



Mellanox OFED for FreeBSD for ConnectX-4/ConnectX-4 Lx/ ConnectX-5 Release Note

Rev 3.5.0



NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT ("PRODUCT(S)") AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES "ASIS" 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) 9703400
Fax: (408) 9703403

© Copyright 2018. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, ConnectIB®, ConnectX®, COREDirect®, GPUDirect®, LinkX®, Mellanox MultiHost®, 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
Release Update History	5
Chapter 1 Introduction	6
1.1 Supported Platforms and Operating Systems	6
1.2 Supported Adapters Firmware Versions	6
Chapter 2 Changes and New Features in Rev 3.5.0	7
Chapter 3 Known Issues	8
Chapter 4 Bug Fixes	11
Chapter 5 Change Log History	13



List of Tables

Table 1:	Release Update History	5
Table 2:	Supported Platforms and Operating Systems	6
Table 3:	Changes and New Features in Rev 3.5.0	7
Table 4:	Known Issues	8
Table 5:	Bug Fixes	11
Table 6:	Change Log History	13

Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 3.5.0	January 7, 2019	Updated the description of the Ethernet Counters entry under Table 3, “Changes and New Features in Rev 3.5.0,” on page 7.
	November 29, 2018	Initial release of this version.

1 Introduction

These are the release notes for Mellanox Technologies' driver for FreeBSD Rev 3.5.0 driver kit for Mellanox ConnectX®-4, ConnectX®-4 Lx, ConnectX®-5, ConnectX®-5 Ex adapter cards supporting the following uplinks to servers:

Uplink/HCAs	Driver Name	Uplink Speed
ConnectX®-4	mlx5	<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE^a, and 100GigE
ConnectX®-4 Lx		<ul style="list-style-type: none"> Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, and 50GigE
ConnectX®-5/ ConnectX®-5 Ex		<ul style="list-style-type: none"> InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, and 100GigE

- a. 56 GbE 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 Platforms and Operating Systems

The following are the supported OSs in Mellanox OFED for FreeBSD for ConnectX-4/ConnectX-4 Lx/ConnectX-5 Rev 3.5.0:

Table 2 - Supported Platforms and Operating Systems

Operating System	Platform
FreeBSD 12	AMD64/x86_64

1.2 Supported Adapters Firmware Versions

Mellanox OFED for FreeBSD Rev 3.5.0 supports the following Mellanox network adapter cards:

Supported Adapters	Current Firmware Rev.
ConnectX®-4	12.24.1000
ConnectX®-4 Lx	14.24.1000
ConnectX-5/ConnectX-5 Ex	16.24.1000

2 Changes and New Features in Rev 3.5.0

Table 3 - Changes and New Features in Rev 3.5.0

Category	Description
Relaxed Ordering	Added support for configuring PCIe packet write ordering via sysctl.
Enhanced Transmission Selection (ETS)	Added support for setting the bandwidth limit as a ratio rather than in bits per second. The ratio must be an integer number between 1 and 100, inclusive. This feature also enables setting a minimal BW guarantee on traffic classes (TCs).
Ethernet Counters	Added support for the following new counters: <ul style="list-style-type: none">• tx_jumbo_packets• rxstat0.bytes• txstat0tc0.bytes

For additional information on the new features, please refer to the User Manual.

3 Known Issues

The following is a list of general limitations and known issues of the various components of this Mellanox OFED for FreeBSD release.

Table 4 - Known Issues

Internal Ref.	Issue
1320335	Description: When Witness is enabled, the following message may appear in logs: “lock order reversal in mlx5_en_rx and in_pcb/tcp_input”.
	Workaround: N/A
	Keywords: Witness, LOR
	Discovered in Release: 3.5.0
1554671	Description: mlx5ib unload fails while OpenSM is running in background.
	Workaround: Kill OpenSM in order for the driver to unload.
	Keywords: mlx5ib, OpenSM, RDMA
	Discovered in Release: 3.5.0
1578093	Description: ibstat tool shows the wrong value of “rate” after unplugging the cable from the HCA.
	Workaround: N/A
	Keywords: ibstate, rate
	Discovered in Release: 3.5.0
1581628	Description: Driver unload hangs while RDMA user space application is running.
	Workaround: Close all RDMA user space running applications prior to unloading the driver.
	Keywords: RDMA, driver unload
	Discovered in Release: 3.5.0
1439351	Description: Link local GIDs are dysfunctional when IPv6 address is configured for the first time.
	Workaround: Set the net device state to “up”. For example: # <code>ifconfig mce0 up</code>
	Keywords: RoCE, IPv6
	Discovered in Release: 3.4.2
1435021	Description: All Rx priority pause counters values increase when Rx global pause is enabled.
	Workaround: Ignore Rx priority pause counters when Rx global pause is enabled.
	Keywords: Rx pause counters, priority
	Discovered in Release: 3.4.2

Table 4 - Known Issues

Internal Ref.	Issue
1434034	Description: RDMA-CM applications do not work when PCP is configured on one side of the connection.
	Workaround: Make sure PCP is configured on both sides of the connection.
	Keywords: RDMA-CM, PCP
	Discovered in Release: 3.4.2
1428828	Description: Extended join multicast API is not supported.
	Workaround: N/A
	Keywords: RDMA, Multicast
	Discovered in Release: 3.4.2
1402958	Description: Interfaces are not loaded after firmware software reset while RDMA traffic is running in the background.
	Workaround: Stop all RDMA applications for the interfaces to be loaded.
	Keywords: Self healing, RDMA
	Discovered in Release: 3.4.2
1313461	Description: When Packet Pacing is enabled in firmware, only one traffic class will be supported by the firmware.
	Workaround: Disable Packet Pacing in the firmware configuration. For example: <pre># cat /tmp/disable_pp.txt MLNX_RAW_TLV_FILE 0x00000004 0x0000010c 0x00000000 0x00000000 # mlxconfig -d pci0:4:0:0 -f /tmp/disable_pp.txt set_raw</pre>
	Keywords: Firmware, Packet Pacing
	Discovered in Release: 3.4.2
1227471	Description: When loading and unloading linuxkpi module, the following error message will appear in the dmesg, indicating that a memory leak has occurred: <pre>"Warning: memory type linux leaked memory on destroy (2 allocations, 64 bytes leaked). Warning: memory type linuxcurrent leaked memory on destroy (7 allocations, 896 bytes leaked)."</pre>
	Workaround: N/A
	Keywords: linuxkpi
	Discovered in Release: 3.4.1

Table 4 - Known Issues

Internal Ref.	Issue
1243940	<p>Description: RDMA applications (user space and kernel space) might hang when restarting the driver during traffic.</p> <p>Workaround: Send a termination signal for the application or make sure to stop the application before restarting the driver.</p> <p>Keywords: RDMA, driver restart</p> <p>Discovered in Release: 3.4.1</p>
-	<p>Description: The following error message may be printed to dmesg when using static configuration via rc.conf: "loopback_route: deletion failed" This is a kernel-related issue.</p> <p>Workaround: N/A</p> <p>Keywords: Static Configuration</p>
-	<p>Description: Choosing a wrong interface media type will cause a “no carrier” status and the physical port will not be active.</p> <p>Workaround: N/A</p> <p>Keywords: Media Type</p>
-	<p>Description: There is no TCP traffic when configuring MTU in the range of 72-100 bytes in ConnectX®-4 Lx.</p> <p>Workaround: N/A</p> <p>Keywords: MTU</p>

4 Bug Fixes

The table below lists the bugs fixed in this release.

Table 5 - Bug Fixes

Internal Ref.	Issue
1498467	Description: Added support for 10G-ER and 10G-LR modules recognition.
	Keywords: SFP module
	Discovered in Release: 3.4.2
	Fixed in Release: 3.5.0
1175757	Description: Added support for running RDMA CM with IPoIB.
	Keywords: RDMA CM, IPoIB
	Discovered in Release: 3.4.1
	Fixed in Release: 3.5.0
1337448/ 1485155/ 1470374	Description: Fixed the issue of when rebooting a virtual machine (VM), the following log message may appear: <code>warning: event(0) on port 0</code>
	Keywords: Virtualization, RDMA
	Discovered in Release: 3.4.2
	Fixed in Release: 3.5.0
1297834	Description: Fixed the issue of when running over VLAN, RDMA loopback traffic used to fail.
	Keywords: RDMA, loopback, VLAN
	Discovered in Release: 3.4.1
	Fixed in Release: 3.4.2
1258718	Description: Fixed the issue of when working in RoCE mode using ConnectX-4 HCAs only, a bandwidth performance degradation used to occur when sending/receiving a message of any size larger than 16K.
	Keywords: RoCE, performance, ConnectX-4
	Discovered in Release: 3.4.1
	Fixed in Release: 3.4.2
1273118/ 1399014	Description: Added support for RDMA multicast traffic.
	Keywords: RDMA, multicast
	Discovered in Release: 3.4.1
	Fixed in Release: 3.4.2

Table 5 - Bug Fixes

Internal Ref.	Issue
765775	Description: Suppressed EEPROM error message/s that used to be received when SFP cages were empty.
	Keywords: EEPROM, SFP
	Discovered in Release: 3.0.0
	Fixed in Release: 3.3.0
854565	Description: Allowed setting software MTU size below the value of 1500.
	Keywords: MTU
	Discovered in Release: 3.0.0
	Fixed in Release: 3.3.0

5 Change Log History

Table 6 - Change Log History

Release	Category	Description
3.4.2	RoCE Packet Sniffing	Added support for RoCE packets sniffing using tcpdump tool.
	VLAN 0 Priority Tagging	Added support for 802.1Q Ethernet frames to be transmitted with VLAN ID set to zero in RoCE mode.
	Differentiated Service Code Point (DSCP)	Added support for classifying and managing network traffic and providing quality of service (QoS) on IP and RoCE networks.
	Trust State	Added support for prioritizing sent/received packets based on packet fields.
	Reset Flow	Added support for a reset mechanism to recover from fatal failures. Upon such failures, a firmware dump for all relevant registers will be triggered, followed by a firmware and driver reset.
	RDMA Mutlicast Support	Added support for sending and receiving RDMA multicast packets.
3.4.1	Explicit Congestion Notification (ECN)	Added support for ECN, which enables end-to-end congestion notifications between two end-points when a congestion occurs.
	Rate Limiting	Added support for users to rate limit a specific Traffic Class.
	Priority Flow Control (PFC)	Added the ability to apply pause functionality to specific classes of traffic on the Ethernet link. Note: Currently, only layer 2 PFC (PCP) is supported.
	Rx Hardware Time-Stamping	Added support for adding high-quality hardware time-stamping on incoming packets.
	Firmware Dump	Added the ability to dump hardware registered data upon demand.
3.3.0	Packet Pacing	Also known as “rate limit”, this feature is now supported at a GA level. Note: This feature is supported in firmware v12.17.1016 and above.
3.0.0	Hardware LRO	Added support for Large Receive Offload (LRO) in the hardware. It increases inbound throughput of high-bandwidth network connections by reducing CPU overhead. Hardware LRO is only supported in ConnectX®-4.
	Completion Based Moderation	Added the option to reset the timer for generating interrupts upon completion generation.
	EEPROM Cable Reading	Added support for EEPROM cable reading via ifconfig and sysctl. EEPROM is only supported in ConnectX®-4.
	Interface Name	Changed the interface name from mlx5en<X> to mce<X>.