pre> # cd /media/CDROM\(vcd\) / # ./install sh: 1: Syntax error: "(" unexpected </pre> <p>Workaround: N/A</p> <p>Keywords: Installation</p>
982144	<p>Description: When offload traffic sniffer is on, the bandwidth could decrease up to 50%.</p> <p>Workaround: N/A</p> <p>Keywords: Offload Traffic Sniffer</p>

Table 6 - Known Issues

Internal Reference Number	Issue
982534	<p>Description: In ConnectX-3, when using a server with page size of 64K, the UAR BAR will become too small. This may cause one of the following issues:</p> <ol style="list-style-type: none"> 1. mlx4_core driver does not load. 2. The mlx4_core driver does load, but calls to <code>ibv_open_device</code> may return ENOMEM errors. <p>Workaround:</p> <ol style="list-style-type: none"> 1. Add the following parameter in the firmware's ini file under [HCA] section: <code>log2_uar_bar_megabytes = 7</code> 2. Re-burn the firmware with the new ini file. <p>Keywords: PPC</p>
981362	<p>Description: On several OSs, setting a number of TC is not supported via the tc tool.</p> <p>Workaround: Set the number of TC via the <code>/sys/class/net/<interface>/qos/tc_num</code> sysfs file.</p> <p>Keywords: Ethernet, TC</p>
979457	<p>Description: When setting IOMMU=ON, a severe performance degradation may occur due to a bug in IOMMU.</p> <p>Workaround: Make sure the following patches are found in your kernel:</p> <ul style="list-style-type: none"> • iommu/vt-d: Fix PASID table allocation • iommu/vt-d: Fix IOMMU lookup for SR-IOV Virtual Functions <p>Note: These patches are already available in Ubuntu 16.04.02 and 17.04 OSs.</p> <p>Keywords: Performance, IOMMU</p>

4 Bug Fixes History

This table lists the bugs fixed in this release.

For the list of old bug fixes, please refer to MLNX_EN Archived Bug Fixes file at: http://www.mellanox.com/pdf/prod_software/MLNX_EN_Archived_Bug_Fixes.pdf

Table 7 - Bug Fixes History

Internal Ref	Issue
1412468	Description: Added support for multi-host connection on mstflint's mstfwreset.
	Keywords: mstfwreset, mstflint, MFT, multi-host
	Discovered in Release: 4.3-1.0.1.0
	Fixed in Release: 4.4-1.0.1.0
1423319	Description: Removed the following prints on server shutdown: mlx5_core 0005:81:00.1: mlx5_enter_error_state:96:(pid1): start mlx5_core 0005:81:00.1: mlx5_enter_error_state:109:(pid1): end
	Keywords: mlx5, fast shutdown
	Discovered in Release: 4.3-1.0.1.0
	Fixed in Release: 4.4-1.0.1.0
1318251	Description: Fixed the issue of when bringing mlx4/mlx5 devices up or down, a call trace in <code>nvme_rdma_remove_one</code> or <code>nvmet_rdma_remove_one</code> may occur.
	Keywords: NVMeoF, mlx4, mlx5, call trace
	Discovered in Release: 4.3-1.0.1.0
	Fixed in Release: 4.4-1.0.1.0
1181815	Description: Fixed an issue where 4K UD packets were dropped when working with 4K MTU on mlx4 devices.
	Keywords: mlx4, 4K MTU, UD
	Discovered in Release: 4.2-1.2.0.0
	Fixed in Release: 4.3-1.0.1.0
1247458	Description: Added support for VLAN Tag (VST) creation on RedHat v7.4 with new <code>iproute2</code> packages (<code>iptool</code>).
	Keywords: SR-IOV, VST, RedHat
	Discovered in Release: 4.2-1.2.0.0
	Fixed in Release: 4.3-1.0.1.0
1229554	Description: Enabled RDMA CM to honor incoming requests coming from ports of different devices.
	Keywords: RDMA CM
	Discovered in Release: 4.2-1.0.0.0
	Fixed in Release: 4.3-1.0.1.0

Table 7 - Bug Fixes History

Internal Ref	Issue
1262257	Description: Fixed an issue where sending Work Requests (WRs) with multiple entries where the first entry is less than 18 bytes used to fail.
	Keywords: ConnectX-5; libibverbs; Raw QP
	Discovered in Release: 4.2-1.2.0.0
	Fixed in Release: 4.3-1.0.1.0
1249358/1261023	Description: Fixed the issue of when the interface was down, ethtool counters ceased to increase. As a result, RoCE traffic counters were not always counted.
	Keywords: Ethtool counters, mlx5
	Discovered in Release: 4.2-1.2.0.0
	Fixed in Release: 4.3-1.0.1.0
1244509	Description: Fixed compilation errors of MLNX_EN over kernel when CONFIG_PTP_1588_CLOCK parameter was not set.
	Keywords: PTP, mlx5e
	Discovered in Release: 4.2-1.2.0.0
	Fixed in Release: 4.3-1.0.1.0
1266802	Description: Fixed an issue where the system used to hang when trying to allocate multiple device memory buffers from different processes simultaneously.
	Keywords: Device memory programming
	Discovered in Release: 4.2-1.0.0.0
	Fixed in Release: 4.3-1.0.1.0
1078887	Description: Fixed an issue where post_list and CQ_mod features in perftest did not function when running the --run_indefinitely flag.
	Keywords: perftest, --run_indefinitely
	Discovered in Release: 4.2-1.0.1.0
	Fixed in Release: 4.2-1.2.0.0
1186260	Description: Fixed the issue where CNP counters exposed under /sys/class/infiniband/mlx5_bond_0/ports/1/hw_counters/ did not aggregate both physical functions when working in RoCE LAG mode.
	Keywords: RoCE, LAG, ECN, Congestion Counters
	Discovered in Release: 4.2-1.0.1.0
	Fixed in Release: 4.2-1.2.0.0
1192374	Description: Fixed wrong calculation of max_device_ctx capability in ConnectX-4, ConnectX-4 Lx, and ConnectX-5 HCAs.
	Keywords: ibv_exp_query_device, max_device_ctx mlx5
	Discovered in Release: 4.2-1.0.1.0
	Fixed in Release: 4.2-1.2.0.0

Table 7 - Bug Fixes History

Internal Ref	Issue
1084791	Description: Fixed the issue where occasionally, after reboot, rpm commands used to fail and create a core file, with messages such as “Bus error (core dumped)”, causing the openibd service to fail to start.
	Keywords: rpm, openibd
	Discovered in Release: 3.4-2.0.0.0
	Fixed in Release: 4.2-1.0.1.0
960642/960653	Description: Added support for min_tx_rate and max_tx_rate limit per virtual function ConnectX-5 and ConnectX-5 Ex adapter cards.
	Keywords: SR-IOV, mlx5
	Discovered in Release: 4.0-1.0.1.0
	Fixed in Release: 4.2-1.0.1.0
866072/869183	Description: Fixed the issue where RoCE v2 multicast traffic using RDMA-CM with IPv4 address was not received.
	Keywords: RoCE
	Discovered in Release: 3.4-1.0.0.0
	Fixed in Release: 4.2-1.0.1.0
1163835	Description: Fixed an issue where ethtool -P output was 00:00:00:00:00:00 when using old kernels.
	Keywords: ethtool, Permanent MAC address, mlx4, mlx5
	Discovered in Release: 4.0-2.0.0.1
	Fixed in Release: 4.2-1.0.1.0
1067158	Description: Replaced a few “GPL only” legacy libibverbs functions with upstream implementation that conforms with libibverbs GPL/BSD dual license model.
	Keywords: libibverbs, license
	Discovered in Release: 4.1-1.0.2.0
	Fixed in Release: 4.2-1.0.1.0
1119377	Description: Fixed an issue where ACCESS_REG command failure used to appear upon RoCE Multihost driver restart in dmesg. Such an error message looked as follows: <pre>mlx5_core 0000:01:00.0: mlx5_cmd_check:705:(pid 20037): ACCESS_REG(0x805) op_mod(0x0) failed, status bad parameter(0x3), syndrome (0x15c356)</pre>
	Keywords: RoCE, multihost, mlx5
	Discovered in Release: 4.1-1.0.2.0
	Fixed in Release: 4.2-1.0.1.0

Table 7 - Bug Fixes History

Internal Ref	Issue
1122937	Description: Fixed an issue where concurrent client requests got corrupted when working in persistent server mode due to a race condition on the server side.
	Keywords: librdmacm, rping
	Discovered in Release: 4.1-1.0.2.0
	Fixed in Release: 4.2-1.0.1.0
1102158	Description: Fixed an issue where client side did not exit gracefully in RTT mode when the server side was not reachable.
	Keywords: librdmacm, rping
	Discovered in Release: 4.1-1.0.2.0
	Fixed in Release: 4.2-1.0.1.0
1038933	Description: Fixed a backport issue where IPv6 procedures were called while they were not supported in the underlying kernel.
	Keywords: iw_cm
	Discovered in Release: 4.0-2.0.0.1
	Fixed in Release: 4.1-1.0.2.0
1064722	Description: Added log debug prints when changing HW configuration via DCB. To enable log debug prints, run: <code>ethtool -s <devname> msglvl hw on/off</code>
	Keywords: DCB, msglvl
	Discovered in Release: 4.0-2.0.0.1
	Fixed in Release: 4.1-1.0.2.0
1022251	Description: Fixed SKB memory leak issue that was introduced in kernel 4.11, and added warning messages to the Soft RoCE driver for easy detection of future SKB leaks.
	Keywords: Soft RoCE
	Discovered in Release: 4.0-2.0.0.1
	Fixed in Release: 4.1-1.0.2.0
1044546	Description: Fixed the issue where a kernel crash used to occur when Rxe device was coupled with a virtual (dummy) device.
	Keywords: Soft RoCE
	Discovered in Release: 4.0-2.0.0.1
	Fixed in Release: 4.1-1.0.2.0
1047617	Description: Fixed the issue where a race condition in the RoCE GID cache used to cause for the loss of IP-based GIDs.
	Keywords: RoCE, GID
	Discovered in Release: 4.0-2.0.0.1
	Fixed in Release: 4.1-1.0.2.0

Table 7 - Bug Fixes History

Internal Ref	Issue
1006768	Description: Fixed the issue where an rdma_cm connection between a client and a server that were on the same host was not possible when working over VLAN interfaces.
	Keywords: RDMACM
	Discovered in Release: 4.0-2.0.0.1
	Fixed in Release: 4.1-1.0.2.0
801807	Description: Fixed an issue where RDMACM connection used to fail upon high connection rate accompanied with the error message: RDMA_CM_EVENT_UNREACHABLE.
	Keywords: RDMACM
	Discovered in Release: 3.0-2.0.1
	Fixed in Release: 4.1-1.0.2.0
869768	Description: Fixed the issue where SR-IOV was not supported in systems with a page size greater than 16KB.
	Keywords: SR-IOV, mlx5, PPC
	Discovered in Release: 4.0-2.0.0.1
	Fixed in Release: 4.1-1.0.2.0
919545	Description: Fixed the issue of when the Kernel becomes out of memory upon driver start, it could crash on SLES 12 SP2.
	Keywords: mlx_5 Eth Driver
	Discovered in Release: 3.4-2.0.0.0
	Fixed in Release: 4.0-2.0.0.1
864063	Description: Fixed the issue of when Spoof-check may have been turned on for MAC address 00:00:00:00:00:00.
	Keywords: mlx4
	Discovered in Release: 3.4-1.0.0.0
	Fixed in Release: 4.0-2.0.0.1
869209	Description: Fixed an issue that caused TCP packets to be received in an out of order manner when Large Receive Offload (LRO) is on.
	Keywords: mlx5_en
	Discovered in Release: 3.3-1.0.0.0
	Fixed in Release: 4.0-2.0.0.1

Table 7 - Bug Fixes History

Internal Ref	Issue
890285	Description: Fixed the issue where memory allocation for CQ buffers used to fail when increasing the RX ring size.
	Keywords: mlx5_core
	Discovered in Release: 3.4-1.0.0.0
	Fixed in Release: 4.0-1.0.1.0
867094	Description: Fixed the issue where MLNX_EN used to fail to load on 4K page Arm architecture.
	Keywords: Arm
	Discovered in Release: 3.4-1.0.0.0
	Fixed in Release: 4.0-1.0.1.0

5 Change Log History

Table 8 - Change Log History

Category	Description
4.3-1.0.1.0	
HCAs: ConnectX-5	
Erasure Coding Offload verbs	Added support for erasure coding offload software verbs (encode/decode/update API) supporting a number of redundancy blocks (m) greater than 4.
HCAs: ConnectX-4/ConnectX-4 Lx/ConnectX-5	
Virtual MAC	Removed support for Virtual MAC feature.
RoCE LAG	Added out of box RoCE LAG support for RHEL 7.2 and RHEL 6.9.
Dropped Counters	Added a new counter <code>rx_steer_missed_packets</code> which provides the number of packets that were received by the NIC, yet were discarded/dropped since they did not match any flow in the NIC steering flow table.
	Added the ability for SR-IOV counter <code>rx_dropped</code> to count the number of packets that were dropped while vport was down.
HCAs: mlx5 Driver	
Reset Flow	Added support for triggering software reset for firmware/driver recovery. When fatal errors occur, firmware can be reset and driver reloaded.
HCAs: ConnectX-4 Lx/ConnectX-5	
Striding RQ with HW Time-Stamping	Added the option to retrieve the HW timestamp when polling for completions from a completion queue that is attached to a multi-packet RQ (Striding RQ).
4.2-1.0.1.0	
HCAs: mlx5 Driver	
Physical Address Memory Allocation	Added support to register a specific physical address range.
HCAs: Innova IPsec EN	
Innova IPsec Adapter Cards	Added support for Mellanox Innova IPsec EN adapter card, that provides security acceleration for IPsec-enabled networks.
HCAs: ConnectX-4/ConnectX-4 Lx/ConnectX-5	
Precision Time Protocol (PTP)	Added support for PTP feature over PKEY interfaces. This feature allows for accurate synchronization between the distributed entities over the network. The synchronization is based on symmetric Round Trip Time (RTT) between the master and slave devices, and is enabled by default.

Table 8 - Change Log History

Category	Description
Virtual MAC	Added support for Virtual MAC feature, which allows users to add up to 4 virtual MACs (VMACs) per VF. All traffic that is destined to the VMAC will be forwarded to the relevant VF instead of PF. All traffic going out from the VF with source MAC equal to VMAC will go to the wire also when Spoof Check is enabled. For further information, please refer to “Virtual MAC” section in MLNX_EN User Manual.
Receive Buffer	Added the option to change receive buffer size and cable length. Changing cable length will adjust the receive buffer's xon and xoff thresholds. For further information, please refer to “Receive Buffer” section in MLNX_EN User Manual.
GRE Tunnel Offloads	Added support for the following GRE tunnel offloads: <ul style="list-style-type: none"> • TSO over GRE tunnels • Checksum offloads over GRE tunnels • RSS spread for GRE packets
NVMeoF	Added support for the host side (RDMA initiator) in RedHat 7.2 and above.
Dropless Receive Queue (RQ)	Added support for the driver to notify the FW when SW receive queues are overloaded.
PFC Storm Prevention	Added support for configuring PFC stall prevention in cases where the device unexpectedly becomes unresponsive for a long period of time. PFC stall prevention disables flow control mechanisms when the device is stalled for a period longer than the default pre-configured timeout. Users now have the ability to change the default timeout by moving to auto mode. For further information, please refer to “PFC Stall Prevention” section in MLNX_EN User Manual.
HCAs: ConnectX-5	
Q-in-Q	Added support for Q-in-Q VST feature in ConnectX-5 adapter cards family.
Virtual Guest Tagging (VGT+)	Added support for VGT+ in ConnectX-4/ConnectX-5 HCAs. This feature is an advanced mode of Virtual Guest Tagging (VGT), in which a VF is allowed to tag its own packets as in VGT, but is still subject to an administrative VLAN trunk policy. The policy determines which VLAN IDs are allowed to be transmitted or received. The policy does not determine the user priority, which is left unchanged. For further information, please refer to “Virtual Guest Tagging (VGT+)” section in MLNX_EN User Manual.
Tag Matching Offload	Added support for hardware Tag Matching offload with Dynamically Connected Transport (DCT).
HCAs: ConnectX-3/ConnectX-3 Pro	
HCAs: All	
CR-DUMP	Added support for the driver to take an automatic snapshot of the device's CR-Space in cases of critical failures. For further information, please refer to “CRDUMP” section in MLNX_EN User Manual.

Table 8 - Change Log History

Category	Description
4.1-1.0.2.0	
HCAs: mlx5 Driver	
RoCE Diagnostics and ECN Counters	Added support for additional RoCE diagnostics and ECN congestion counters under <code>/sys/class/infiniband/mlx5_0/ports/1/hw_counters/</code> directory. For further information, refer to the Understanding mlx5 Linux Counters and Status Parameters Community post.
rx-fcs Offload (eth-tool)	Added support for rx-fcs ethtool offload configuration. Normally, the FCS of the packet will be truncated by the ASIC hardware before sending it to the application socket buffer (skb). Ethtool allows to set the rx-fcs not to be truncated, but to pass it to the application for analysis. For more information and usage, refer to Understanding ethtool rx-fcs for mlx5 Drivers Community post.
DSCP Trust Mode	Added the option to enable PFC based on the DSCP value. Using this solution, VLAN headers will no longer be mandatory for use. For further information, refer to the HowTo Configure Trust Mode on Mellanox Adapters Community post.
RoCE ECN Parameters	ECN parameters have been moved to the following directory: <code>/sys/kernel/debug/mlx5/<PCI BUS>/cc_params/</code> For more information, refer to the HowTo Configure DCQCN (RoCE CC) for ConnectX-4 (Linux) Community post.
Flow Steering Dump Tool	Added support for <code>mlx_fs_dump</code> , which is a python tool that prints the steering rules in a readable manner.
Secure Firmware Updates	Firmware binaries embedded in <code>MLNX_EN</code> package now support Secure Firmware Updates. This feature provides devices with the ability to verify digital signatures of new firmware binaries, in order to ensure that only officially approved versions are installed on the devices. For further information on this feature, refer to Mellanox Firmware Tools (MFT) User Manual.
PeerDirect	Added the ability to open a device and create a context while giving PCI peer attributes such as name and ID. For further details, refer to the PeerDirect Programming Community post.
Probed VFs	Added the ability to disable probed VFs on the hypervisor. For further information, see HowTo Configure and Probe VFs on mlx5 Drivers Community post.
Local Loopback	Improved performance by rendering Local loopback (unicast and multicast) disabled by mlx5 driver by default while local loopback is not in use. The mlx5 driver keeps track of the number of transport domains that are opened by user-space applications. If there is more than one user-space transport domain open, local loopback will automatically be enabled.
1PPS Time Synchronization (at alpha level)	Added support for One Pulse Per Second (1PPS), which is a time synchronization feature that allows the adapter to send or receive 1 pulse per second on a dedicated pin on the adapter card. For further information on this feature, refer to the HowTo Test 1PPS on Mellanox Adapters Community post.

Table 8 - Change Log History

Category	Description
Fast Driver Unload	Added support for fast driver teardown in shutdown and kexec flows.
HCAs: ConnectX-5/ConnectX-5 Ex	
NVMeoF Target Offload	Added support for NVMe over fabrics (NVMeoF) offload, an implementation of the new NVMeoF standard target (server) side in hardware. For further information on NVMeoF Target Offload, refer to HowTo Configure NVMeoF Target Offload .
HCAs: All	
RDMA CM	Changed the default RoCE mode on which RDMA CM runs to RoCEv2 instead of RoCEv1. RDMA_CM session requires both the client and server sides to support the same RoCE mode. Otherwise, the client will fail to connect to the server. For further information, refer to RDMA CM and RoCE Version Defaults Community post.
Lustre	Added support for Lustre file system open-source project.
4.0-2.0.0.1	
PCIe Error Counting	[ConnectX-4/ConnectX-4 Lx] Added the ability to expose physical layer statistical counters to ethtool.
Standard ethtool	[ConnectX-4/ConnectX-4 Lx] Added support for flow steering and rx-all mode.
SR-IOV Bandwidth Share for Ethernet/RoCE (beta)	[ConnectX-4/ConnectX-4 Lx] Added the ability to guarantee the minimum rate of a certain VF in SR-IOV mode.
Adapter Cards	Added support for ConnectX-5 and ConnectX-5 Ex HCAs.
NFS over RDMA (NFSoRDMA)	Removed support for NFSoRDMA drivers. These drivers are no longer provided along with the MLNX_EN package.