



# Red Hat Enterprise Linux (RHEL) 7.6 Driver Release Notes

---

RHEL 7.6

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT ("PRODUCT(S)") AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES "AS-IS" 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](http://www.mellanox.com)  
Tel: (408) 970-3400  
Fax: (408) 970-3403

© Copyright 2019. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, Connect-IB®, ConnectX®, CORE-Direct®, GPUDirect®, LinkX®, Mellanox Multi-Host®, Mellanox Socket Direct®, UFM®, and Virtual Protocol Interconnect® are registered trademarks of Mellanox Technologies, Ltd.

For the complete and most updated list of Mellanox trademarks, visit <http://www.mellanox.com/page/trademarks>.

All other trademarks are property of their respective owners.

# Table of Contents

<b>Table of Contents</b> .....	<b>3</b>
<b>List Of Tables</b> .....	<b>4</b>
<b>Chapter 1 Overview</b> .....	<b>5</b>
1.1 Supported HCAs Firmware Versions .....	5
1.2 SR-IOV Support .....	6
1.3 RoCE Support .....	6
1.4 VXLAN Support .....	6
1.5 DPDK Support .....	6
1.6 ASAP2 Open vSwitch Hardware Offloads Support .....	6
<b>Chapter 2 Changes and New Features</b> .....	<b>8</b>
<b>Chapter 3 Certifications</b> .....	<b>9</b>
3.1 RHEL NIC Qualification .....	9
<b>Chapter 4 Known Issues</b> .....	<b>10</b>

## 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:	DPDK Support .....	6
Table 7:	ASAP2 Open vSwitch Hardware Offloads Support .....	6
Table 8:	Changes and New Features .....	8
Table 9:	Known Issues .....	10

# 1 Overview

These are the release notes of Red Hat Enterprise Linux (RHEL) 7.6 Driver Release Notes. This document provides instructions on drivers for Mellanox Technologies ConnectX® based adapter cards with Red Hat Enterprise Linux (RHEL) 7.6 Inbox Driver environment.

This version supports the following uplinks to servers:

**Table 1 - Supported Uplinks to Servers**

Uplink/HCAs	Uplink Speed	Supported Driver
ConnectX®-5	<ul style="list-style-type: none"> <li>InfiniBand: SDR, QDR, FDR, FDR10, EDR</li> <li>Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE<sup>a</sup>, and 100GigE</li> </ul>	mlx5_core (includes the ETH functionality as well), mlx5_ib
ConnectX®-4	<ul style="list-style-type: none"> <li>InfiniBand: SDR, QDR, FDR, FDR10, EDR</li> <li>Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE<sup>a</sup>, and 100GigE</li> </ul>	mlx5_core (includes the ETH functionality as well), mlx5_ib
ConnectX®-4 Lx	<ul style="list-style-type: none"> <li>Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, and 50GigE</li> </ul>	mlx5_core (includes the ETH functionality as well)
Connect-IB®	<ul style="list-style-type: none"> <li>InfiniBand: SDR, QDR, FDR10, FDR</li> </ul>	mlx5_core, mlx5_ib
ConnectX®-3/ ConnectX®-3 Pro	<ul style="list-style-type: none"> <li>InfiniBand: SDR, QDR, FDR10, FDR</li> <li>Ethernet: 10GigE, 40GigE and 56GigE<sup>a</sup></li> </ul>	mlx4_core, mlx4_en, mlx4_ib
ConnectX®-2	<ul style="list-style-type: none"> <li>InfiniBand: SDR, DDR</li> <li>Ethernet: 10GigE, 20GigE</li> </ul>	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) 7.6 Driver supports the following Mellanox network adapter cards firmware versions:

**Table 2 - Supported HCAs Firmware Versions**

HCA	Recommended Firmware Rev.
Connect-IB®	10.16.1200
ConnectX®-5	16.24.1000
ConnectX®-4 Lx	14.24.1000
ConnectX®-4	12.24.1000
ConnectX®-3 Pro	2.42.5000
ConnectX®-3	2.42.5000
ConnectX®-2	2.9.1000

## 1.2 SR-IOV Support

**Table 3 - SR-IOV Support**

Driver	Support
mlx4_core, mlx4_en, mlx4_ib	ETH Infiniband - Technical Preview <sup>a</sup>
mlx5_core (includes ETH functionality), mlx5_ib	ETH Infiniband - Technical Preview <sup>a</sup>

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)

## 1.5 DPDK Support

**Table 6 - DPDK Support**

Driver	Support
mlx4	Yes
mlx5	Yes

## 1.6 ASAP<sup>2</sup> Open vSwitch Hardware Offloads Support

**Table 7 - ASAP<sup>2</sup> Open vSwitch Hardware Offloads Support**

Driver	Support
mlx4	No
mlx5	Yes <sup>a</sup>

a. Technical Preview is not fully supported production feature.

## 2 Changes and New Features

*Table 8 - Changes and New Features*

Driver/ Component	Feature/Change	Description
mlx5	Striding RQ/CQE Zipping User Controlled	A Multi-Packet receive WQ comes to improve performance and memory footprint. It reduces memory footprint by allowing each packet to consume a different number of strides instead of the whole WR.
	Support CQ Moderation	Enhances performance by moderating the number of cookies needed to create an event instead of application having to suffer from event per cookie.
	OVS Enhancements	<b>[ConnectX-5 and above]</b> Offloads TC VLAN push/pop using them as a hardware action during Flow Steering.
mlx4	Support CQ Moderation	Enhances performance by moderating the number of cookies needed to create an event instead of application having to suffer from event per cookie.
rdma-core	Version Update	Updated to version rdma-core-17.2-3.el7
mstflint	Version Update	Updated to version mstflint-4.9.0-3.el7
VMA	Version Update	Update to version libvma-8.6.10-1.el7

## 3 Certifications

### 3.1 RHEL NIC Qualification

As of RHEL 7.5, RHEL NIC qualification has passed successfully as described in:

[https://github.com/ctrautma/RHEL\\_NIC\\_QUALIFICATION/](https://github.com/ctrautma/RHEL_NIC_QUALIFICATION/)

Covering:

- ConnexX-3, ConnexX-4, ConnectX-5 adapter cards
- Kernel Data Path, OVS, OVS-DPDK
- PVP (Physical-Virtual-Physical), Performance (VSPerf), OVS functionality



## 4 Known Issues

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

**Table 9 - Known Issues (Sheet 1 of 2)**

Internal Ref.	RedHat Ref.	Description
1646807	1615865	<b>Description:</b> Due to a missing kernel support, creating 127 VFs on a ConnectX-5 adapter card fails. The failure may occur even with lower number of VFs (higher then 60).
		<b>Workaround:</b> N/A
		<b>Keywords:</b> ConnectX-5, VFs
1511227	-	<b>Description:</b> The POP datapath flows are run in the software while the Push flows in the hardware when sending packets over VST VXLAN using OVS.
		<b>Workaround:</b> Enable hw-tc-offload on uplink and representor For example: <ul style="list-style-type: none"> <li>• <code>ethtool -K enp139s0f0_0 hw-tc-offload on</code></li> <li>• <code>ethtool -K enp139s0f0 hw-tc-offload on</code></li> </ul>
		<b>Keywords:</b> OVS VXLAN
1591745	1624551	<b>Description:</b> Disabling SR-IOV or unbinding PFs per VFs while hairpin rules are active will result in Kernel crash.
		<b>Workaround:</b> Remove hairpin rules prior to disabling SR-IOV or unbinding PFs.
		<b>Keywords:</b> Kernel crash, hairpin
1510748	-	<b>Description:</b> RoCE is not functional in ConnectX4-Lx adapter cards when running in switchdev mode (for example: RDMA_CM, ibstat).
		<b>Workaround:</b> Disable encap on the eSwitch. For example: <code>devlink dev eswitch set pci/0000:24:00.0 encap disable</code>
		<b>Keywords:</b> OVS RDMA-CM ConnectX4-Lx
1284047	-	<b>Description:</b> Bandwidth degradations due to Page Table Isolation (PTI) Intel's CPU security fix.
		<b>Workaround:</b> PTI can be disabled in one of the following manners: <ul style="list-style-type: none"> <li>• Disable it during the runtime by writing 0 to <code>/sys/kernel/debug/x86/pti_enabled</code>.</li> <li>• Add "nopti" or "pti=off" to grub.conf</li> </ul>
		<b>Keywords:</b> Performance
1336618	-	<b>Description:</b> On rare occasions, under heavy traffic and loading/unloading the mlx4_en, mlx4_ib and ml4_core drivers may cause VPD access failure.
		<b>Workaround:</b> N/A
		<b>Keywords:</b> mlx4, Firmware

**Table 9 - Known Issues (Sheet 2 of 2)**

Internal Ref.	RedHat Ref.	Description
-	1189428	<p><b>Description:</b> kdump over mlx5 driver is not supported. However, it is supported on Ethernet interfaces but it requires preserving enough memory to support it.</p> <p><b>Workaround:</b> As kdump is supported on Ethernet interfaces, make sure to preserve enough memory by adding <code>crashkernel=512M</code> to the kernel command line.</p> <p><b>Keywords:</b> kdump, mlx5</p>
-	-	<p><b>Description:</b> Atomic Operations in Connect-IB® are limited on little-endian machines (e.g. x86).</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Atomic Operations in Connect-IB®</p>
-	1450329	<p><b>Description:</b> Occasionally, under special circumstances, Virtual Function can be removed.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> mlx5, SR-IOV</p>
-	1460579	<p><b>Description:</b> Bond creation on mlx4 Virtual Function interfaces might fail with error -524.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> mlx4, SR-IOV</p>
-	1462591	<p><b>Description:</b> InfiniBand mlx5 SR-IOV is not supported using OpenSM.</p> <p><b>Workaround:</b> Use the SM on the Mellanox Switch/UFM/mlnx_opensm</p> <p><b>Keywords:</b> mlx5, SR-IOV, OpenSM</p>