



Connect. Accelerate. Outperform.™

Mellanox DPDK Release Notes

Rev 2.2_4.2

www.mellanox.com

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 2016. Mellanox Technologies LTD. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, CloudX logo, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, EZchip®, EZchip logo, EZappliance®, EZdesign®, EZdriver®, EZsystem®, GPUDirect®, InfiniHost®, InfiniScale®, Kotura®, Kotura logo, Mellanox Federal Systems®, Mellanox Open Ethernet®, Mellanox ScalableHPC®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, Open Ethernet logo, PhyX®, SwitchX®, Tiler®, Tiler logo, TestX®, The Generation of Open Ethernet logo, UFM®, 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>

Contents

Document Revision History	5
1 Overview	6
1.1 System Requirements	6
2 Changes and Major New Features in Rev 2.2_4.2	7
3 mlx4 and mlx5 PMD Drivers Features	8
4 Known Issues	9
4.1 General Known Issues	9
4.2 mlx4 PMD Known Issues.....	9
4.3 mlx5 PMD Known Issue	10
5 Bug Fixes	14
5.1 ConnectX-4 PMD Bug Fixes.....	14
5.2 ConnectX-3 PMD Bug Fixes.....	16
5.3 DPDK v2.2 Bug Fixes	17
6 Change Log History	19

List of Tables

Table 1: Document Revision History	5
Table 2: System Requirements	6
Table 3: Changes and Major New Features.....	7
Table 4: General Known Issues	9
Table 5: mlx4 PMD Known Issues	9
Table 6: ConnectX-4 PMD Bug Fixes	14
Table 7: DPDK v2.2 Bug Fixes.....	17
Table 8: Change Log History.....	19

Document Revision History

Table 1: Document Revision History

Release	Date	Description
Rev 2.2_4.2	October 09, 2016	Initial release of this DPDK version.

1 Overview

These are the release notes for mlx4 and mlx5 DPDK Poll-Mode Driver (PMD) for Mellanox ConnectX®-3/ConnectX®-3 Pro and ConnectX®-4/ ConnectX®-4 Lx Ethernet adapters.

1.1 System Requirements

Table 2: System Requirements

Specification	Value
Network Adapter Cards	ConnectX®-3 / ConnectX®-3 Pro / ConnectX®-4 / ConnectX®-4 Lx network adapter card. (These must be configured to work in ETH mode.)
Firmware	<ul style="list-style-type: none"> • ConnectX-4: 12.17.1010 • ConnectX-4 Lx: 14.17.1010 • ConnectX 3/ ConnectX-3 Pro: 2.36.5000 and above
Linux Driver Stack	MLNX_OFED- 3.4-1.0.0.0 / MLNX_EN-3.4-1.0.0.0
ESXi 6.0 Driver	ConnectX-4 / ConnectX-4 Lx: 4.15.6-22
ESXi 5.5 U1 Driver	ConnectX-4 / ConnectX-4 Lx: 4.5.6-22
Operating Systems and Kernels	All OSs supported by MLNX_OFED
Minimum memory requirements	16 GB RAM
Transport	Ethernet
CPU Arch	x86, Power8

2 Changes and Major New Features in Rev 2.2_4.2

Table 3: Changes and Major New Features

Driver	Changes
ConnectX-4 PMD, mlx5	<ul style="list-style-type: none">• Added support for MLNX_OFED 3.4-x.x.x.x• Added support for MLNX_EN 3.4-x.x.x.x• Vectorized logic on TX path to improve CPU utilization• Bug fixes for mlx5

3 mlx4 and mlx5 PMD Drivers Features

Feature	mlx4 PMD	mlx5 PMD
Supported NICs	ConnectX®-3 ConnectX®-3 Pro	ConnectX®-4 ConnectX®-4 Lx
PCI mapping	Function per device	Function per port
KVM SR-IOV	Yes	Yes
ESX 6.0 SR-IOV	Yes	Yes
ESX 5.5 SR-IOV	No	Yes
Scattering/gathering RX/TX packets	Yes	Yes
Multiple RX (with RSS/RCA) and TX queues	Yes	Yes
IPv4, TCP IPv4, UDP IPv4 RSS	Yes	Yes
IPv6 RSS	Yes	Yes
VXLAN RSS	Yes	According to the outer header
Number of RSS queues	Power of 2	Any
Get and Set RSS key per flow type (rss_hf)	No	Yes
Multiple MAC addresses	Yes	Yes
VLAN filtering	Yes	Yes
Link state information	Yes	Yes
Software counters/statistics	Yes	Yes
Start/stop/close operations	Yes	Yes
Multiple physical ports host adapter	Yes	Yes
Promiscuous mode	Yes	Yes
Multicast Promiscuous	Yes	Yes
Checksums hardware offloading	Yes	Yes
Checksum VXLAN hardware offloading	Yes	No
Flow Director	No	Yes
RX VLAN stripping	No	Yes
TX VLAN insertion	No	Yes

4 Known Issues

4.1 General Known Issues

Table 4: General Known Issues

Subject	Description	Workaround
InfiniBand	Mellanox PMDs does not support InfiniBand.	N/A
Support for DPDK integrated shared library	Mellanox PMDs does not support DPDK integrated shared library	PMD compiled as shared library can be used
Hardware counters	Hardware counters are not implemented.	N/A
The Primary and Secondary multi process model	The Primary and Secondary multi process model is currently supported in TX only.	Use multithreaded model instead of multi process in case mutli process RX is needed
Active - Passive Bonding mode	Bond PMD does not configure the needed MAC address to Bond slaves in case of failover. Active-Passive Bonding mode is available in case of KVM SR-IOV when 2 VMs are configured with the same MAC.	
Bond LAG interface is down	Bond PMD does not notice when Mellanox PMD's ports are up	Run ifconfig down and up on the bond slaves of the Mellanox interfaces

4.2 mlx4 PMD Known Issues

Table 5: mlx4 PMD Known Issues

Subject	Description	Workaround
Multicast / Broadcast Self-loopback on SR-IOV VF	When a Multicast or Broadcast packet is sent to the SR-IOV VF, it is returned back to the port.	N/A
Performance degradation when SGE_NUM =4	Performance degradation might occur for small packets when PMD is compiled with SGE_NUM = 4 compared to the performance when SGE_NUM = 1	If scattered packets are not used, compile PMD with SGE_NUM = 1
Hardware checksum offloading on ConnectX-3	Hardware checksum offloading is not supported in ConnectX-3	Use ConnectX3-Pro if hardware checksum offloading is needed
Sending abnormal packets	Sending abnormal packets, smaller then 16B and bigger than the configured MTU causes PMD queues to enter an error state.	run dev_stop and dev_start

Subject	Description	Workaround
ESX 6.0 and DMFS A0 mode	DMFS A0 mode is not supported in ESX 6.0 VM	N/A
ESX 6.0 and HW checksum offloading	RX hardware checksum offloads are not supported in ESX 6.0 VM	N/A
ESX 6.0 and “add MAC”	Only one additional MAC can be added in additional to the VF’s MAC in ESX 6.0 VM	N/A
VXLAN hardware checksum offloading	VXLAN hardware checksum offloading is supported only if the used kernel supports it	N/A
L4 checksum report with hardware checksum offloading	When the packet does not include TCP/UDP header and the hardware checksum offloading is enabled, L4 checksum will be reported as bad	N/A
RSS	RSS hash key and options cannot be modified.	N/A
Promiscuous and SR-IOV	Promiscuous mode does not work when SR-IOV is enabled.	N/A
Promiscuous and SR-IOV	In testpmd: Although promiscuous mode fails in SR-IOV, it still shows as enabled This is DPDK implementation bug.	N/A
VLAN filtering	VLAN filtering is supported only with non-optimized steering mode.	See QSG, RX VLAN Filter section
Number of configured RSS queues must be power of 2	Number of configured RSS queues must be power of 2.	Use only power of 2 RSS queues
Broadcast packets and VLAN filter on KVM VM	Broadcast packets with VLAN are not received when VLAN filter is configured on KVM VM	N/A
RSS and VXLAN traffic	VXLAN traffic is not being spread among the RSS queues	N/A
MLNX OFED 2.4	MLNX_OFED 2.4 and below is not supported	It is recommended to use MLNX_OFED 3.3

4.3 mlx5 PMD Known Issue

Table 4: mlx5 PMD Known Issues

Subject	Description	Workaround
OVS support	Using OVS 2.5.0 requires several additional patches for it to function properly	Use OVS 2.6.0 branch instead. For further information, refer to the Mellanox Community.

Subject	Description	Workaround
UDP port 4791	Packets that have UDP destination port 4791 will be silently dropped by the NIC which may cause to low performance of 0% packet loss tests. This port is a reserved port and should not be used.	Use a different port as the UDP destination port
Auto inline	Auto inline feature is supported only on a single port ConnectX-4 adapter card (is not supported on ConnectX-4 Lx) and cannot be used for 0 packet loss performance	Use auto-inline only with supported mode
LCS events	When using default firmware configuration, LCS events are not received when <code>rte_eth_dev_set_link_down/up</code> are called.	Run <code>mlxconfig -d /dev/mst/mt4115_pciconf0 set KEEP_ETH_LINK_UP_P1=0 KEEP_ETH_LINK_UP_P2=0</code> Reboot the server.
LCS events on VM	LCS events are not received on KVM/ESX VM when <code>rte_eth_dev_set_link_down/up</code> are called.	N/A
VXLAN HW Checksum offloading	VXLAN HW Checksum offloading is not supported	Will be added in next releases
VXLAN RSS	VXLAN traffic is spread between RSS queues according to the outer header	Will be added in next releases
Flow director VXLAN filter	Flow director VXLAN filter is not supported	Will be added in next releases
The dest MAC of the sent packet	When destination MAC of a sent packet is the same as the port's MAC, the packet is returned to the port and is not sent to the network. Note: On ConnectX-4 Lx, the packet is sent when <code>txq_mpw_en</code> is set to 1 (the default configuration).	N/A
RSS with non-power of 2 RX queues performance	When the number of RX queues is not power of 2, the RETA table size is automatically configured to 256 to achieve the best flows spreading. Still in some cases traffic spreading between the queues can be non-equal	Use power of 2 number of queues if possible
Maximum supported size of TX/RX queue	The maximum supported number of descriptors in TX/RX queues is equal to 32K. When <code>txq_inline</code> parameter is used the	Use maximum number of TX and RX descriptors as explained.

Subject	Description	Workaround
	<p>number is smaller and is equal to $32K/(txq_inline/32 + 1)$.</p> <p>When auto-inline is used, the maximum supported number of descriptor is 6K</p> <p>If the number of descriptors of TX or RX queues is higher than the above maximum, segmentation fault will occur</p>	
Performance of small messages when MTU > mbuf size	Performance degradation of small messages might occur when MTU is higher than the mbuf size	N/A
Broadcast packets on KVM VM	Broadcast packets are not being received on KVM VM	N/A
Adding MAC address on KVM VM	Adding MAC address on KVM VM is not supported	N/A
VLAN packets on KVM VM	VLAN packets are received but are not sent when DPDK is running on KVM VM	N/A
RX VLAN Stripping has unexpected behavior	<p>DPDK has a flag per port to indicate the VLAN stripping state (On/Off), that updated when using <code>rte_eth_dev_set_vlan_offload()</code> API.</p> <p>This flag is not updated when calling the <code>rte_eth_dev_set_vlan_strip_on_queue()</code> API (VLAN stripping per queue), and can cause some confusion when mixing between those two APIs.</p>	Use only one API to configure RX stripping
TX VLAN insertion offloading with ConnectX-4 LX	TX VLAN offloading insertion is supported only when <code>txq_mpw_en</code> is set to 0. By default is not supported since MPW is enabled	Set <code>txq_mpw_en=0</code> . Please check the Quick Start Guide for more information regarding command line options.
Changing MTU during traffic	Changing MTU during traffic results in segmentation fault.	Stop the ports before changing the MTU
Flow Control settings on KVM VM	Flow Control settings are not allowed on VM	Set Flow Control settings on the Hypervisor
Allmulticast on KVM and ESX VM	Allmulticast on VM is not supported	N/A
Adding unicast MAC on ESX VM is not supported	ESX VM supports only a single unicast MAC	N/A

Subject	Description	Workaround
MTU configuration on KVM VM	When configuring the MTU on the VM make sure the MTU on the hypervisor is the same as on the VM	N/A
Flow director on ESX VM	Flow director is not supported on ESX VM	N/A

5 Bug Fixes

5.1 ConnectX-4 PMD Bug Fixes

Table 6: ConnectX-4 PMD Bug Fixes

Internal Ref.	Issue
-	<p>Description: Fixed an issue related to receiving packets smaller than the MTU size and bigger than the mbuf size</p> <p>Key Words: Small mbuf size</p> <p>Discovered in Release: MLNX_DPDK 2.2_2.7</p> <p>Fixed in Release: MLNX_DPDK 2.2_4.2</p>
-	<p>Description: Fixed data corruption in case of mbuf allocation failure</p> <p>Key Words: Data corruption</p> <p>Discovered in Release: MLNX_DPDK 2.2_2.7</p> <p>Fixed in Release: MLNX_DPDK 2.2_3.9</p>
-	<p>Description: Fixed link status report when LSC events were configured</p> <p>Key Words: LCS events</p> <p>Discovered in Release: MLNX_DPDK 2.2_2.7</p> <p>Fixed in Release: MLNX_DPDK 2.2_3.9</p>
-	<p>Description: Fixed RX VLAN offload capability report. This capability was implemented but not reported</p> <p>Key Words: VLAN offload capability report</p> <p>Discovered in Release: MLNX_DPDK 2.2_2.7</p> <p>Fixed in Release: MLNX_DPDK 2.2_3.9</p>
-	<p>Description: Modified the function priv_set_flags (priv, keep, flags) behavior. Now the "keep" mask is no longer ignored.</p> <p>Key Words: Setting interface flags</p> <p>Discovered in Release: MLNX_DPDK 2.2_1.6</p> <p>Fixed in Release: MLNX_DPDK 2.2_3.9</p>
-	<p>Description: Fixed an issue that caused VLAN filters removal when two VLAN filters were configured and one was removed. In this case, the second was also removed by mistake.</p> <p>Key Words: VLAN filter</p> <p>Discovered in Release: MLNX_DPDK 2.2_1.6</p> <p>Fixed in Release: MLNX_DPDK 2.2_3.9</p>
-	<p>Description: Fixed inconsistent return value in Flow Director</p> <p>Key Words: Flow director return error code</p> <p>Discovered in Release: MLNX_DPDK 2.2_1.6</p>

Internal Ref.	Issue
	Fixed in Release: MLNX_DPDK 2.2_2.7
-	<p>Description: By default, testpmd configures CRC to be scattered by the HW, but it does not strip the CRC header. ConnectX-4 adds CRC header for every packet that is sent. In case of IO forwarding on a loopback setup, the forwarded packet increases each send by 4B</p> <p>Key Words: testpmd application and CRC stripping</p> <p>Discovered in Release: MLNX_DPDK 2.2._1.6</p> <p>Fixed in Release: MLNX_DPDK 2.2_2.7</p>
-	<p>Description: Fixed an issue that resulted in the following error when working with MTU > 2048 and the PMD is compiled with SGE_NUM=4: mlx5: got completion with error:</p> <p>Key Words: Jumbo frames</p> <p>Discovered in Release: MLNX_DPDK 2.2._1.4</p> <p>Fixed in Release: MLNX_DPDK 2.2_1.6</p>
-	<p>Description: Avoid packet duplication with overlapping Flow Director rules. For example the case where Flow Director mask is zero and the value is not</p> <p>Key Words: Flow Director: zero values</p> <p>Discovered in Release: MLNX_DPDK 2.2._1.4</p> <p>Fixed in Release: MLNX_DPDK 2.2_1.6</p>
-	<p>Description: Fixed an error which occurred when the software configured FCS stripping or HW padding which were not supported by the firmware.</p> <p>Key Words: HW capabilities check</p> <p>Discovered in Release: MLNX_DPDK 2.2._1.4</p> <p>Fixed in Release: MLNX_DPDK 2.2_1.6</p>
-	<p>Description: Fixed a faulty L4 checksum which was reported when the packet is L3</p> <p>Key Words: L4 checksum of L3 packet</p> <p>Discovered in Release: MLNX_DPDK 2.1._1</p> <p>Fixed in Release: MLNX_DPDK 2.2_1</p>
-	<p>Description: Fixed a faulty L3 and L4 checksums which were reported as bad in case packet is L2</p> <p>Key Words: L3 and L4 checksum of L2 packet</p> <p>Discovered in Release: MLNX_DPDK 2.1._1</p> <p>Fixed in Release: MLNX_DPDK 2.2_1</p>
-	<p>Description: Fixed an issue that showed some features as not supported although they were when DPDK was compiled with -j option</p> <p>Key Words: Parallel builds compilation</p>

Internal Ref.	Issue
	Discovered in Release: MLNX_DPDK 2.1._1
	Fixed in Release: MLNX_DPDK 2.2_1
-	Description: Fixed an issue causing IPv6 resolution protocol to not work in case all-multicast options were not configured
	Key Words: IPv6 resolution protocol
	Discovered in Release: MLNX_DPDK 2.1._1
	Fixed in Release: MLNX_DPDK 2.2_1

5.2 ConnectX-3 PMD Bug Fixes

Table 6: ConnectX-3 PMD Bug Fixes

Internal Ref.	Issue
-	Description: Modified the function <code>priv_set_flags</code> (<code>priv</code> , <code>keep</code> , <code>flags</code>) behavior. Now the "keep" mask is no longer ignored.
	Key Words: Setting interface flags
	Discovered in Release: MLNX_DPDK 2.2_1.6
	Fixed in Release: MLNX_DPDK 2.2_3.9
-	Description: Fixed compilation issues on RedHat 6.7.
	Key Words: Compilation issues
	Discovered in Release: MLNX_DPDK 2.2_2.7
	Fixed in Release: MLNX_DPDK 2.2_3.9
-	Description: Fixed compilation issues with glibc 2.20 and ICC.
	Key Words: Compilation issues
	Discovered in Release: MLNX_DPDK 2.2_2.7
	Fixed in Release: MLNX_DPDK 2.2_3.9
-	Description: Fixed DPDK on Windriver
	Key Words: Windriver OS support
	Discovered in Release: MLNX_OFED 3.2-x.x.x.x
	Fixed in Release: MLNX_OFED 3.3.x.x.x.x
-	Description: Fixed TX checksum offloading on KVM VM.
	Key Words: TX checksum offloading on KVM
	Discovered in Release: MLNX_OFED 3.2-x.x.x.x
	Fixed in Release: MLNX_OFED 3.3.x.x.x.x
-	Description: Fixed Virtual Machines (VM) with enabled Multicast Promiscuous mode do not receive the multicast traffic that other VMs on the same hypervisor are registered to.
	Key Words: Multicast Promiscuous mode and SR-IOV

Internal Ref.	Issue
	Discovered in Release: MLNX_OFED 2.4-x.x.x.x
	Fixed in Release: MLNX_OFED 3.0-x.x.x.x and above
-	Description: Fixed QP setup failure on Debian 3.16.7-ckt7-1
	Key Words: QP setup failure
	Discovered in Release: DPDK 2.1 inbox
	Fixed in Release: MLNX_DPDK 2.2_1.1
-	Description: Fixed an unpredicted behavior due to “cow” mechanism when an application used fork() or system call that were not related to the networking.
	Key Words: Memory corruption when fork() is used in the application
	Discovered in Release: DPDK 2.0 inbox
	Fixed in Release: MLNX_DPDK 2.0_2.8.4
-	Description: Fixed a wrong speed reporting when a 10Gb/s port was set as 1Gb/s (10Gb/s instead of 1Gb/s)
	Key Words: 1Gb/s port’s link reported as 10Gb/s
	Discovered in Release: DPDK 2.0 inbox
	Fixed in Release: MLNX_DPDK 2.0_2.8.4
-	Description: Fixed a multiple RX VLAN filters issue preventing multiple RX VLAN filters to function properly although all were configured. In this case, only the first one functioned properly.
	Key Words: Multiple RX VLAN filters
	Discovered in Release: DPDK 2.0 inbox
	Fixed in Release: MLNX_DPDK 2.0_2.8.4

5.3 DPDK v2.2 Bug Fixes

Table 7: DPDK v2.2 Bug Fixes

Internal Ref.	Issue
-	Description: Fixed an inconsistency issue of byte order in ethdev library when setting fdir's mask and flow entry causing the inputs of mask to be in big endian format
	Key Words: Byte order consistency of flow director
	Discovered in Release: DPDK 2.2
	Fixed in Release: MLNX_DPDK 2.2_3.9
-	Description: Fixed compilation issues on RedHat 6.7.
	Key Words: Compilation issues
	Discovered in Release: DPDK 2.2
	Fixed in Release: MLNX_DPDK 2.2_3.9

Internal Ref.	Issue
-	Description: Multiple changes in all DPDK libs to support latest gcc compilers when pedantic flag is enabled.
	Key Words: Power8 compilation with DEBUG mode enabled
	Discovered in Release: DPDK 2.2
	Fixed in Release: MLNX_DPDK 2.2_3.9
-	Description: Added support for 100 Gb/s link rate
	Key Words: ConnectX-4 100G/b port link speed
	Discovered in Release: DPDK 2.2
	Fixed in Release: MLNX_DPDK 2.2_1.6
-	Description: Fixed compilation failure with <code>-pedantic</code> on Power8
	Key Words: Compilation on Power8
	Discovered in Release: DPDK 2.2 inbox
	Fixed in Release: MLNX_DPDK 2.2_1.4
-	Description: Add support for 512 entries in RETA table
	Key Words: RETA table
	Discovered in Release: DPDK 2.2 inbox
	Fixed in Release: MLNX_DPDK 2.2_1.4

6 Change Log History

Table 8: Change Log History



NOTE: The changes are from the previous DPDK version on dpdk.org and not between two MLNX_DPDK versions. DPDK (dpdk.org) Release Notes can be found at: http://dpdk.org/doc/guides/rel_notes/

Driver	Changes
Rev 2.2_3.9 from Rev 2.2_2.7	
ConnectX-3 PMD, mlx4	<ul style="list-style-type: none"> Added link up/down callback functions support
DPDK v2.2	<ul style="list-style-type: none"> Backport changes/fixes from DPDK 16.07 mbuf: add prefetch helpers: Enables the drivers to prefetch the RX or TX part of the mbuf, regardless of the cache line size. Note: Some architectures (ex: Power8) have a cache line size of 128 bytes, so the drivers should not expect that prefetching the second part of the mbuf with <code>rte_prefetch0 (&m->cacheline1)</code> to be valid. mbuf: add raw allocation function: Added a new new public function in <code>rte_mbuf</code> to allocate a raw mbuf (uninitialized)
Rev 2.2_2.7 from Rev 2.2_1.6	
ConnectX-4 PMD, mlx5	<ul style="list-style-type: none"> Major CPU usage performance optimizations Performance improvements for messages $\geq 256B$ Automatic scatter gather setting: RX according to MTU and TX is set to HW maximum Added support for flow director drop queue Replaced compilation options with command line parameters Added support for ESX 6.0 and 5.5 SR-IOV for ConnectX-4 HW TX VLAN insertion only (removed SW VLAN insertion) Bug fixes
Rev 2.2_1.6 from Rev 2.2_1.4	
ConnectX-4 PMD, mlx5	<ul style="list-style-type: none"> DPDK bug fixes
ConnectX-3 PMD, mlx4	<ul style="list-style-type: none"> Bug fixes
Rev 2.2_1.4 from DPDK 2.2	
ConnectX-4 PMD, mlx5	<ul style="list-style-type: none"> Added support for flow director filters (<code>RTE_FDIR_MODE_PERFECT</code> and <code>RTE_FDIR_MODE_PERFECT_MAC_VLAN</code> modes) Single core and max performance optimizations Added support for RX VLAN stripping

Driver	Changes
	<ul style="list-style-type: none"> • Added support for TX VLAN insertion (supported only with MLNX_OFED-3.2-1.5.0.1 and above) • Added callbacks to support link (up / down) changes • Added support for HW packet padding (supported only with MLNX_OFED-3.2-1.5.0.1 and above) • Added support for scattering FCS (supported only with MLNX_OFED-3.2-1.5.0.1 and above) • Added TX support for secondary processes • Added support for MLNX_OFED 3.2-1.0.1.1 and higher • DPDK fixes (see section Table 7: DPDK v2.2 Bug Fixes)
ConnectX-4 PMD, mlx4	<ul style="list-style-type: none"> • Added support for VMware ESX 6.0 U1
Rev 2.1_1.1 from DPDK 2.1 inbox	
ConnectX-4 PMD, mlx5	<ul style="list-style-type: none"> • Added support for ConnectX@-4 and ConnectX@-4 LX NICs • Added support for MLNX_OFED 3.1-1.0.0 • DPDK fixes (see section Table 7: DPDK v2.2 Bug Fixes)
Rev 2.1_1.1 from MLNX DPDK 2.0_2.8.4	
ConnectX-3 PMD, mlx4	<ul style="list-style-type: none"> • Removed support for VMware ESX 5.5 • Added support for Accelerated verbs – PMD does not include libibverbs and libmlx4