



# Red Hat Enterprise Linux (RHEL) 7.9 Driver Release Notes

RHEL 7.9

# Table of Contents

<b>1</b>	<b>Overview</b> .....	<b>4</b>
	Supported HCAs Firmware Versions.....	5
	SR-IOV Support.....	5
	RoCE Support .....	6
	VXLAN Support .....	6
	DPDK Support .....	6
	ASAP <sup>2</sup> Open vSwitch Hardware Offloads Support .....	6
<b>2</b>	<b>Changes and New Features</b> .....	<b>7</b>
<b>3</b>	<b>Known Inbox-Related Issues</b> .....	<b>8</b>

## List of Tables

Table 1: Supported Uplinks to Servers .....	4
Table 2: Supported HCAs Firmware Versions .....	5
Table 3: SR-IOV Support .....	5
Table 4: RoCE Support .....	6
Table 5: VXLAN Support .....	6
Table 6: DPDK Support .....	6
Table 7: ASAP <sup>2</sup> Open vSwitch Hardware Offloads Support .....	6
Table 8: Changes and New Features .....	7

# 1 Overview

These are the release notes of Red Hat Enterprise Linux (RHEL) 7.9 Driver Release Notes. This document provides instructions on drivers for Mellanox Technologies ConnectX® based adapter cards with Red Hat Enterprise Linux (RHEL) 7.9 Inbox Driver environment.

This version supports the uplinks to servers described in the table below.

**Table 1: Supported Uplinks to Servers**

HCA's	Uplink Speed	Supported Driver
ConnectX®-6	<ul style="list-style-type: none"> <li>InfiniBand: SDR, EDR, HDR</li> <li>Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 100GigE, 200 GigE (Alpha: Force mode)</li> </ul>	mlx5_core (includes the ETH functionality as well), mlx5_ib
BlueField® <sup>a</sup>	<ul style="list-style-type: none"> <li>Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, and 100GigE</li> </ul>	mlx5_core (includes the ETH functionality as well), mlx5_ib
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>b</sup> and 100 GigE</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>b</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)
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>b</sup></li> </ul>	mlx4_core, mlx4_en, mlx4_ib
Connect-IB®	<ul style="list-style-type: none"> <li>InfiniBand: SDR, QDR, FDR10, FDR</li> </ul>	mlx5_core, mlx5_ib

a. BlueField is supported as a standard ConnectX-5 Ethernet NIC only.

b. 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.

## Supported HCAs Firmware Versions

Red Hat Enterprise Linux (RHEL) 7.9 driver supports the following Mellanox network adapter cards firmware versions.

**Table 2: Supported HCAs Firmware Versions**

HCA	Recommended Firmware Rev.
ConnectX®-6	20.27.2008
BlueField®	18.27.2008
ConnectX®-5	16.27.2008
ConnectX®-4 Lx	14.27.2008
ConnectX®-4	12.27.2008
ConnectX®-3 Pro	2.42.5000
ConnectX®-3	2.42.5000
Connect-IB®	10.16.1200

## 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 a fully supported production feature.**

## RoCE Support

Table 4: RoCE Support

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

## VXLAN Support

Table 5: VXLAN Support

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

## DPDK Support

Table 6: DPDK Support

Driver	Support
mlx4	Yes
mlx5	Yes

## 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 a fully supported production feature.

---

## 2 Changes and New Features

Table 8: Changes and New Features

Component	Feature/Change	Description
mlx4, mlx5	General update	The driver's base Upstream kernel was not updated in this RHEL version; however, applicable bug fixes were backported.
RDMA user-space	rdma-core	Updated the rdma-core package to version v22.4-5.el7.

## 3 Known Inbox-Related Issues

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

Internal Ref.	Bugzilla Ref.	Description
2299041	1874101	<p><b>Description:</b> Race condition in the <code>mlx5e_rep_neigh_update</code> work queue sometimes triggers a kernel panic.</p> <p>When offloading encapsulation actions over the <code>mlx5</code> device using the <code>SwitchDev</code> in-kernel driver model in the Single Root I/O Virtualization (SR-IOV) capability, a race condition can happen in the <code>mlx5e_rep_neigh_update</code> work queue. Consequently, the system terminates unexpectedly with the kernel panic and the following message appears.</p> <pre>Workqueue: mlx5e mlx5e_rep_neigh_update [mlx5_core]</pre> <p><b>Workaround:</b> Make sure to update your system to the latest errata kernel, <code>kernel-3.10.0-1160.2.1.el7</code>, in which this issue has been fixed.</p> <p><b>Keywords:</b> OVS, SR-IOV, encap, <code>SwitchDev</code></p>
-	1816660	<p><b>Description:</b> When the <code>NUM_OF_VFS</code> parameter configured in the Firmware (using the <code>mstconfig</code> tool) is higher than 64, VF LAG mode will not be supported while deploying OVS offload.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> <code>ConnectX-5</code>, VF LAG, <code>ASAP<sup>2</sup></code>, <code>SwitchDev</code></p>
-	1816660	<p><b>Description:</b> An internal firmware error occurs either when attempting to disable single-root input/output virtualization, or when unbinding PF using a function (such as <code>ifdown</code> and <code>ip link</code>) under the following condition:</p> <p>Being in VF LAG mode in an OVS offload deployment, where at least one VF of any PF is still bound on the host or attached to a VM.</p> <p><b>Workaround:</b> Unbind or detach VFs before you perform these actions as follows.</p> <ol style="list-style-type: none"> <li>1. Shutdown and detach any VMs</li> <li>2. Remove VF LAG bond interface from OVS</li> <li>3. Unbind VFs, perform for each configured VF: <pre># echo &lt;VF PCIe BDF&gt; &gt; /sys/bus/pci/drivers/mlx5_core/ unbind</pre> </li> <li>4. Disable SR-IOV, perform for each PF: <pre># echo 0 &gt; /sys/class/net/&lt;PF&gt;/device/sriov_numvfs</pre> </li> </ol> <p><b>Keywords:</b> <code>ConnectX-5</code>, VF LAG, <code>ASAP<sup>2</sup></code>, <code>SwitchDev</code></p>



## Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation [“NVIDIA”] makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer [“Terms of Sale”]. NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

## Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

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

## Copyright

© 2020 NVIDIA Corporation. All rights reserved.

