

Red Hat Enterprise Linux (RHEL) 7.7 Driver Release Notes

RHEL 7.7

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 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

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	10

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	10

1 Overview

These are the release notes of Red Hat Enterprise Linux (RHEL) 7.7 Driver Release Notes. This document provides instructions on drivers for Mellanox Technologies ConnectX® based adapter cards with Red Hat Enterprise Linux (RHEL) 7.7 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

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) 7.7 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

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 (Sheet 1 of 2)

Driver/ Component	Feature/Change	Description
General	ConnectX®-6 adapter support	Added ConnectX-6 device ID for hardware enablement and testing.
	BlueField® adapter support	Added BlueField device ID for hardware enablement and testing - at Technical Preview level (verified with ConnectX-5 functionality of BlueField)
mlx5	vNIC steering drop statistics support	Added a counter called “Rx steering missed dropped packets” - counts packets which were dropped due to a miss on NIC Rx steering rules. This counter will be shown on ethtool as a counter called rx_steer_missed_packets.
	Implement fragmented completion queue (CQ)	Added a new scheme of fragmented CQ to allocate fragmented buffers rather than contiguous ones. The aim of this scheme is to avoid malfunction when the memory is fragmented or the system is low on memory.
	IB Device Memory support (MEMIC)	User space applications are now allowed to use the internal device memory in order to improve performance.
	Tunnel protocol (MPLSoGRE, MPLSoUDP, VXLan) RX/TX decap/encap offload	The ability to create flow actions which can change packet headers is now exposed.
	Flow counter support	Added support for flow counters on top of generic verbs counters interface.
	DEVX interface	DEVX enables direct access from the user space to the mlx5 device driver by using the KABI mechanism. It minimizes the dependency between user-space and kernel, such that user-space functionality can be added with minimal to none kernel changes.
	RoCE ICRC counter exposure	Adding support to query the counter that counts the RoCE packets with corrupted ICRC (Invariant Cyclic Redundancy Code).
	IB device name	Allowing to rename the user visible IB device name from vendor specific name (e.g. mlx5_0) to any other name.

Table 8 - Changes and New Features (Sheet 2 of 2)

Driver/ Component	Feature/Change	Description
mlx5	OVS offload enhancements	<ul style="list-style-type: none"> • Added IB representor when in switchdev mode • Support for setup/match of tos and ttl TC offloading of IP tunnels • Direct VF Mirroring • Remote Mirroring - added support for additional mirroring output in SwitchDev mode. The mirroring port may be a local or a remote VF, using VXLAN or GRE encapsulations. • Slow path performance • Batched reading of flow counters • QinQ support - QinQ, also known as Stacked VLAN or Double VLAN, allows multiple VLAN tags in an Ethernet frame to achieve traffic isolation. • VF LAG - added support for High Availability and load balancing for Virtual Functions of different physical ports in SwitchDev SR-IOV mode. • GRE HW offloading - ability to offload TC filters set on GRE interfaces. • VXLAN/GRE Tunneling over VLAN - added support for VXLAN and GRE tunnel encap/decap offload over Ethernet tagged packets. • VLAN Rewrite - added support for offloading VLAN ID modify operation, allowing the user to replace the VLAN tag of the incoming frame with a user-specified VLAN tag value.
mlx4	Configuration parameters setting	Configuration parameters setting can now be performed via devlink
rdma-core	Version Update	Updated to version rdma-core-22.1-3.el7
mstflint	Version Update	Updated to version mstflint-4.11.0-5.el7
VMA	Version Update	Update to version libvma-8.7.5-1.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.	RedHat Ref.	Description
1511227	-	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.
		Workaround: Enable hw-tc-offload on uplink and representor For example: <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>
		Keywords: OVS VXLAN
1510748	-	Description: RoCE is not functional in ConnectX4-Lx adapter cards when running in switchdev mode (for example: RDMA_CM, ibstat).
		Workaround: Disable encap on the eSwitch. For example: <code>devlink dev eswitch set pci/0000:24:00.0 encap disable</code>
		Keywords: OVS RDMA-CM ConnectX4-Lx
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 <code>/sys/kernel/debug/x86/pti_enabled</code>. • Add "nopti" or "pti=off" to grub.conf
		Keywords: Performance
1336618	-	Description: On rare occasions, under heavy traffic and loading/unloading the <code>mlx4_en</code> , <code>mlx4_ib</code> and <code>ml4_core</code> 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: <code>kdump</code> over <code>mlx5</code> driver is not supported. However, it is supported on Ethernet interfaces but it requires preserving enough memory to support it.
		Workaround: As <code>kdump</code> is supported on Ethernet interfaces, make sure to preserve enough memory by adding <code>crashkernel=512M</code> to the kernel command line.
		Keywords: <code>kdump</code> , <code>mlx5</code>

Table 9 - Known Issues (Sheet 2 of 2)

Internal Ref.	RedHat Ref.	Description
-	1462591	<p>Description: InfiniBand mlx5 SR-IOV is not supported using OpenSM.</p> <p>Workaround: Use the SM on the Mellanox Switch/UFM/mlnx_opensm</p> <p>Keywords: mlx5, SR-IOV, OpenSM</p>