Red Hat Enterprise Linux (RHEL) Driver Release Notes

RHEL 8.3

Table of Contents

List of Tables

1 Overview

These are the release notes of Red Hat Enterprise Linux (RHEL) 8.3 Driver Release Notes. This document provides instructions on drivers for Mellanox Technologies ConnectX® based adapter cards with Red Hat Enterprise Linux (RHEL) 8.3 Inbox Driver environment.

This version supports the following uplinks to servers:

Table 1: Supported Uplinks to Servers

HCAs	Uplink Speed	Supported Driver
ConnectX®-6 Dx	• Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 100GigE and 200 GigE	mlx5_core (includes the ETH functionality as well), mlx5_ib
ConnectX®-6	• InfiniBand: SDR, EDR, HDR • Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE and 100GigE	mlx5_core (includes the ETH functionality as well), mlx5_ib
BlueField®[a]	• Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, and 100GigE	mlx5_core (includes the ETH functionality as well)
Innova™ IPsec EN	• Ethernet: 10GigE, 40GigE	mlx5_core (includes the ETH functionality as well)
ConnectX®-5	• 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	• 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	• Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, and 50GigE	mlx5_core (includes the ETH functionality as well)
ConnectX®-3/ ConnectX®-3 Pro	• InfiniBand: SDR, QDR, FDR10, FDR • Ethernet: 10GigE, 40GigE and 56GigE[b]	mlx4_core, mlx4_en, mlx4_ib
Connect-IB®	• 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.

Supported HCAs Firmware Versions

Red Hat Enterprise Linux (RHEL) 8.3 driver supports the following Mellanox network adapter cards firmware versions:

Table 2: Supported HCAs Firmware Versions

HCA	Recommended Firmware Rev.
ConnectX®-6 Dx	22.28.2006
ConnectX®-6	20.28.2006
BlueField® (Technical Preview)	18.28.2006
ConnectX®-5	16.28.2006
ConnectX®-4 Lx	14.28.2006
ConnectX®-4	12.28.2006
ConnectX®-3 Pro	2.42.5000
ConnectX®-3	2.42.5000
Innova™ IPsec EN	14.22.1002
Connect-IB®	10.16.1002

SR-IOV Support

Table 3: SR-IOV Support

Driver	Support	Notes
mlx4_core, mlx4_en, mlx4_ib	Eth InfiniBand: Technical Preview	Running InfiniBand (IB) SR-IOV requires IB Virtualization support on the OpenSM (Session Manager). This capability is supported only on OpenSM provided by Mellanox, that is not available Inbox. This support can be achieved by running the highest-priority OpenSM on a Mellanox switch in an IB fabric. The switch SM can support this feature by enabling the virt flag (# ib sm virt enable). **Note**: This capability is not tested over Inbox environment and considered Tech Preview.
mlx5_core (includes ETH functionality), mlx5_ib	Eth InfiniBand: Technical Preview	

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	Mellanox PMD is enabled by default.
mlx5	Mellanox PMD is enabled by default.

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

Component	Feature/Change	Description
mlx5	Connection Tracking Offload	Added support for offloading TC filters containing connection tracking matches and actions.
	Kernel Software Steering Remote Mirroring	Added kernel support for OVS remote mirroring to allow hardware traffic mirroring to multiple ports.
	Rx Reporter in Devlink Health	Added support for monitoring and recovering from errors that occur on the Rx queue, such as CQE errors and timeout.
	RoCE Disablement via devlink	Added the option to disable RoCE traffic handling. This enables forwarding of traffic over UDP port 4791 that is handled as RoCE traffic when RoCE is enabled. When RoCE is disabled, there is no GID table, only Raw Ethernet QP type is supported and RoCE traffic is handled as regular Ethernet traffic. Use the `devlink` utility for disabling RoCE support.
	kTLS TX Support for ConnectX-6 Dx	Added support for hardware offload encryption of kTLS TX traffic to improve performance.
	Devlink Health State Notifications	Added support for receiving notifications on devlink health state changes when an error is reported or recovered by one of the reporters. These notifications can be seen using the userspace `devlink monitor` command.
	Flow Autogroup Default via Devlink	Added a devlink parameter to control the number of large groups in an auto-grouped flow table. The default value is 15, and the range is between 1 and 1024.
	General Driver Update	Aligned the mlx5 driver to the Linux upstream kernel driver version 5.6
mlx4	General Driver Update	Aligned the mlx4 driver to the Linux upstream kernel driver version 5.6
rdma-core	RDMA user-space	Updated the RDMA package to version 29.0-3.el8
mstflint	mstflint user-space	Updated mstflint package to version 4.14.0-1.el8
libvma	VMA	Updated VMA package to version 9.0.2-1.el8
ucx	UCX	Updated UCX package to version 1.8.0-1.el8

3 Certifications

RHEL NIC Qualification

RHEL 8.0, Successfully passed RHEL NIC qualification has passed successfully as described in: https://github.com/ctrautma/RHEL_NIC_QUALIFICATION/tree/8.0-Beta

Covering:

- ConnextX-4 Lx and ConnectX-5 adapter cards
- OVS functional, OVS non-offload, OVS-offload, OVS-DPDK

4 Known Inbox-Related Issues

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

Internal Ref.	Bugzilla Ref.	Description
2345747	1890261	**Description:** RHEL installer fails to start when InfiniBand network interfaces are configured using installer boot options.
		Workaround: Create a new installation media including the updated Anaconda and NetworkManager packages, using the Lorax tool. For more information on how to do so, please see here.
		Keywords: PXE, IPoIB, InfiniBand
	1816660	**Description:** When the NUM_OF_VFS parameter configured in the Firmware (using the mstconfig tool) is higher than 64, VF LAG mode will not be supported while deploying OVS offload.
		Workaround: N/A
		Keywords: ConnectX-5, VF LAG, ASAP2, SwitchDev
	1816660	**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 ifdown and ip link) 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.

		Workaround: Unbind or detach VFs before you perform these actions as follows.
		1. Shutdown and detach any VMs.
		2. Remove VF LAG bond interface from OVS.
		3. Unbind VFs, perform for each configured VF:
		`# echo <VF PCIe BDF> > /sys/bus/pci/drivers/mlx5_core/unbind`
		4. Disable SR-IOV, perform for each PF:
		`# echo 0 > /sys/class/net/<PF>/device/sriov_numvfs`

		Keywords: ConnectX-5, VF LAG, ASAP2, SwitchDev
1284047	-	**Description:** BW degradations due to PTI (Page Table Isolation) in Intel's CPU security fix.
		Workaround: PTI can be disabled in run time by writing 0 to /sys/ kernel/debug/x86/pti_enabled. Another option is adding "nopti" or "pti=off" to grub.conf.
		Keywords: Performance

Internal Ref.	Bugzilla Ref.	Description
1610281	-	**Description:** Setting speed to 56Gb/s on ConnectX-4 causes FW syndrome (0x1a303e).
		Workaround: N/A
		Keywords: ConnectX-4, syndrome
1609804	-	**Description:** Kernel panic during MTU change under stress traffic.
		Workaround: N/A
		Keywords: Panic, MTU
1578022	-	**Description:** OVS offload: fragmented traffic is not offload. When sending traffic with packets bigger than MTU, traffic runs but is not offloaded.
		Workaround: N/A
		Keywords: OVS offload, fragmentation

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

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com