
Connect. Accelerate. Outperform.™

# MLNX_EN for Linux Release Notes

Rev 3.0-1.0.1

www.mellanox.com


Mellanox Technologies
350 Oakmead Parkway Suite 100
Sunnyvale, CA 94085
U.S.A.
www.mellanox.com
Tel: (408) 970-3400
Fax: (408) 970-3403

Mellanox Technologies, Ltd.
Hakidma 26
Ofer Industrial Park
Yokneam 2069200
Israel
www.mellanox.com
Tel: +972 (0)74 723 7200
Fax: +972 (0)4 959 3245

# Table of Contents

# List of Tables

# Release Update History

***Table 1 - Release Update History***

| Release | Date | Description |
|---|---|---|
| Rev 3.0-1.0.1 | June 15, 2015 | Initial version |

# 1 Introduction

These are the release notes for Mellanox Technologies' MLNX_EN for Linux Rev 3.0-1.0.1driver kit for Mellanox adapter cards supporting the following uplinks to servers:

- ConnectX®-4:
  - Ethernet (Beta): 10GigE, 25GigE, 40GigE, 50GigE and 100GigE
- ConnectX®-3/ConnectX®-3 Pro:
  - Ethernet: 10GigE, 40GigE and 56GigE[1]
- PCI Express 2.0: 2.5 or 5.0 GT/s
- PCI Express 3.0: 8 GT/s

## 1.1 Main Features in This Release

MLNX_EN for Linux Rev 3.0-1.0.1 provides the following new features:

***Table 2 - MLNX_EN for Linux Rev 3.0-1.0.1 Main Features***

| Feature | Description |
|---|---|
| HCAs | Added support for ConnectX®-4 Single/Dual-Port Adapter supporting up to 100Gb/s. |
| Ignore Frame Check Sequence (FCS) Errors | Upon receiving packets, the packets go through a checksum validation process for the FCS field. If the validation fails, the received packets are dropped. Using this feature, enables you to choose whether or not to drop the frames in case the FCS is wrong and use the FCS field for other info. |
| ethtool | Updated ethtool to incorporate ConnectX®-4 adapter card functionalities. |

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

## 1.2 Unsupported Functionality/Features

The following are the unsupported functionalities/features in MLNX_OFED Rev 3.0-1.0.1:

- pm_qos API - Power Management
- Adaptive Interrupt Moderation Algorithm
- Virtual Guest Tagging (VGT+)

1. 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.

## 1.3 Supported Platforms and Operating Systems

The following are the supported OSs in MLNX_EN for Linux Rev 3.0-1.0.1:

***Table 3 - Supported Platforms and Operating Systems***

| Operating System | Platform |
|---|---|
| RHEL/CentOS 6.4 | x86_64/PPC64 (Power 7) |
| RHEL/CentOS 6.5 | x86_64/PPC64 (Power 7) |
| RHEL/CentOS 6.6 | x86_64/PPC64 (Power 7) |
| RHEL/CentOS 7.0 | x86_64/PPC64 (Power 7) |
| RHEL/CentOS 7.1 | x86_64/PPC64 (Power 7)/PPC64le (Power 8)/ ARM64 (ARM is at beta level) |
| SLES11 SP2 | x86_64/PPC64 (Power 7) |
| SLES11 SP3 | x86_64/PPC64 (Power 7) |
| SLES12 | x86_64/PPC64le (Power 8) |
| OEL 6.3 | x86_64 |
| OEL 6.4 | x86_64 |
| OEL 6.5 | x86_64 |
| OEL 6.6 | x86_64 |
| Citrix XenServer Host 6.5 | i686 |
| Fedora 19 | x86_64/PPC64 (Power 7) |
| Fedora 21 | x86_64 |
| Ubuntu 12.04 | x86_64 |
| Ubuntu 14.04 | x86_64/PPC64le (Power 8)/ARM64 (ARM is at beta level) |
| Ubuntu 14.10 | x86_64/PPC64le (Power 8) |
| Debian 7.6 | x86_64 |
| Debian 8.0 | x86_64 |
| kernel 3.16[a] - 4.0 | |

a. This kernel is supported only when using the Operating Systems stated in the table above.

## 1.4 Supported Hypervisors

The following are the supported hypervisors in MLNX_EN Rev 3.0-1.0.1:

- KVM: RedHat 6.5, 6.6, 7.0, 7.1 Ubuntu 14.04, Sles11SP3
- Xen4.2
- XenServer6.5

## 1.5 Supported HCAs Firmware Versions

MLNX_OFED Rev 3.0-1.0.1 supports the following Mellanox network adapter cards firmware versions:

***Table 4 - Supported HCAs Firmware Versions***

| HCA | Recommended Firmware Rev. | Additional Firmware Rev. Supported |
|---|---|---|
| ConnectX®-4 | 12.0100.6440 | - |

***Table 4 - Supported HCAs Firmware Versions***

| HCA | Recommended Firmware Rev. | Additional Firmware Rev. Supported |
|---|---|---|
| ConnectX®-3 Pro | Rev 2.34.5000 | Rev 2.33.5100 |
| ConnectX®-3 | Rev 2.34.5000 | Rev 2.33.5100 |
| ConnectX®-2 | Rev 2.9.1000 | Rev 2.9.1000 |

For official firmware versions please see:
http://www.mellanox.com/content/pages.php?pg=firmware_download

# 2 Changes in Rev 3.0-1.0.1 From Rev 2.4-1.0.0.1

***Table 5 - Changes in v3.0-1.0.1***

| Category | Description |
|---|---|
| HCAs | Added support for ConnectX®-4 Single/Dual-Port Adapter supporting 100Gb/s. |
| ethtool | Updated ethtool to incorporate ConnectX®-4 adapter card functionalities. |

# 3 Known Issues

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

## 3.1 Ethernet Known Issues

***Table 6 - Ethernet Known Issues***

| Index | Description | Workaround |
|---|---|---|
| 1. | Ethernet PV VLAN Guest transparent Tagging (VGT) is only supported in openvswitch and not in standard Linux vBridges and libvirt<br>For more information please see : http://libvirt.org/formatnetwork.html (Setting VLAN tag section) | - |
| 2. | Changing the ring size on 32-bit system may result in failure due to lack of memory. Therefore, mlx4_en will not be able to vmap enough memory and the below message will be printed in dmesg:<br>`vmap allocation for size 528384 failed: use vmalloc=<size> to increase size`<br>`In this case user can enlarge the vmalloc memory by adding vmalloc=<size> to grub.conf`<br>`Default vmalloc setting is 128M. It is recommended to add each time 64M of memory until desired ring size can be allocated.`<br>Please note, that in case vmalloc size is too big, the OS will fail to boot, so please use caution when adding additional memory.<br>For more info refer to: http://www.mythtv.org/wiki/Common_Problem:_vmalloc_too_small | - |
| 3. | Kernel panic might occur during fio splice in kernels before 2.6.34-rc4. | Use kernel v2.6.34-rc4 which provides the following solution:<br>`baff42a net: Fix oops from tcp_collapse() when using splice()` |
| 4. | When creating more than 125 VLANs and SR-IOV mode is enabled, a kernel warning message will be printed indicating that the native VLAN is created but will not work with RoCE traffic.<br><br>kernel warning: `mlx4_core 0000:07:00.0: vhcr command ALLOC_RES (0xf00) slave:0 in_param 0x7e in_mod=0x107, op_mod=0x1 failed with error:0, status -28` | |
| 5. | In PPC systems when QoS is enabled a harmless Kernel DMA mapping error messages might appear in kernel log (iommu related issue). | - |

***Table 6 - Ethernet Known Issues (Continued)***

| Index | Description | Workaround |
|---|---|---|
| 6. | The default priority to TC mapping assigns all priorities to TC0. This configuration achieves fairness in transmission between priorities but may cause undesirable PFC behavior where pause request for priority "n" affects all other priorities. | Run: `mlnx_qos -i <dev> -p 0,1,2,3,4,5,6,7 -s ets,ets,ets,ets,ets,ets,ets,ets -t 12,13,12,13,12,13,12,13` This needs to be applied every time after loading the `mlx4_en` driver. |
| 7. | Transmit timeout might occur on RH6.3 as a result of lost interrupt (OS issue). In this case, the following message will be shown in dmesg: `do_IRQ: 0.203 No irq handler for vector (irq -1)` | - |
| 8. | In ConnectX®-2, RoCE UD QP does not include VLAN tags in the Ethernet header | |
| 9. | VXLAN may not be functional when configured over Linux bridge in RH7.0 or Ubuntu14.04. The issue is within the bridge modules in those kernels. In Vanilla kernels above 3.16 issues is fixed. | - |
| 10. | In RH6.4, ping may not work over VLANs that are configured over Linux bridge when the bridge has a mlx4_en interface attached to it. | - |
| 11. | The interfaces LRO needs to be set to "OFF" manually when there is a bond configured on Mellanox interfaces with a Bridge over that bond. | Run: `ethtool -K ethX lro off` |
| 12. | On SLES12, the bonding interface over Mellanox Ethernet slave interfaces does not get IP address after reboot. | 1. Set "STARTMODE=hotplug" in the bonding slave's ifcfg files. More details can be found in the SUSE documentations page: https://www.suse.com/documentation/sles-12/book_sle_admin/?page=/documentation/sles-12/book_sle_admin/data/sec_bond.html<br>2. Enable the "nanny" service to support hot-plugging: Open the `"/etc/wicked/common.xml"` file. Change: `"<use-nanny>false</use-nanny>"` to `"<use-nanny>true</use-nanny>"`<br>3. Run: `# systemctl restart wickedd.service wicked` |
| 13. | ethtool -x command does not function in SLES OS. | - |
| 14. | Ethertype proto 0x806 not supported by ethtool. | - |

## 3.2 mlx5 Driver (ConnectX®-4) Known Issues

***Table 7 - mlx5 Driver Known Issues***

| Index | Description | Workaround |
|---|---|---|
| 1. | EEH events that arrive while the mlx5 driver is loading may cause the driver to hang. | - |
| 2. | The mlx5 driver can handle up to 5 EEH events per hour. | If more events are received, cold reboot the machine. |
| 3. | In PPC systems, when working with ConnextX®-4 adapter card configured as Ethernet, driver load fails with BAD INPUT LENGTH. dmesg: `command failed, status bad input length(0x50), syndrome 0x9074aa` | - |
| 4. | Changing the link speed is not supported in Ethernet driver when connected to a ConnectX-4 card. | - |
| 5. | Changing the RX queues number is not supported in Ethernet driver when connected to a ConnectX-4 card. | - |
| 6. | Error counters such as: CRC error counters, RX out range length error counter, are missing in the ConnectX-4 Ethernet driver. | - |

## 3.3 Resiliency Known Issues

***Table 8 - Resiliency Known Issues***

| Index | Description | Workaround |
|---|---|---|
| 1. | Reset Flow can run on XenServer 6 only after the active user space applications running verbs are terminated. | - |
| 2. | SR-IOV non persistent configuration (such as VGT, VST, Host assigned GUIDs, and QP0-enabled VFs) may be lost upon Reset Flow. | Reset Admin configuration post Reset Flow |
| 3. | Upon Reset Flow or after running restart driver, Ethernet VLANs are lost. | Reset the VLANs using the `ifup` command. |
| 4. | Restarting the driver or running `connectx_port_config` when Reset Flow is running might result in a kernel panic | - |
| 5. | Networking configuration (e.g. VLANs, IPv6) should be statically defined in order to have them set after Reset Flow as of after restart driver. | - |
| 6. | The qemu version (2.0) provided inbox with Ubuntu 14.04 does not work properly when more than 2 VMs are run over an Ubuntu 14.04 Hypervisor. | - |
| 7. | Attempting to attach a PF to a VM when SR-IOV is already enabled on that PF may result in a kernel panic. | - |
| 8. | On various combinations of Hypervisor/OSes and Guest/OSes, an issue might occur when attaching/detaching VFs to a guest while that guest is up and running. | Attach/detach VFs to/from a VM only while that VM is down. |

## 3.4 Memory Consumption Known Issues

***Table 9 - Memory Consumption Known Issues***

| Index | Description | Workaround |
|---|---|---|
| 1. | "Out of memory" issues may rise during drivers load depending on the values of the driver module parameters set (e.g. log_num_cq) | - |

## 3.5 Ethernet Performance Counters Known Issues

***Table 10 - Ethernet Performance Counters Known Issues***

| Index | Description | Workaround |
|---|---|---|
| 1. | In a system with more than 61 VFs, the 62nd VF and onwards is assigned with the SINKQP counter, and as a result will have no statistics, and loopback prevention functionality for SINK counter. | - |
| 2. | Since each VF tries to allocate 2 more QP counter for its RoCE traffic statistics, in a system with less than 61 VFs, if there is free resources it receives new counter otherwise receives the default counter which is shared with Ethernet. In this case RoCE statistics is not available. | - |
| 3. | In ConnectX®-3, when we enable function-based loopback prevention for Ethernet port by default (i.e., based on the QP counter index), the dropped self-loopback packets increase the IfRxErrorFrames/Octets counters. | - |

## 3.6 Hardware Known Issues

***Table 11 - Hardware Known Issues***

| Index | Description | Workaround |
|---|---|---|
| 1. | In ConnectX®-2, if the driver load succeeds, the informative message below is presented conveying the below limitations: `If port type is ETH then the maximum priority for VLAN tagged is 3 mlx4_core 0000:0d:00.0: command SET_PORT (0xc) failed: in_param=0x120064000, in_mod=0x2, op_mod=0x0, fw status = 0x40` | - |

## 3.7 Operating Systems Known Issues

***Table 12 - Operating Systems Issues***

| Index | Description | Workaround |
|---|---|---|
| 1. | RHEL 5.X and SLES 10 SPX are currently not supported | - |

## 3.8 SR-IOV Known Issues

***Table 13 - SR-IOV Issues***

| Index | Description | Workaround |
|---|---|---|
| 1. | Enabling SR-IOV requires appending the `intel_iommu=on` option to the relevant OS in file `/boot/grub/grub.conf/`. Without that SR-IOV cannot be loaded. | - |
| 2. | When using legacy VMs with MLNX_EN 2.x hypervisor, you may need to set the `'enable_64b_cqe_eqe'` parameter to zero on the hypervisor.<br>It should be set in the same way that other module parameters are set for mlx4_core at module load time.<br>For example, add `"options mlx4_core enable_64b_cqe_eqe=0"` as a line in the file `/etc/modprobe.d/mlx4_core.conf`. | - |
| 3. | Attaching or detaching a Virtual Function on SLES11 SP3 to a guest Virtual Machine while the `mlx4_core` driver is loaded in the Virtual Machine may cause a kernel panic in the hypervisor. | Unload the `mlx4_core` module in the hypervisor before attaching or detaching a function to or from the guest. |
| 4. | Enabling SR-IOV requires appending the `"intel_iommu=on"` option to the relevant OS in file `/boot/grub/grub.conf/`. Without that SR-IOV cannot be loaded. | - |
| 5. | The known PCI BDFs for all VFs in kernel command line should be specified by adding `xen-pciback.hide`<br>For further information, please refer to http://wiki.xen.org/wiki/Xen_PCI_Passthrough | - |
| 6. | The qemu version (2.0) provided inbox with Ubuntu 14.04 does not work properly when more than 2 VMs are run over an Ubuntu 14.04 Hypervisor. | - |
| 7. | Attempting to attach a PF to a VM when SR-IOV is already enabled on that PF may result in a kernel panic. | - |
| 8. | On various combinations of Hypervisor/OSes and Guest/OSes, an issue might occur when attaching/detaching VFs to a guest while that guest is up and running. | Attach/detach VFs to/from a VM only while that VM is down. |
| 9. | On ConnectX®-3 HCAs with firmware version 2.32.5000 and later, SR-IOV VPI mode works only with Port 1 = ETH and Port 2 = IB. | - |
| 10. | Occasionally, the `lspci \| grep Mellanox` command shows incorrect or partial information due to the current `pci.ids` file on the machine. | 1. Locate the file:<br>`$locate pci.ids`<br>2. Manually update the file according to the latest version available online at:<br>https://pci-ids.ucw.cz/v2.2/pci.ids<br>This file can also be downloaded. |
| 11 | SR-IOV is not supported in XenServer 6.5. | - |
| 12 | SR-IOV is not supported in AMD architecture. | - |

## 3.9 Quality of Service Known Issues

***Table 14 - Quality of Service Issues***

| Index | Description | Workaround |
|---|---|---|
| 1. | QoS is not supported in XenServer, Debian 6.0 and 6.2 with uek kernel | - |
| 2. | When QoS features are not supported by the kernel, `mlnx_qos` tool may crash. | - |

## 3.10 Flow Steering Known Issues

***Table 15 - Flow Steering Issues***

| Index | Description | Workaround |
|---|---|---|
| 1. | Flow Steering is disabled by default in firmware version < 2.32.5100. | To enable it, set the parameter below as follow: `log_num_mgm_entry_size should set to -1` |
| 2. | IPv4 rule with source IP cannot be created in SLES 11.x OSes. | - |
| 3. | RFS does not support UDP. | - |

## 3.11 Installation/Driver Start Known Issues

***Table 16 - Installation/Driver Start Issues***

| Index | Description | Workaround |
|---|---|---|
| 1. | Driver restart fails after mlnx_en installation if MLNX_OFED was previously installed on the system. | 1. Reboot the server after mlnx_en installation.<br>2. Answer **'n'** to the following question at the end of the installation: `"Do you wish to reload the driver now? (y/n) [y] "` n<br>3. Run `"apt-get remove --purge mlnx-ofed-kernel-utils"`<br>4. Reboot the server after mlnx_en installation. |
| 2. | Firmware update fails during mlx4_en standalone installation on guest machine connect to HCA in Passthrough mode. | - |
| 3. | On Debian 7.x, the Inbox driver is loaded instead of mlx4_en standalone drivers | Add the following lines to `"/etc/rc.local"`:<br>• `/sbin/modprobe -r mlx4_fc > /dev/null 2>&1`<br>• `/sbin/modprobe -r mlx4_en > /dev/null 2>&1`<br>• `/sbin/modprobe -r mlx4_ib > /dev/null 2>&1`<br>• `/sbin/modprobe -r mlx4_core > /dev/null 2>&1`<br>• `# load the new driver`<br>• `/sbin/modprobe mlx4_en > /dev/null 2>&1` |

*Table 16 - Installation/Driver Start Issues (Continued)*

| Index | Description | Workaround |
|---|---|---|
| 4. | mlx4_en standalone drivers are not available after installation on SLSE11 SP3 PPC64. | Run the installation with the "--disable-kmp" flag. |
| 5. | On RHEL 7, MLNX_EN driver v2.3-1.0.0 installation fails the first time. However, the driver is insatalled successfully on the second attempt. | To avoid failure on the first attempt, prior to running the installation, run:<br>`make -C /usr/src/kernels/$(uname -r) kernelrelease` |

## 3.12 Performance Known Issues

*Table 17 - Performance Known Issues*

| Index | Description | Workaround |
|---|---|---|
| 1. | On machines with irqbalancer daemon turned off, the default Ethernet interrupts will be routed to a single core which may cause overload and software/hardware lockups. | Execute the following script as root:<br>`set_irq_affinity.sh <interface or Eth device> [2nd interface or Eth device]` |
| 2. | Out of the box throughput performance in Ubuntu14.04 is not optimal and may achieve results below the line rate in 40GE link speed. | For additional performance tuning, please refer to Performance Tuning Guide. |
| 3. | UDP receiver throughput may be lower then expected, when running over mlx4_en Ethernet driver.<br>This is caused by the adaptive interrupt moderation routine, which sets high values of interrupt coalescing, causing the driver to process large number of packets in the same interrupt, leading UDP to drop packets due to overflow in its buffers. | Disable adaptive interrupt moderation and set lower values for the interrupt coalescing manually.<br>`ethtool -C <eth>X adaptive-rx off rx-usecs 64 rx-frames 24`<br><br>Values above may need tuning, depending the system, configuration and link speed. |
| 4. | Performance degradation might occur when bonding Ethernet interfaces on Centos 6.5 | - |

# 4 Bug Fixes History.

***Table 18 - Fixed Bugs List***

| # | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|
| 1. | LRO fixes and improvements for jumbo MTU. | 2.3-2.0.1 | 2.4-1.0.0.1 |
| 2. | Fixed a crash occurred when changing the number of rings (ethtool set-channels) when interface connected to netconsole. | 2.2-1.0.1 | 2.4-1.0.0.1 |
| 3. | Fixed ping issues with IP fragmented datagrams in MTUs 1600-1700. | 2.2-1.0.1 | 2.4-1.0.0.1 |
| 4. | The default priority to TC mapping assigns all priorities to TC0. This configuration achieves fairness in transmission between priorities but may cause undesirable PFC behavior where pause request for priority "n" affects all other priorities. | 2.3-1.0.1 | 2.4-1.0.0.1 |
| 5. | Fixed a crashing issue on non-SR-IOV systems related to VxLAN with dynamic port configuration. | 2.3-1.0.0 | 2.3-2.0.1 |
| 6. | Fixed "sleeping while atomic" error occurred when the driver ran many firmware commands simultaneously. | 2.3-1.0.0 | 2.3-2.0.1 |
| 7. | Fixed dmesg warnings: "NOHZ: local_softirq_pending 08". | 2.3-1.0.0 | 2.3-2.0.1 |
| 8. | Fixed erratic report of hardware clock which caused bad report of PTP hardware Time Stamping. | 2.3-1.0.0 | 2.3-2.0.1 |
| 9. | Fixed counter index allocation for VFs which enables Ethernet port statistics. | 2.3-1.0.0 | 2.3-2.0.1 |
| 10. | Fixed kernel panic on Debian-6.0.7 which occurred when the number of TX channels was set above the default value | 2.1-1.0.0 | 2.2-1.0.1 |
| 11. | Fixed wrong calculation of packet true-size reporting in LRO flow | 2.1-1.0.0 | 2.2-1.0.1 |
| 12. | Fixed a crash incidence which occurred when enabling Ethernet Time-stamping and running VLAN traffic | 2.0-3.0.0 | 2.2-1.0.1 |
| 13. | Fixed ALB bonding mode failure when enslaving Mellanox interfaces | 2.0-3.0.0 | 2.1-1.0.0 |
| 14. | Fixed leak of mapped memory | 2.0-3.0.0 | 2.1-1.0.0 |
| 15. | Fixed TX timeout issues | 1.5.10 | 2.0.-3.0.0 |
| 16. | Fixed Ethtool status report for Virtual Functions | 1.5.10 | 2.0.-3.0.0 |
| 17. | Fixed an issue of VLAN traffic over Virtual Machine in para-virtualized mode | 1.5.10 | 2.0.-3.0.0 |
| 18. | Fixed Ethtool operation crash while interface down | 1.5.10 | 2.0.-3.0.0 |
| 19. | Fixed driver unloading after installation.<br>Currently, the driver is reloaded after installation | 1.5.9 | 1.5.10 |
| 20. | Fixed an issue related to the configuration files.<br>Configuration files are no longer modified. | 1.5.9 | 1.5.10 |

***Table 18 - Fixed Bugs List***

| # | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|
| 21. | Fixed MTU size calculation | 1.5.9 | 1.5.10 |
| 22. | Fixed race in join/leave multicast group flow | 1.5.9 | 1.5.10 |
| 23. | Restored RX buffer pointer in case of failure | 1.5.9 | 1.5.10 |
| 24. | Fixed IPMI | 1.5.9 | 1.5.10 |
| 25. | Fixed BlueFlame race | 1.5.9 | 1.5.10 |
| 26. | Added sanity check for module parameters | 1.5.9 | 1.5.10 |
| 27. | Fixed lost of connectivity bug for iperf long run | 1.5.9 | 1.5.10 |
| 28. | Fixed bad TX completion notification request logic, could lead to TX timeout | 1.5.8.3 | 1.5.9 |
| 29. | Fixed PowerPC connectivity loss | 1.5.8.3 | 1.5.9 |
| 30. | Fixed possible transmit timeout under heavy traffic load | 1.5.8.2 | 1.5.8.3 |
| 31. | Fixed error flows in case of MSI-X assignment failures | 1.5.7.2 | 1.5.8.2 |
| 32. | Fixed loss of multicast traffic during addition of multicast addresses | 1.5.7.2 | 1.5.8.2 |
| 33. | Enabled legal MTU configuration prior to interface bring-up | 1.5.7.2 | 1.5.8.2 |
| 34. | Fixed WoL configuration behavior for single ported devices | 1.5.7 | 1.5.7.2 |
| 35. | Fixed the assignment of valid flow control parameters during interface initialization | 1.5.7 | 1.5.7.2 |
| 36. | Fixed qdisc behavior when transmitting packets with VLAN ID 0 | 1.5.7 | 1.5.7.2 |
| 37. | Fixed badly placed counters on ethool -S report | 1.5.7 | 1.5.7.2 |
| 38. | Fixed the crash occurred during initialization if an invalid mac address was burned on the card | 1.5.6 | 1.5.7 |
| 39. | Fixed a bug where interface link state was not updated correctly in certain scenarios | 1.5.6 | 1.5.7 |
| 40. | Fixed a crash occurred when LRO enabled via ethtool but disabled in module parameter | 1.5.6 | 1.5.7 |
| 41. | Fixed a warning when using netconsole | 1.5.6 | 1.5.7 |
| 42. | Fixed a crash occurred during driver initialization due to multicast list update | 1.5.6 | 1.5.7 |
| 43. | Fixed a bug that caused the driver not to work with MSI-X on RedHat EL4 systems | 1.5.1.3 | 1.5.6 |
| 44. | Fixed the driver functionality in case the device is enslaved to bond with VLANs over it | 1.5.1.3 | 1.5.6 |
| 45. | Fixed the crash occurred when sending malformed TCP packets | 1.5.1.3 | 1.5.6 |

***Table 18 - Fixed Bugs List***

| # | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|
| 46. | If the two ports of a ConnectX/ConnectX-2 device are bonded together and one port is closed, then the other port will cease carrying traffic. Fixed in this release. | 1.4.2 | 1.5.1.3 |
| 47. | Multiple error messages when working under bonding on some operating systems | 1.4.2 | 1.5.1.3 |
| 48. | Fixed a bug avoiding the driver to work with MSI-X on RedHat EL4 systems | | 1.4.2 |
| 49. | Enabled the driver to load on systems with 32 cores and higher | | 1.4.2 |
| 50. | The driver is being stuck if the HW/FW stops responding, reset is done instead | | 1.4.2 |
| 51. | Fixed recovery flows from memory allocation failures | | 1.4.2 |
| 52. | Corrected checksum reports on SLES11 systems | | 1.4.2 |
| 53. | When the system is low on memory, the driver now allocates smaller RX rings | | 1.4.2 |
| 54. | The driver now retries to obtain MSI-X vectors if the initial request is rejected by the OS | | 1.4.2 |

# 5 Change Log History

***Table 19 - Change Log History***

| Release | Category | Description |
|---|---|---|
| Rev 2.4-1.0.0.1 | mlx4_en net-device Ethtool | Added support for Ethtool speed control and advertised link mode. |
| | | Added ethtool txvlan control for setting ON/OFF hardware TX VLAN insertion: `ethtool -k txvlan [on/off]` |
| | | Ethtool report on port parameters improvements. |
| | | Ethernet TX packet rate improvements. |
| | Virtualization | VXLAN fixes and performance improvements. |
| | Ethernet net-device | New adaptive interrupt moderation scheme to improve CPU utilization. |
| Rev 2.3-2.0.1 | Reset Flow | Added support for Enhanced Error Handling for PCI (EEH), a recovery strategy for I/O errors that occur on the PCI bus. |
| | VXLAN | Bug Fixes, see Section 4, “Bug Fixes History.”, on page 14 |
| Rev 2.3-1.0.0 | Ethernet | Added support for arbitrary UDP port for VXLAN.<br>From upstream 3.15-rc1 and onward, it is possible to use arbitrary UDP port for VXLAN.<br>This feature requires firmware version 2.32.5100 or higher.<br>Additionally, the following kernel configuration option `CONFIG_MLX4_EN_VXLAN=y` must be enabled. |
| | | MLNX_EN no longer changes the OS sysctl TCP parameters. |
| Rev 2.2-1.0.1 | Reset Flow | Reset Flow is not activated by default. It is controlled by the mlx4_core`'internal_err_reset'` module parameter. |
| | Ethernet | Ethernet VXLAN support for kernels 3.12.10 or higher |
| | | Power Management Quality of Service: when the traffic is active, the Power Management QoS is enabled by disabling the CPU states for maximum performance. |
| | | Ethernet PTP Hardware Clock support on kernels/OSes that support it |
| | Performance | Out of the box performance improvements:<br>• Use of affinity hints (based on NUMA node of the device) to indicate the IRQ balancer daemon on the optimal IRQ affinity<br>• Improvement in buffers allocation schema (based on the hint above)<br>• Improvement in the adaptive interrupt moderation algorithm |

***Table 19 - Change Log History***

| Release | Category | Description |
|---|---|---|
| Rev 2.1-1.0. | mlnx_en | Added reporting autonegotiation support |
| | | Added Transmit Packet Steering (XPS) support |
| | | Added reporting 56Gbit/s link speed support |
| | | Added Receive Flow Steering (RFS) support in UDP |
| | | Added Low Latency Socket (LLS) support |
| | | Added check for dma_mapping errors |
| Rev 2.0-3.0.0 | Operating Systems | Additional OS support:<br>• SLES11SP3<br>• Fedora16, Fedora17 |
| | Hardware | Added ConnectX-3 Pro support |
| Rev 1.5.10 | General | Section 4, “Bug Fixes History.”, on page 14 |
| Rev 1.5.9 | Operating Systems | Added support for kernel.org 3.5 |
| | Performance | Improved latency by optimizing RX repost mechanism |
| Rev 1.5.8.3 | Operating Systems | Added support for RHEL6.3 |
| Rev 1.5.8.2 | Operating Systems | Added support for new kernels: 3.1, 3.2, 3.3 |
| Rev 1.5.8.2 | Performance | Moved to interrupt mode to handle TX completions |
| | | Added IRQ affinity control scripts (please see README file for more details) |
| | | Optimized Numa aware memory allocations |
| | | Optimized interrupt usage for TX/RX completions |
| | Installation | Added KMP compliant installation process |
| | Linux Tools | Added support for Ethtool |
| Rev 1.5.7.2 | Operating Systems | Added support for new OS's:<br>• RHEL6.2<br>• RHEL5.8<br>• SLES11SP2 |
| | Performance | Added recording RX queue for GRO packets |
| | | Added the usage of Toeplitz hash function for RSS calculation |
| | Reports/Statistics | Enabled RXHASH report on supported systems |

***Table 19 - Change Log History***

| Release | Category | Description |
|---|---|---|
| Rev 1.5.7 | Operating Systems | Added support for new OS's:<br>• RHEL6.1<br>• RHEL5.5<br>• RHEL5.7<br>• kernel.org (2.6.37, 2.6.38, 2.6.39, 3.0)<br>• RHEL6.1 KVM |
| | Performance | Improved performance on PPC systems (Using GRO where LRO is not efficient) |
| | | Added IPv6 support to LRO |
| | | Incremented number of TX and RX queues |
| | | Enabled NAPI usage at any given time |
| | | Enabled TX completions spread among multiple MSI-X vectors |
| | | Improved small packets packet rate |
| | | Added 40GigE support (including Ethtool report) |
| | | Added NUMA support |
| | | Added general performance improvements |
| Rev 1.5.6 | Operating Systems | Added support for new OS's:<br>• RHEL6.0<br>• RHEL5.6<br>• SLES11SP1<br>• kernel.org (2.6.35, 2.6.36) |

***Table 19 - Change Log History***

| Release | Category | Description |
|---|---|---|
| Rev 1.5.6 | Performance | Added blue flame support for kernels > 2.6.28 (improves TX latency by 0.4 usec) |
| | | Added RX acceleration feature that supports recvmsg and recvmmsg system calls. See MLNX_EN_Linux_README for further details. |
| | | Added option to use interrupts for TX completion (polling is the default) |
| | | Added option to disable NAPI (enabled by default) |
| | | Added support for control number of RX rings from module parameter |
| | | Added interrupt vector per each RX ring. See /proc/interrupts |
| | | Adaptive moderation improvements |
| | | Added system tuning option to achieve better performance (idle loop polling) |
| | Linux Tools | Added hardware revision report via Ethtool |
| | Multicast Filtering | Added exact match multicast filtering |
| | Driver Load | Link is brought up upon driver load |
| Rev 1.5.1.3 | Operating Systems | Added support for new OS's:<br>• RHEL5.5<br>• kernel.org (2.6.16 - 2.6.32) |
| | Performance | Added UDP RSS support (on ConnectX-2 HW only) |
| | | Improved VLAN tagging performance |
| | Linux Tools | Ethtool -e support |