



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

Rev 3.4.1



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) 970-3400
Fax: (408) 970-3403

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

Mellanox®, Mellanox logo, Accelio®, BridgeX®, CloudX logo, CompustorX®, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, EZchip®, EZchip logo, EZappliance®, EZdesign®, EZdriver®, EZsystem®, GPUDirect®, InfiniHost®, InfiniBridge®, InfiniScale®, Kotura®, Kotura logo, Mellanox CloudRack®, Mellanox CloudXMellanox®, Mellanox Federal Systems®, Mellanox HostDirect®, Mellanox Multi-Host®, Mellanox Open Ethernet®, Mellanox OpenCloud®, Mellanox OpenCloud Logo®, Mellanox PeerDirect®, Mellanox ScalableHPC®, Mellanox StorageX®, Mellanox TuneX®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, NPS®, Open Ethernet logo, PhyX®, PlatformX®, PSIPHY®, SiPhy®, StoreX®, SwitchX®, Tiler®, Tiler logo, TestX®, TuneX®, The Generation of Open Ethernet logo, UFM®, Unbreakable Link®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

For the most updated list of Mellanox trademarks, visit <http://www.mellanox.com/page/trademarks>

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 3.4.1.	7
Chapter 3 Known Issues	8
Chapter 4 Bug Fixes.	10
Chapter 5 History of Changes and New Features	11



List of Tables

Table 1:	Release Update History	5
Table 2:	Supported Platforms and Operating Systems	6
Table 3:	Changes and Fixes in Rev 3.4.1	7
Table 4:	Known Issues	8
Table 5:	Bug Fixes	10
Table 6:	History of Changes and New Features	11



Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 3.4.1	March 1, 2018	Initial release of this version

1 Introduction

These are the release notes for Mellanox Technologies' driver for FreeBSD Rev 3.4.1 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 (at beta level): 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 (at beta level): SDR, QDR, FDR, FDR10, EDR Ethernet: 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, and 100GigE

- a. 56 GbE is a Mellanox proprietary 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 Rev 3.4.1:

Table 2 - Supported Platforms and Operating Systems

Operating System	Platform
FreeBSD 12 Current (rev 329338)	AMD64/x86_64

1.2 Supported Adapters Firmware Versions

Mellanox OFED for FreeBSD Rev 3.4.1 supports the following Mellanox network adapter cards:

Supported Adapters	Current Firmware Rev.
ConnectX®-4	12.22.1002
ConnectX®-4 Lx	14.22.1002
ConnectX-5/ConnectX-5 Ex	16.22.1002

2 Changes and New Features in Rev 3.4.1

Table 3 - Changes and Fixes in Rev 3.4.1

Category	Description
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.

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
1175757	Description: Running RDMA CM with IPoIB is not supported.
	Workaround: N/A
	Keywords: RDMA CM, IPoIB
	Discovered in Release: 3.4.1
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
1243940	Description: RDMA applications might hang when restarting the driver during traffic.
	Workaround: Send a termination signal for the application or make sure to stop the application before restarting the driver.
	Keywords: RDMA, driver restart
	Discovered in Release: 3.4.1
1258718	Description: When working in RoCE mode using ConnectX-4 HCAs only, a bandwidth performance degradation occurs when sending/receiving a message of any size larger than 16K.
	Workaround: N/A
	Keywords: RoCE, performance, ConnectX-4
	Discovered in Release: 3.4.1
1273118	Description: RDMA multicast traffic is not supported.
	Workaround: N/A
	Keywords: RDMA, multicast
	Discovered in Release: 3.4.1

Table 4 - Known Issues

Internal Ref.	Issue
1297834	Description: When running over VLAN, RDMA loopback traffic fails.
	Workaround: N/A
	Keywords: RDMA, loopback, VLAN
	Discovered in Release: 3.4.1
-	Description: Kernel panic may occur when creating/destroying a VLAN in a loop during traffic.
	Workaround: N/A
	Keywords: VLAN
-	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.
	Workaround: N/A
	Keywords: Static Configuration
-	Description: LAG creation and destruction may sometimes cause both interfaces to have the same MAC address. This is a kernel-related issue.
	Workaround: N/A
	Keywords: LAG
-	Description: Rapidly restarting the interface using netif start/stop may cause an infinite number of messages printing to dmesg: "mce0: a looped back NS message is detected during DAD for fe80:6::e61d:2dff:feal:e974. Another DAD probes are being sent"
	Workaround: N/A
	Keywords: netif restart
-	Description: Choosing a wrong interface media type will cause a "no carrier" status and the physical port will not be active.
	Workaround: N/A
	Keywords: Media Type
-	Description: There is no TCP traffic when configuring MTU in the range of 72-100 bytes in ConnectX®-4 Lx.
	Workaround: N/A
	Keywords: MTU

4 Bug Fixes

The table below lists the bugs fixed in this release.

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 History of Changes and New Features

Table 6 - History of Changes and New Features

Release	Category	Description
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>.