



Red Hat Enterprise Linux (RHEL) 7.3 Driver Release Notes

RHEL 7.3

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 2017. 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®, 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 |
| List Of Tables | 4 |
| Chapter 1 Overview | 5 |
| 1.1 Supported HCAs Firmware Versions | 5 |
| 1.2 SR-IOV Support | 6 |
| 1.3 RoCE Support | 6 |
| 1.4 VXLAN Support | 6 |
| Chapter 2 Changes and New Features | 7 |
| Chapter 3 Known Issues | 8 |

List Of Tables

| | | |
|----------|--|---|
| Table 1: | Supported Uplinks to Servers | 5 |
| Table 2: | Supported HCAs Firmware Versions | 5 |
| Table 3: | SR-IOV Support | 6 |
| Table 4: | RoCE Support | 6 |
| Table 5: | VXLAN Support | 6 |
| Table 6: | Changes and New Features | 7 |
| Table 7: | Known Issues | 8 |

1 Overview

These are the release notes of Red Hat Enterprise Linux (RHEL) 7.3 Driver Release Notes. This document provides instructions on drivers for Mellanox Technologies ConnectX® based adapter cards with Red Hat Enterprise Linux (RHEL) 7.3 Inbox Driver environment.

This version supports the following uplinks to servers:

Table 1 - Supported Uplinks to Servers

| Uplink/HCAs | Uplink Speed | Supported Driver |
|---------------------------------|---|---|
| ConnectX®-4 | <ul style="list-style-type: none"> Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE^a, and 100GigE | mlx5_core (includes the ETH functionality as well), mlx5_ib |
| ConnectX®-4 Lx | <ul style="list-style-type: none"> Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, and 50GigE | mlx5_core (includes the ETH functionality as well) |
| Connect-IB® | <ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR10, FDR | mlx5_core, mlx5_ib |
| ConnectX®-3/ ConnectX®-3 Pro | <ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR10, FDR Ethernet: 10GigE, 40GigE and 56GigE^a | mlx4_core, mlx4_en, mlx4_ib |
| ConnectX®-2 | <ul style="list-style-type: none"> InfiniBand: SDR, DDR Ethernet: 10GigE, 20GigE | mlx4_core, mlx4_en, mlx4_ib |

a. 56GbE 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 HCAs Firmware Versions

Red Hat Enterprise Linux (RHEL) 7.3 Driver RHEL 7.3 supports the following Mellanox network adapter cards firmware versions:

Table 2 - Supported HCAs Firmware Versions

| HCA | Recommended Firmware Rev. | Additional Firmware Rev. Supported |
|-----------------|---------------------------|------------------------------------|
| Connect-IB® | 10.16.1006 | N/A |
| ConnectX®-4 Lx | 14.17.1010 | N/A |
| ConnectX®-4 | 12.17.1010 | N/A |
| ConnectX®-3 Pro | 2.40.5000 | N/A |
| ConnectX®-3 | 2.40.5000 | N/A |
| ConnectX®-2 | 2.9.1000 | N/A |

1.2 SR-IOV Support

Table 3 - SR-IOV Support

| Driver | Support |
|---|--------------------------------|
| mlx4_core, mlx4_en, mlx4_ib | Technical Preview ^a |
| mlx5_core (includes ETH functionality), mlx5_ib | Technical Preview ^a |

a. Technical Preview is not fully supported production feature.

1.3 RoCE Support

Table 4 - RoCE Support

| Driver | Support |
|-------------------|---------|
| mlx4 - RoCE v1/v2 | Yes |
| mlx5 - RoCE v1/v2 | Yes |

1.4 VXLAN Support

Table 5 - VXLAN Support

| Driver | Support |
|----------------------|-------------------|
| mlx4 - VXLAN offload | Yes |
| mlx5 - VXLAN offload | Yes (without RSS) |

2 Changes and New Features

Table 6 - Changes and New Features

| Feature/Change | Description |
|----------------|--|
| mlx5 | <ul style="list-style-type: none"> • Ethernet SRIOV • eSwitch - SR-IOV (VST, VGT, vf statistics) • Flow steering V1/V2 • RoCE v1/v2 • Multicast loopback prevention • RX/TX checksum offload • mlx5e Time Stamping & PTP • Drop counters • Raw Ethernet QP • CSUM in RX flow: IPoIB offloads with ConnectX-4 • Trap rules (sniffing mode) • Scatter FCS • Re-registration of MRs • Checksum complete • Cross-channel synchronization • Atomic operations • QoS DCBNL IEE ETS/PFC/ETS rate limit • WoL support via ethtool • VXLAN Support without RSS • Ethtool Vlan strip on/off • rx-fcs enable/disable • Ethtool EEPROM query • Ethtool LED blinking • Counters redesign • QoS TC/UP mapping support • Queue Counters • Mlx5e out of buffer counter • Striding RQ & fragmented WQEs • Use only close NUMA node for default RSS • Pport counters • Memory windows allocation • CQ tasklets • RSS Toeplitz hash function • XOR hash function • HW LRO • Change link speed • RX/TX pause setting (via ethtool) • Arbitrary sg list support • Offload arithmetic hardware capabilities |
| mlx4 | <ul style="list-style-type: none"> • RoCE LAG • RoCE VF LAG for RoCE V1 • Trap rule (sniffing mode) • Multicast loopback prevention • RoCE v2 • QinQ • QCN • TX Queue max rate |

3 Known Issues

The following table describes known issues in this release and possible workarounds.

Table 7 - Known Issues (Sheet 1 of 2)

| Internal Ref. | Description |
|---------------|---|
| BZ 1366935 | Description: Anaconda does not present IPoIB interfaces in TUI when installing the OS over PXE IPoIB. |
| | WA: N/A |
| | Keywords: PXE, IPoIB |
| 847472 | Description: Forcing off the VM while the ConnectIB® device is attached to that VM as a pass-through will cause error messages that will be reported in the dmesg. |
| | WA: Shutdown the VM instead of forcing off |
| | Keywords: Virtualization |
| BZ 1366256 | Description: Unfree resources after Anaconda installation will cause the machine to boot with asserts in mlx5_core driver. |
| | WA: N/A |
| | Keywords: Anaconda, mlx5 |
| 831038 | Description: On rare occasions, under heavy traffic and loading/unloading mlx4_en, mlx4_ib and ml4_core may cause VPD access failure. |
| | WA: N/A |
| | Keywords: mlx4 |
| BZ 1189428 | Description: kdump over NFS for mlx5 driver is not supported. |
| | WA: N/A |
| | Keywords: kdump, mlx5 |
| 634952 | Description: On HP P70, attaching mlx5 device as pass-through to VM is unsupported. When trying to do so, the following message is displayed: <i>“Device is ineligible for IOMMU domain attach due to platform RMRR requirement.”</i> |
| | WA: N/A |
| | Keywords: Virtualization, mlx5 |
| | Description: Atomic Operations in Connect-IB® are limited on little-endian machines (e.g. x86). |
| | WA: N/A |
| | Keywords: Atomic Operations, Connect-IB® |

Table 7 - Known Issues (Sheet 2 of 2)

| Internal Ref. | Description |
|---------------------------|--|
| BZ 1371075 | <p>Description: kernel oops related to mlx5 drivers might occur when booting with ConnectX-4 NIC.</p> <p>WA: Update firmware to 12.12.780 and newer from the Mellanox website (http://www.mellanox.com/page/firmware_table_ConnectX4IB)</p> <p>Keywords: Atomic Operations, Connect-IB®</p> |
| BZ 1177729, BZ 1252687 | <p>Description: ConnectX-2® Ethernet VF cannot be attached to the Guest.</p> <p>WA: N/A</p> <p>Keywords: Virtualization, ConnectX®</p> |