Red Hat Enterprise Linux (RHEL) 7.5-ALT Driver Release Notes

RHEL 7.5-ALT

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

Table of Contents

List Of Tables

1 Overview

These are the release notes of Red Hat Enterprise Linux (RHEL) 7.5-ALT Driver for Arm and PPC architectures tested on IBM-Power8 and Cavium-Arm.

This document provides instructions on drivers for Mellanox Technologies ConnectX® based adapter cards with Red Hat Enterprise Linux (RHEL) 7.5-ALT 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	• InfiniBand: SDR, QDR, FDR, FDR10, EDR • Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE[a], and 100GigE	mlx5_core (includes the ETH functionality as well), mlx5_ib
ConnectX®-4	• InfiniBand: SDR, QDR, FDR, FDR10, EDR • Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE[a], and 100GigE	mlx5_core (includes the ETH functionality as well), mlx5_ib
ConnectX®-4 Lx	• Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, and 50GigE	mlx5_core (includes the ETH functionality as well)
Connect-IB®	• InfiniBand: SDR, QDR, FDR10, FDR	mlx5_core, mlx5_ib
ConnectX®-3/ ConnectX®-3 Pro	• InfiniBand: SDR, QDR, FDR10, FDR • Ethernet: 10GigE, 40GigE and 56GigE[a]	mlx4_core, mlx4_en, mlx4_ib
ConnectX®-2	• 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) 7.5 Driver RHEL 7.5-ALT 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®	10.16.1200	N/A
ConnectX®-5	16.22.1002	N/A
ConnectX®-4 Lx	14.22.1002	N/A
ConnectX®-4	12.22.1002	N/A
ConnectX®-3 Pro	2.42.5000	N/A
ConnectX®-3	2.42.5000	N/A

Table 2 - Supported HCAs Firmware Versions

HCA	Recommended Firmware Rev.	Additional Firmware Rev. Supported
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	• 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

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	Technical Preview[a]
mlx5 - VXLAN offload	Yes (without RSS)

a. Technical Preview is not fully supported production feature.

1.5 Open vSwitch Hardware Offloads Support

Table 6 - Open vSwitch Hardware Offloads Support

Driver	Support
mlx4	No
mlx5	Yes

2 Changes and New Features

Table 7 - Changes and New Features

Driver	Feature/Change	Description
mlx5	RoCE Diagnostics and ECN Counters	Added support for additional RoCE diagnostics and ECN congestion counters under /sys/class/infiniband/mlx5_0/ports/1/hw_counters/ directory.
	Local Loopback	Improved performance by rendering Local loopback (unicast and multicast) disabled by mlx5 driver by default while local loopback is not in use. The mlx5 driver keeps track of the number of transport domains that are opened by user-space applications. If there is more than one user-space transport domain open, local loopback will automatically be enabled.
	Enhanced IPoIB	Added support for Enhanced IPoIB feature, which enables better utilization of features supported in ConnectX-4 adapter cards, by optimizing IPoIB data path and thus, reaching peak performance in both bandwidth and latency. Enhanced IPoIB is enabled by default
	Precision Time Protocol (PTP)	Added support for PTP feature in IPoIB offloaded devices. This feature allows for accurate synchronization between the distributed entities over the network. The synchronization is based on symmetric Round Trip Time (RTT) between the master and slave devices. The feature is enabled by default.
	Tag Matching Offload	Added support for hardware Tag Matching offload with Dynamically Connected Transport (DCT).
	GRE Tunnel Offloads	Added support for the following GRE tunnel offloads: • TSO over GRE tunnels • Checksum offloads over GRE tunnels • RSS spread for GRE packets
	Dropless Receive Queue (RQ)	Added support for the driver to notify the FW when SW receive queues are overloaded.
	OVS enhancements	• Ipv6 support for VXLAN offloads (encap/decap) • Encapsulation flow enhancement • Expose min-inline mode to libmlx5 • mlx5 offloading of TC pedit (header re-write) action • mlx5 neigh update • Add support for drop action for steering rules • Header re-write • Flow counters
mlx4	RSS Support	RSS (Receive Side Scaling) technology allows to spread incoming traffic between different receive descriptor queues. Assigning each queue to different CPU cores allows to better load balance the incoming traffic and improve performance.

3 Known Issues

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

Table 8 - Known Issues

Internal Ref.	Bugzilla Ref.	Description
1336723	-	**Description:** NVMe IO may fail due to a block layer merge BIOS that does not comply to the virtual boundary.
		Workaround: Disable all merge tries, run "`echo 2 > /sys/block/<dev>/queue/nomerges`".
		Keywords: mlx5, NVMe
1336728	1468598	**Description:** When using a large block size (over 1G) iSER, memreg may fail and cause an IO error.
		Workaround: N/A
		Keywords: mlx5, iSER