

SUSE Linux Enterprise Server (SLES) 15 SP1 Inbox Driver Release Notes

SLES 15 SP1

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, Mellanox Open Ethernet®, LinkX®, Mellanox Spectrum®, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, ONE SWITCH. A WORLD OF OPTIONS®, Open Ethernet logo, Spectrum logo, Switch-IB®, SwitchX®, 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	5
1.2 SR-IOV Support	6
1.3 RoCE Support	6
1.4 VXLAN Support	6
1.5 DPDK Support	6
1.6 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	5
Table 3:	SR-IOV Support.	6
Table 4:	RoCE Support	6
Table 5:	VXLAN Support.	6
Table 6:	DPDK Support.	6
Table 7:	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 SUSE Linux Enterprise Server (SLES) 15 SP1 Inbox Driver. This document provides instructions on drivers for Mellanox Technologies ConnectX® based adapter cards with SUSE Linux Enterprise Server (SLES) 15 SP1 Inbox Driver environment.

This version supports the following uplinks to servers:

Table 1 - Supported Uplinks to Servers

HCA	Uplink Speed	Supported Driver
ConnectX®-5	<ul style="list-style-type: none"> 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	<ul style="list-style-type: none"> 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	<ul style="list-style-type: none"> Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, and 50GigE 	mlx5_core (includes the ETH functionality as well)
Innova™ IPsec EN	<ul style="list-style-type: none"> Ethernet: 10GigE, 40GigE 	mlx5_core (includes the ETH functionality as well)
Connect-IB®	<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR10, FDR 	mlx5_core, mlx5_ib
ConnectX®-3/ConnectX®-3 Pro	<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR10, FDR Ethernet: 10GigE, 40GigE and 56GigE^a 	mlx4_core, mlx4_en, mlx4_ib
BlueField	<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE, and 100GigE 	mlx5_core (includes the ETH functionality as well), mlx5_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

SUSE Linux Enterprise Server (SLES) 15 SP1 Inbox Driver SLES 15 SP1 supports the following Mellanox network adapter cards firmware versions:

Table 2 - Supported HCAs Firmware Versions

HCA	Recommended Firmware Rev.
Connect-IB®	10.16.1200
BlueField (technical preview)	18.25.1020
ConnectX®-5	16.25.1020

Table 2 - Supported HCAs Firmware Versions

HCA	Recommended Firmware Rev.
Innova™ IPsec EN	14.25.1020
ConnectX®-4 Lx	14.25.1020
ConnectX®-4	12.25.1020
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

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 Open vSwitch Hardware Offloads Support

Table 7 - Open vSwitch Hardware Offloads Support

Driver	Support
mlx4	No
mlx5	Yes

2 Changes and New Features

Table 8 - Changes and New Features

Feature/Change		Description
General	BlueField SmartNIC	Added support for Mellanox BlueField system-on-a-chip (SoC) device, optimized for NVMe storage systems, Network Functions Virtualization (NFV), security systems, and embedded appliances. Currently supported as a regular NIC (with functionality equivalent to that of ConnectX-5).
mlx5	Drop rules counters	Added the support for new counters for drop rules configured in eSwitch egress and ingress ACL.
	Explicit congestion notification for RoCE	Added the support for ECN into mlx5 driver for RDMA RoCE traffic and allows to ECN parameter configuration for mlx5 driver.
	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.
	dscp trust state and dscp2prio mapping	Added the option to enable PFC based on the DSCP value. Using this solution, VLAN headers will no longer be mandatory for use.
	OVS Enhancements	[ConnectX-5 and above] Offloads TC VLAN push/pop using them as a hardware action during Flow Steering.
	IB representor in switchdev mode	Adding support for an IB (RAW Ethernet only) device representor which is created when the user switches to switchdev mode.
	IB Device Memory support (MEMIC)	This capability can decrease latency of a trading operation where data is already allocated in the device memory, saving the PCI round trip when performing the following: <ul style="list-style-type: none"> • Atomic operations on semaphores remotely • Modification by the device on received traffic that should be transmitted directly after this modification.
	Introduce DEVX interface	DEVX enables direct access from the user space area to the mlx5 device driver by using the KABI mechanism.
	UDP GSO	This feature enables UDP GSO support
	Q-in-Q Support	QinQ, also known as Stacked VLAN or Double VLAN, allows multiple VLAN tags in an Ethernet frame to achieve traffic isolation.
mlx4	rdma-core	Updated to version rdma-core-22.1-2.10
	mstflint	Updated to version mstflint-4.11.0-2.3
	libvma	Updated to version libvma-8.7.1-4.24
	dpdk	Updated to version dpdk-18.11-2.29

3 Known Inbox-Related Issues

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

Table 9 - Known Issues

Internal Ref.	Description
1284059	<p>Description: BW degradations due to Page Table Isolation (PTI) Intel's CPU security fix.</p> <p>Workaround: Disable the PTI at run time by:</p> <ul style="list-style-type: none"> • writing 0 to <code>/sys/kernel/debug/x86/pti_enabled</code>. <p>or</p> <ul style="list-style-type: none"> • adding "nopti" or "pti=off" to <code>grub.conf</code> <p>Keywords: Performance</p>