

Red Hat Enterprise Linux (RHEL) 7.8 Driver Release Notes

RHEL 7.8

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
Tel: (408) 970-3400
Fax: (408) 970-3403

© Copyright 2020. 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

Table of Contents	3
List Of Tables	4
Chapter 1 Overview	5
1.1 Supported HCAs Firmware Versions	6
1.2 SR-IOV Support	6
1.3 RoCE Support	6
1.4 VXLAN Support	7
1.5 DPDK Support	7
1.6 ASAP2 Open vSwitch Hardware Offloads Support.	7
Chapter 2 Changes and New Features	8
Chapter 3 Known Issues	9

List Of Tables

Table 1:	Supported Uplinks to Servers	5
Table 2:	Supported HCAs Firmware Versions	6
Table 3:	SR-IOV Support.	6
Table 4:	RoCE Support	6
Table 5:	VXLAN Support.	7
Table 6:	DPDK Support.	7
Table 7:	ASAP2 Open vSwitch Hardware Offloads Support	7
Table 8:	Changes and New Features.	8
Table 9:	Known Issues	9

1 Overview

These are the release notes of Red Hat Enterprise Linux (RHEL) 7.8 Driver Release Notes. This document provides instructions on drivers for Mellanox Technologies ConnectX® based adapter cards with Red Hat Enterprise Linux (RHEL) 7.8 Inbox Driver environment.

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

Table 1 - Supported Uplinks to Servers

Uplink/HCAs	Uplink Speed	Supported Driver
ConnectX®-6	<ul style="list-style-type: none"> InfiniBand: SDR, EDR, HDR Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 100GigE, 200GbE (Alpha: Force Mode) 	mlx5_core (includes the ETH functionality as well), mlx5_ib
BlueField® ^a	<ul style="list-style-type: none"> Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, and 100GigE 	mlx5_core (includes the ETH functionality as well)
ConnectX®-5	<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE^b, and 100GigE 	mlx5_core (includes the ETH functionality as well), mlx5_ib
ConnectX®-4	<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE^b, 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)
ConnectX®-3/ ConnectX®-3 Pro	<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR10, FDR Ethernet: 10GigE, 40GigE and 56GigE^b 	mlx4_core, mlx4_en, mlx4_ib
Connect-IB®	<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR10, FDR 	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.

1.1 Supported HCAs Firmware Versions

Red Hat Enterprise Linux (RHEL) 7.8 driver supports the following Mellanox network adapter cards firmware versions:

Table 2 - Supported HCAs Firmware Versions

HCA	Recommended Firmware Rev.
ConnectX®-6	20.26.1040
BlueField®	18.26.1040
ConnectX®-5	16.26.1040
ConnectX®-4 Lx	14.26.1040
ConnectX®-4	12.26.1040
ConnectX®-3 Pro	2.42.5000
ConnectX®-3	2.42.5000
Connect-IB®	10.16.1200

1.2 SR-IOV Support

Table 3 - SR-IOV Support

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

a. Technical Preview is not fully supported production feature.

1.3 RoCE Support

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 ASAP² Open vSwitch Hardware Offloads Support

Table 7 - ASAP² Open vSwitch Hardware Offloads Support

Driver	Support
mlx4	No
mlx5	Yes ^a

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	Device out of Buffer Counter	Added as new ethtool counter that counts the packets dropped due to full internal receive queue. This counter is shown on “ethtool -S” output as a new counter named <code>dev_internal_queue_oob</code> .
	Flow Counter Preallocation	Added flow counters bulk allocation and pool, to improve the performance of flow counter acquisition.
	VLAN Push/Pop Offload (VGT)	Allowed offload of VLAN push/pop operations on both transmitted and received packets.
	VF LAG Load Balancing	Added support for load balancing over VF LAG configuration.
	General mlx5 Driver Update	No driver’s base Upstream kernel was updated in this RHEL version, however, applicable bug fixes were backported, in addition to the features listed in this table.
mlx4	General mlx4 Driver Update	No driver’s base Upstream kernel was updated in this RHEL version, however, applicable bug fixes were backported.
rdma-core	Version Update	Updated rdma-core version to v22.4-1.el7.
mstflint		Updated mstflint version to v4.13.3-2.el7.
VMA		Updated VMA version to v8.7.5-1.el7.

3 Known Issues

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

Table 9 - Known Issues (Sheet 1 of 2)

Internal Ref.	Bugzilla Ref.	Description
-	1816660	<p>Description: 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>Workaround: N/A</p> <p>Keywords: ConnectX-5, VF LAG, ASAP², SwitchDev</p>
-	1816660	<p>Description: 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: 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>Workaround: Unbind or detach VFs before you perform these actions as follows.</p> <ol style="list-style-type: none"> 1. Shutdown and detach any VMs 2. Remove VF LAG bond interface from OVS 3. Unbind VFs, perform for each configured VF: <pre># echo <VF PCIe BDF> > /sys/bus/pci/drivers/mlx5_core/unbind</pre> 4. Disable SR-IOV, perform for each PF: <pre># echo 0 > /sys/class/net/<PF>/device/sriov_numvfs</pre> <p>Keywords: ConnectX-5, VF LAG, ASAP², SwitchDev</p>
1511227	-	<p>Description: The POP datapath flows are run in the software while the Push flows in the hardware when sending packets over VST VXLAN using OVS.</p> <p>Workaround: Enable <code>hw-tc-offload</code> on uplink and representor</p> <p>For example:</p> <ul style="list-style-type: none"> • <code>ethtool -K enp139s0f0_0 hw-tc-offload on</code> • <code>ethtool -K enp139s0f0 hw-tc-offload on</code> <p>Keywords: OVS VXLAN</p>
1510748	-	<p>Description: RoCE is not functional in ConnectX4-Lx adapter cards when running in switchdev mode (for example: <code>RDMA_CM</code>, <code>ibstat</code>).</p> <p>Workaround: Disable <code>encap</code> on the eSwitch.</p> <p>For example: <code>devlink dev eswitch set pci/0000:24:00.0 encap disable</code></p> <p>Keywords: OVS RDMA-CM ConnectX4-Lx</p>

Table 9 - Known Issues (Sheet 2 of 2)

Internal Ref.	Bugzilla Ref.	Description
1284047	-	Description: Bandwidth degradations due to Page Table Isolation (PTI) - Intel's CPU security fix.
		Workaround: PTI can be disabled in one of the following manners: <ul style="list-style-type: none"> • Disable it during the runtime by writing 0 to /sys/kernel/debug/x86/pti_enabled. • Add "nopti" or "pti=off" to grub.conf
		Keywords: Performance
1336618	-	Description: On rare occasions, under heavy traffic and loading/unloading the mlx4_en, mlx4_ib and ml4_core drivers may cause VPD access failure.
		Workaround: N/A
		Keywords: mlx4, Firmware
1775867	1727593	Description: Syndrome (0x563e2f) followed by kernel panic during VF mirroring might be observe.
		Workaround: N/A
		Keywords: VF Mirroring, syndrome (0x563e2f)
-	1189428	Description: kdump over mlx5 driver is not supported. However, it is supported on Ethernet interfaces but it requires preserving enough memory to support it.
		Workaround: As kdump is supported on Ethernet interfaces, make sure to preserve enough memory by adding <code>crashkernel=512M</code> to the kernel command line.
		Keywords: kdump, mlx5
-	1462591	Description: InfiniBand mlx5 SR-IOV is not supported using OpenSM.
		Workaround: Use the SM on the Mellanox Switch/UFM/mlnx_opensm
		Keywords: mlx5, SR-IOV, OpenSM