



Mellanox ConnectX[®]-4 Lx Firmware Release Notes

Rev 14.23.1020

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®, LinkX®, 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®, Tilera®, Tilera 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

Release Update History	5
Chapter 1 Overview	6
1.1 Supported Devices	6
1.2 Supported Cables and Modules	7
1.2.1 Validated and Supported 1GbE Cables	7
1.2.2 Validated and Supported 10GbE Cables	7
1.2.3 Validated and Supported 25GbE Cables	9
1.2.4 Validated and Supported 40GbE Cables	10
1.2.5 Validated and Supported 50GbE Cables	11
1.2.6 Validated and Supported 100GbE Cables	11
1.2.7 Validated and Supported FDR Cables	12
1.3 Tested Switches	13
1.3.1 Tested 10GbE Switches	13
1.3.2 Tested 40GbE Switches	14
1.3.3 Tested 100GbE Switches	14
1.4 Tools, Switch Firmware and Driver Software	14
1.5 Supported FlexBoot, UEFI	15
1.6 Revision Compatibility	15
Chapter 2 Changes and New Features in Rev 14.23.1020	16
Chapter 3 Known Issues	17
Chapter 4 Bug Fixes History	21
Chapter 5 Firmware Changes and New Feature History	29
Chapter 6 FlexBoot Changes and New Features	38
6.1 FlexBoot Known Issues	41
6.2 FlexBoot Bug Fixes History	46
Chapter 7 UEFI Changes and Major New Features	48
7.1 UEFI Known Issues	48
Chapter 8 Unsupported Features and Commands	49
8.1 Unsupported Features	49
8.2 Unsupported Commands	49
Chapter 9 Supported Non-Volatile Configurations	50

List of Tables

Table 1:	Release Update History	5
Table 2:	Supported Devices	6
Table 3:	Validated and Supported 1GbE Cables	7
Table 4:	Validated and Supported 10GbE Cables	7
Table 5:	Validated and Supported 25GbE Cables	9
Table 6:	Validated and Supported 40GbE Cables	10
Table 7:	Validated and Supported 50GbE Cables	11
Table 8:	Validated and Supported 100GbE Cables	11
Table 9:	Validated and Supported FDR Cables	12
Table 10:	Tested 10GbE Switches	13
Table 11:	Tested 40GbE Switches	14
Table 12:	Tested 100GbE Switches	14
Table 13:	Tools, Switch Firmware and Driver Software	14
Table 14:	Supported FlexBoot, UEFI	15
Table 15:	Changes and New Features in Rev 14.23.1020	16
Table 16:	Ethernet Rate Limit per VF in RoCE Mode Limitations	17
Table 17:	Ethernet Rate Limit per VF in InfiniBand Mode Limitations	17
Table 18:	Known Issues	17
Table 19:	Bug Fixes History	21
Table 20:	Firmware Changes and New Feature History	29
Table 21:	FlexBoot Changes and New Features	38
Table 22:	FlexBoot Known Issues	41
Table 23:	FlexBoot Bug Fixes History	46
Table 24:	UEFI Changes and New Features	48
Table 25:	UEFI Known Issues	48
Table 26:	Unsupported Features	49
Table 27:	Per-physical Port Settings	50
Table 28:	Global Settings	50
Table 29:	Per host/function Settings	50
Table 30:	Per host Settings	50

Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 14.23.1020	August 19, 2018	Added the following tables: - Table 16, “Ethernet Rate Limit per VF in RoCE Mode Limitations,” on page 17 - Table 17, “Ethernet Rate Limit per VF in InfiniBand Mode Limitations,” on page 17
	July 30, 2018	Initial version of this firmware release. This version introduces new Changes and Features (see Section 2, “Changes and New Features in Rev 14.23.1020” , on page 16) and Bug Fixes (see Section 4, “Bug Fixes History” , on page 21).

1 Overview

These are the release notes for the ConnectX®-4 Lx adapters firmware Rev 14.23.1020. This firmware supports the following protocols:

- Ethernet - 1GigE, 10GigE, 25GigE, 40GigE, 50GigE
- PCI Express 3.0, supporting backwards compatibility for v2.0 and v1.1

1.1 Supported Devices

This firmware supports the devices and protocols listed in [Table 2](#)

Table 2 - Supported Devices (Sheet 1 of 2)

Device Part Number	PSID	Device Name	FlexBoot	UEFI x86	UEFI ARM	Enable/disable exprom feature
MCX4111A-ACAT	MT_2410110034	ConnectX®-4 Lx EN network interface card, 25GbE single-port SFP28, PCIe3.0 x8, tall bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4111A-XCAT	MT_2410110004	ConnectX®-4 Lx EN network interface card, 10GbE single-port SFP28, PCIe3.0 x8, tall bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4121A-ACAT	MT_2420110034	ConnectX®-4 Lx EN network interface card; 25GbE dual-port SFP28; PCIe3.0 x8; ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4121A-XCAT	MT_2420110004	ConnectX®-4 Lx EN network interface card, 10GbE dual-port SFP28, PCIe3.0 x8, tall bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4131A-BCAT	MT_2430110027	ConnectX®-4 Lx EN network interface card, 40GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4131A-GCAT	MT_2430110032	ConnectX®-4 Lx EN network interface card, 50GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4411A-ACAN	MT_2450111034	ConnectX®-4 Lx EN network interface card for OCP; 25GbE single-port SFP28; PCIe3.0 x8; ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4411A-ACQN	MT_2450112034	ConnectX®-4 Lx EN network interface card for OCP with Host Management, 25GbE single-port SFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4421A-ACAN	MT_2470111034	ConnectX®-4 Lx EN network interface card for OCP; 25GbE dual-port SFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4421A-ACQN	MT_2470112034	ConnectX®-4 Lx EN network interface card for OCP with Host Management, 25GbE dual-port SFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4421A-XCQN	MT_2470110004	ConnectX®-4 Lx EN network interface card for OCP with Host Management, 10GbE dual-port SFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Enabled)	Not Present	Not Available

Table 2 - Supported Devices (Sheet 2 of 2)

Device Part Number	PSID	Device Name	FlexBoot	UEFI x86	UEFI ARM	Enable/disable exprom feature
MCX4431A-GCAN	MT_2490111032	ConnectX®-4 Lx EN network interface card for OCP, with Host Management, 50GbE single-port QSFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4431M-GCAN	MT_2510111032	ConnectX®-4 Lx EN network interface card for OCP with Multi-Host and Host Management, 50GbE single-port QSFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists

1.2 Supported Cables and Modules

Please refer to the LinkX™ Cables and Transceivers web page

(<http://www.mellanox.com/products/interconnect/cables-configurator.php>) for the list of supported cables.

1.2.1 Validated and Supported 1GbE Cables

Table 3 - Validated and Supported 1GbE Cables

Speed	Cable OPN #	Description
1GB/S	MC3208011-SX	Mellanox Optical module, SX, 850nm
1GB/S	MC3208411-T	Mellanox® module, ETH 1GbE, 1Gb/s, SFP, Base-T, up to 100m

1.2.2 Validated and Supported 10GbE Cables

Table 4 - Validated and Supported 10GbE Cables

Speed	Cable OPN #	Description
10GB/S	CAB-SFP-SFP-1M	Arista 10GBASE-CR SFP+ Cable 1 Meter
10GB/S	CAB-SFP-SFP-3M	Arista 10GBASE-CR SFP+ Cable 3 Meter
10GB/S	CAB-SFP-SFP-5M	Arista 10GBASE-CR SFP+ Cable 5 Meter
10GB/S	MC2309124-004	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 4M
10GB/S	MC2309124-005	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 5M
10GB/S	MC2309130-001	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 1M
10GB/S	MC2309130-002	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 2M
10GB/S	MC2309130-003	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 3M
10GB/S	MC2309130-00A	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 0.5M
10GB/S	MC2609125-004	Mellanox Passive Copper Hybrid Cable ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 4M

Table 4 - Validated and Supported 10GbE Cables

Speed	Cable OPN #	Description
10GB/S	MC2609125-005	Mellanox Passive Copper Hybrid Cable ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 5M
10GB/S	MC2609130-001	Mellanox Passive Copper Hybrid Cable ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 1M
10GB/S	MC2609130-002	Mellanox Passive Copper Hybrid Cable ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 2M
10GB/S	MC2609130-003	Mellanox Passive Copper Hybrid Cable ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 3M
10GB/S	MC2609130-0A1	Mellanox Passive Copper Hybrid Cable ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 1.5M
10GB/S	MC3309124-004	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 4M
10GB/S	MC3309124-005	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 5M
10GB/S	MC3309124-006	Mellanox® Passive Copper Cable, ETH 10GbE, 10Gb/s, SFP+, 6m
10GB/S	MC3309124-007	Mellanox® Passive Copper Cable, ETH 10GbE, 10Gb/s, SFP+, 7m
10GB/S	MC3309130-001	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 1M
10GB/S	MC3309130-002	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 2M
10GB/S	MC3309130-003	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 3M
10GB/S	MC3309130-00A	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 0.5M
10GB/S	MC3309130-0A1	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 1.5M
10GB/S	MC3309130-0A2	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 2.5M
10GB/S	MFM1T02A-LR-F	Mellanox Optical Module ETH 10GBE 10GB/S SFP+ LC-LC 1310NM LR up to 10KM
10GB/S	MFM1T02A-SR-F	Mellanox Optical Module ETH 10GBE 10GB/S SFP+ LC-LC 850NM SR up to 300M
10GB/S	SFP-10G-SR	Cisco 10GBASE-SR SFP+ transceiver module for MMF, 850-nm wavelength, LC duplex connector
10GB/S	SFP-H10GB-CU1M	Cisco 1-m 10G SFP+ Twinax cable assembly, passive
10GB/S	SFP-H10GB-CU3M	Cisco 3-m 10G SFP+ Twinax cable assembly, passive
10GB/S	SFP-H10GB-CU5M	Cisco 5-m 10G SFP+ Twinax cable assembly, passive
10GB/S	FTLX1471D3BCL-ME	10GBASE-LR SFP+ 1310nm 10km DOM Transceiver Module
10GE	BN-QS-SP-CBL-5M	40G QSFP+ to 4xSFP+ DAC Breakout Direct Attach Cable 5m
10GE	1-2053783-1	SFP-H10G-CU1M
10GE	BN-QS-SP-CBL-5M	40G QSFP+ to 4xSFP+ DAC Breakout Direct Attach Cable 5m
10GE	CAB-SFP-SFP-1M	Arista Compatible 10G SFP+ Passive Cable 1m
10GE	CAB-SFP-SFP-5M	Arista Compatible 10G SFP+ Passive Cable 5m
10GE	FTLX8570D3BCL-C2	Cisco FET-10G 10-2566-02 FTLX8570D3BCL-C2 10Gbps Fabric Extender SFP+ Module
10GE	L45593-D178-B50	QSFP-4SFP10G-CU5M

1.2.3 Validated and Supported 25GbE Cables



The 25GbE cables can be supported in ConnectX-4 adapter cards only when connected to the MAM1Q00A-QSA28 module.

Table 5 - Validated and Supported 25GbE Cables

Speed	Cable OPN #	Description
25GbE	FTLF8536P4BCL	Finisar SFP+ transceivers 25Gb/s
25GbE	LTF8507-PC07	Hisense active fiber cable, 25GbE
25GbE	MCP2M00-A001	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1m
25GbE	MCP2M00-A002	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2m
25GbE	MCP2M00-A003	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 3m
25GbE	MCP2M00-A003AP	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 3m, 26AWG
25GbE	MCP2M00-A005E26L	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 5m, Black, 26AWG, CA-L
25GbE	MCP2M00-A00A	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 0.5m
25GbE	MCP2M00-A01A	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1.5m
25GbE	MCP2M00-A01A	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1.5m
25GbE	MCP2M00-A02A	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2.5m
25GbE	MCP7F00-A001	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 4X25GBS QSFP28 TO 4XSFP28 1M
25GbE	MCP7F00-A002	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 4X25GBS QSFP28 TO 4XSFP28 2M
25GbE	MCP7F00-A003	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 4X25GBS QSFP28 TO 4XSFP28 3M
25GbE	MCP7F00-A003-AM	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 3M 30AWG
25GbE	MCP7F00-A005AM	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 4X25GBS QSFP28 to 4XSFP28 5M
25GbE	MCP7F00-A01A	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 4X25GBS QSFP28 to 4XSFP28 1.5M
25GbE	MCP7F00-A02A	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 4X25GBS QSFP28 to 4XSFP28 2.5M
25GbE	MFA2P10-Axxx	Mellanox® active optical cable 25GbE, SFP28, up to 100m
25GbE	MMA2L20-AR	Mellanox® optical transceiver, 25GbE, 25Gb/s, SFP28, LC-LC, 1310nm, LR up to 10km
25GbE	MMA2P00-AS	Mellanox® transceiver, 25GbE, SFP28, LC-LC, 850nm, SR, up to 100m
25GbE	SFP-H25G-CU1M	25GBASE-CR1 Copper Cable 1-meter
25GbE	SFP-H25G-CU2M	25GBASE-CR1 Copper Cable 2-meter
25GbE	SFP-H25G-CU3M	25GBASE-CR1 Copper Cable 3-meter

Table 5 - Validated and Supported 25GbE Cables

Speed	Cable OPN #	Description
25GbE	MMA2P00-AS	Mellanox® transceiver, 25GbE, SFP28, LC-LC, 850nm, SR, up to 100m
25GbE	MFA2P10-A100	Mellanox® active optical cable 25GbE, SFP28, 100m
25GbE	MFA7A50-C030	Mellanox Active Fiber Hybrid Solution ETH 100GBE TO 4X25GBE QSFP28 TO 4XSFP28 30M
25GbE	MCP2M00-A005E26L	Mellanox Passive Copper Cable, ETH, UP TO 25GB/S, SFP28, 5M, BLACK, 26AWG, CA-L
25GbE	FTLF1436P3BCL	25GE LR 10km SFP28 Optical Transceiver
25GbE	FTLF8536P4BCL	25GE SR SFP+ Optical Transceiver
25GbE	LTF8507-PC07	SFP28 AOC 25G Ethernet SFP28 Active Optical Cable

1.2.4 Validated and Supported 40GbE Cables

Table 6 - Validated and Supported 40GbE Cables

Speed	Cable OPN #	Description
NA	MAM1Q00A-QSA	Mellanox® cable module, ETH 10GbE, 40Gb/s to 10Gb/s, QSFP to SFP+
NA	MAM1Q00A-QSA28	Mellanox® cable module, ETH 25GbE, 100Gb/s to 25Gb/s, QSFP28 to SFP28
40GbE	MC2210126-004	Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 4m
40GbE	MC2210126-005	Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 5m
40GbE	MC2210128-003	Mellanox Passive Copper Cable ETH 40GBE 40GbE QSFP 3M
40GbE	MC2210130-001	Mellanox Passive Copper Cable ETH 40GBE 40GbE QSFP 1M
40GbE	MC2210130-002	Mellanox Passive Copper Cable ETH 40GBE 40GbE QSFP 2M
40GbE	MC2210130-00A	Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 0.5m
40GbE	MC2210130-00B	Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 0.75m
40GbE	MC2210310-XXX	Mellanox Active Fiber Cable ETH 40GBE 40GbE QSFP from 3M up to 100M
40GbE	MC2210411-SR4L	Mellanox Optical Module 40GbE QSFP MPO 850NM UP TO 30M
40GbE	MC2210411-SR4	Mellanox Optical Module 40GbE QSFP MPO 850NM UP TO 100M
40GbE	MC2210411-SR4E	Mellanox Optical Module 40GbE QSFP MPO 850NM UP TO 300M
40GbE	QSFP-40G-SR-BD	Cisco 40GBASE-SR-BiDi, duplex MMF
40GbE	QSFP-40G-SR4	Cisco 40GBASE-SR4, 4 lanes, 850 nm MMF
40GbE	QSFP-H40G-ACU10M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 10-meter, active
40GbE	QSFP-H40G-AOC10M	Cisco 40GBASE-AOC QSFP direct-attach Active Optical Cable, 10-meter
40GbE	QSFP-H40G-CU1M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 1-meter, passive
40GbE	QSFP-H40G-CU3M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 3-meter, passive
40GbE	QSFP-H40G-CU5M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 5-meter, passive

Table 6 - Validated and Supported 40GbE Cables

Speed	Cable OPN #	Description
40GE	00D5811-N13445C	IBM Cable DAC 40GbE QSFP+ to QSFP+ Passive Copper 7m
40GE	L45593-D118-D10	QSFP-H40G-CU1M

1.2.5 Validated and Supported 50GbE Cables

Table 7 - Validated and Supported 50GbE Cables

Speed	Cable OPN #	Description
50GbE	MCP7H00-G001	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 2X50GBS QSFP28 TO 2XQSFP28 1M
50GbE	MCP7H00-G002	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 2X50GBS QSFP28 TO 2XQSFP28 2M
50GbE	MCP7H00-G003	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 2X50GBS QSFP28 TO 2XQSFP28 3M
50GbE	MCP7H00-G01A	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 2X50GBS QSFP28 TO 2XQSFP28 1.5M
50GbE	MCP7H00-G02A	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 2X50GBS QSFP28 TO 2XQSFP28 2.5M
50GbE	MFA7A20-C020	Mellanox® active fiber hybrid solution, ETH 100GbE to 2x50GbE, QSFP28 to 2xQSFP28, 20m

1.2.6 Validated and Supported 100GbE Cables

Table 8 - Validated and Supported 100GbE Cables

Speed	Cable OPN #	Description
100GbE	CBL-00195-02	100GbE QSFP28 to QSFP28 copper cable 3M
100GbE	MCP1600-C001	Mellanox Passive Copper Cable ETH 100GBE 100GBS QSFP LSZH 1M
100GbE	MCP1600-C002	Mellanox Passive Copper Cable ETH 100GBE 100GBS QSFP LSZH 2M
100GbE	MCP1600-C003	Mellanox Passive Copper Cable ETH 100GBE 100GBS QSFP LSZH 3M
100GbE	MCP1600-C005E26L	Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP28, 5m, Black, 26AWG, CA-L
100GbE	MCP1600-C00A	Mellanox Passive Copper Cable ETH 100GBE 100GBS QSFP LSZH 0.5M
100GbE	MCP1600-C01A	Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP, LSZH, 1.5m
100GbE	MCP1600-C02A	Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP, LSZH, 2.5m
100GbE	MCP1600-C03A	Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP, PVC, 3.5m 26AWG
100GbE	MCP7F00-A005R26L	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 5m, Colored, 26AWG, CA-L

Table 8 - Validated and Supported 100GbE Cables

Speed	Cable OPN #	Description
100GbE	MFA1A00-C003	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 3m
100GbE	MFA1A00-C005	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 5m
100GbE	MFA1A00-C010	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 10m
100GbE	MFA1A00-C015	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 15m
100GbE	MFA1A00-C020	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 20m
100GbE	MFA1A00-C030	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 30m
100GbE	MFA1A00-C050	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 50m
100GbE	MFA1A00-C100	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 100m
100GbE	MFS1200-C005	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 5m
100GbE	MFS1200-C010	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 10m
100GbE	MFS1200-C015	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 15m
100GbE	MFS1200-C020	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 20m
100GbE	MFS1200-C030	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 30m
100GbE	MFS1200-C050	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 50m
100GbE	MFS1200-C100	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 100m
100GbE	MMA1B00-C100_B	Mellanox® transceiver, up to 100GbE, QSFP28, MPO, 850nm, up to 100m OM3
100GbE	MMA1B00-C100D	Mellanox® Transceiver, 100GbE, QSFP28, MPO, 850nm, up to 100m
100GbE	MMS1C00-C500	Mellanox® transceiver, 100GbE, QSFP28, MPO, 1550nm PSM4, up to 2km
100GbE	MMS1C00-C500	Mellanox® transceiver, 100GbE, QSFP28, MPO, 1550nm PSM4, up to 2km

1.2.7 Validated and Supported FDR Cables

Table 9 - Validated and Supported FDR Cables

Speed	Cable OPN #	Description
FDR	MC2207126-004	Mellanox Passive Copper Cable VPI UP TO 56GB/S QSFP 4M

Table 9 - Validated and Supported FDR Cables

Speed	Cable OPN #	Description
FDR	MC2207128-003	Mellanox Passive Copper Cable VPI UP TO 56GB/S QSFP 3M
FDR	MC2207128-0A2	Mellanox Passive Copper Cable VPI UP TO 56GB/S QSFP 2.5M
FDR	MC2207130-001	Mellanox Passive Copper Cable VPI UP TO 56GB/S QSFP 1M
FDR	MC2207130-002	Mellanox Passive Copper Cable VPI UP TO 56GB/S QSFP 2M
FDR	MC2207130-00A	Mellanox Passive Copper Cable VPI UP TO 56GB/S QSFP 0.5M
FDR	MC2207130-0A1	Mellanox Passive Copper Cable VPI UP TO 56GB/S QSFP 1.5M
FDR	MC2207310-100	Mellanox Active Fiber Cable VPI UP TO 56GB/S QSFP from 3M up to 100M
FDR	MC2207310-XXX	Mellanox Active Fiber Cable VPI UP TO 56GB/S QSFP from 3M up to 100M
FDR	MC2207312-XXX	Mellanox Active Fiber Cable VPI UP TO 56GB/S QSFP from 3M up to 300M
FDR	MC220731V-XXX	Mellanox® Active Fiber Cable, VPI, up to 56Gb/s, QSFP, up to 100m
FDR	MC2207411-SR4L	Mellanox Optical Module IB FDR 56GB/S QSFP MPO 850NM up to 30M
FDR	MCP170L-F001	Mellanox® Passive Copper Cable, VPI, up to 56Gb/s, QSFP, LSZH, 1m
FDR	MCP170L-F002	Mellanox® Passive Copper Cable, VPI, up to 56Gb/s, QSFP, LSZH, 2m
FDR	MCP170L-F003	Mellanox® Passive Copper Cable, VPI, up to 56Gb/s, QSFP, LSZH, 3m

1.3 Tested Switches

1.3.1 Tested 10GbE Switches

Table 10 - Tested 10GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
10/40GbE	N/A	3064	48-port 10Gb/40Gb Switch	Cisco
10/40GbE	N/A	7050Q	16-port 40Gb Switch	Arista
10/40GbE	N/A	7050S	48-port 10Gb/40Gb Switch	Arista
10GbE	N/A	5548	Cisco 10GB ETH switch	Cisco
10GbE	N/A	G8264	BNT 10/40GB ETH switch	BNT
10GbE	N/A	QFX3500	Juniper 10/40GB ETH switch	Juniper
10GbE	N/A	S4810P-AC	48-port 10Gb/40Gb Switch	Force10
10GbE	SwitchX®	SX1016X-1BFR	64-Port 10GbE Switch System	Mellanox

1.3.2 Tested 40GbE Switches

Table 11 - Tested 40GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
10/40GbE	N/A	3064	48-port 10Gb/40Gb Switch	Cisco
10/40GbE	N/A	7050Q	16-port 40Gb Switch	Arista
10/40GbE	N/A	7050S	48-port 10Gb/40Gb Switch	Arista
40GbE	N/A	3132Q	Cisco 40GB ETH switch	Cisco
40GbE	N/A	7050QX	32-port 40Gb Switch	Arista
40GbE	N/A	G8316	BNT 40GB RackSwitch G8316	BNT
40GbE	N/A	S6000	32-port 40Gb Switch	Dell
40GbE	SwitchX®	SX1036B-1BFR	36-Port 40/56GbE Switch System	Mellanox

1.3.3 Tested 100GbE Switches

Table 12 - Tested 100GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
100GbE	N/A	7060CX	32-port 100Gb Switch	Arista
100GbE	N/A	93180YC-EX	48 x 10/25-Gbps fiber ports and 6 x 40/100-Gbps Quad Small Form-Factor Pluggable 28 (QSFP28) ports	Cisco
100GbE	N/A	C3232C	High-Density, 100 Gigabit Ethernet Switch	Cisco
100GbE	N/A	CE8860-4C-EI	24x10GE (SFP+) or 25GE (SFP28) and 2x100GE switch	Huawei
100GbE	Spectrum	SN2410-CB2F	48-port 25GbE + 8-port 100GbE Open Ethernet ToR Switch System	Mellanox
100GbE	Spectrum	SN2700-CS2R	32-port Non-blocking 100GbE Open Ethernet Spine Switch System	Mellanox
100GbE	Spectrum	SN2740-CB2F1	32-port Non-blocking 100GbE Open Ethernet Spine Switch System	Mellanox
100GbE	N/A	Wedge 100-32X R04	32-port 100G QSFP28 - Leaf/Spine Switch, power-to-port airflow, DC Power	Edgecore

1.4 Tools, Switch Firmware and Driver Software

Firmware Rev 14.23.1020 is tested with the following tools, Switch firmware, and driver software:

Table 13 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_OFED	4.4-1.0.1.0/4.3-1.0.1.0/

Table 13 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_EN (MLNX_OFED based code)	4.4-1.0.1.0/4.3-1.0.1.0/
WinOF-2	2.0/1.90
MFT	4.10.0/4.9.0
VMware	- ESXi 6.5 v4.16.13.5 - ESXi 6.0 v4.15.13.2
MLNX-OS	- SwitchX: 3.6.4930 - Spectrum: 3.6.4930
SwitchX®/SwitchX®-2 Firmware	9.4.4040
Spectrum™ Firmware	13.1530.0136
Linux Inbox Drivers	- Ubuntu 14.04.4 - Ubuntu 15.10 - Ubuntu 16.04 - Ubuntu 16.04.1 - Ubuntu 16.10 - SLES12.1 - SLES12.2 - RHEL6.8 - RHEL7.2 - RHEL7.3
Windows Inbox Drivers	Windows Server 2016

1.5 Supported FlexBoot, UEFI



Please be aware that not all firmware binaries contain FlexBoot or UEFI, support may vary between cards (see [Section 1.1, “Supported Devices”, on page 6.](#))

Firmware Rev 14.23.1020 supports the following FlexBoot:

Table 14 - Supported FlexBoot, UEFI

Expansion ROM	Supported Version
FlexBoot	3.5.504
UEFI	14.16.15

1.6 Revision Compatibility

Firmware Rev 14.23.1020 complies with the following programmer’s reference manual:

- Mellanox Adapters Programmer’s Reference Manual (PRM), Rev 0.47 or later*, which has Command Interface Revision 0x5. The command interface revision can be retrieved by means of the QUERY_FW command and is indicated by the field *cmd_interface_rev*.

2 Changes and New Features in Rev 14.23.1020

Table 15 - Changes and New Features in Rev 14.23.1020

Feature/Change	Description
Virtual Functions (VF)	Increased the number of VFs that can work with full VMQoS (8 TC) per PFs as follow: - in dual port devices to: 0-21, 33-45 VFs (22-32 VFs has single TC) - in single port devices to: 0-64 VFs
InfiniBand	Added support for IPoIB non-default Partition Keys (PKeys). Now the PKey values can be modified in the PKey table without the need of recreating the IPoIB (underlay) QPs.
SR-IOV in Multi-Host/Socket-Direct	[Beta] Added support for SR-IOV in Multi-Host/Socket-Direct.
Virtualization	Reduced firmware's memory consumption to increase the supported number of VFs per PF to up to 100.
Tools/Driver Version	Added support for QUERY_DRIVER_VERSION command. This command allows the PF driver to query its VFs driver version which was set by the SET_DRIVER_VERSION command.
Resiliency	Shutting Down RDMA QPs with Excessive Retransmissions is a mechanism used to detect excessive retransmissions for an RC connection, and to close the connection in response to it. If the number of retransmissions due to a Local Ack Timeout, NAK-Sequence Error, or Implied NAK, during a specified period, exceeds the specified threshold, the QP will be handled as if the IB spec defined Retry Count was exceeded.
Diagnostic Counters	Added new diagnostic counters to evaluate the number of ICMC hits and misses for particular resources.
Bug Fixes	See Section 4, "Bug Fixes History", on page 21

3 Known Issues

The following table describes known issues in this firmware release and possible workarounds.
For a list of old firmware Known Issues, please see ConnectX-4Lx Firmware Archived Known Issues file

(http://www.mellanox.com/pdf/firmware/ConnectX4Lx-Firmware_Archived_Known_Issues.pdf)

Table 16 - Ethernet Rate Limit per VF in RoCE Mode Limitations

Adapter Card	Dual Port Device				Single Port Device	
	w/o LAG (TOTAL_VFS>32)		With LAG (TOTAL_VFS<32)		w/o LAG	
	w/o QoS	Full QoS	w/o QoS	Full QoS	w/o QoS	Full QoS

Table 17 - Ethernet Rate Limit per VF in InfiniBand Mode Limitations

Adapter Card	Dual Port Device		Single Port Device	
	w/o LAG		w/o LAG	
	w/o QoS	Full QoS	w/o QoS	Full QoS

Table 18 - Known Issues (Sheet 1 of 4)

Internal Ref.	Issue
1332714	Description: The maximum “read” size of MTRC_STDB is limited to 272 Bytes.
	Workaround: Set the MTRC_STDB.read_size to the maximum value of 0x110=272 Bytes
	Keywords: Access register, MTRC_STDB, tracer to dmesg, fwtrace to dmesg
	Discovered in Version: 14.23.1020
1408994	Description: FTE with both forward (FWD) and encapsulation (ENCAP) actions is not supported in the SX NIC Flow Table.
	Workaround: N/A
	Keywords: SX NIC Flow Table
	Discovered in Version: 14.23.1020

Table 18 - Known Issues (Sheet 2 of 4)

Internal Ref.	Issue																									
1350794	Description: Encapsulation / Decapsulation support in steering has the following limitations: - Encapsulation / Decapsulation can be open on the FDB only if all VFs are non active. - Encapsulation / Decapsulation supports single mode only: FDB / NIC. Opening tables of both types is not supported. - RoCE is supported only when no Encapsulation / Decapsulation is opened on the FDB. - Encapsulation / Decapsulation per device support:<table><thead><tr><th></th><th></th><th>NIC</th><th>FDB</th></tr></thead><tbody><tr><td rowspan="2">ConnectX-4</td><td>encap</td><td>NO</td><td>YES</td></tr><tr><td>decap</td><td>NO</td><td>NO</td></tr><tr><td rowspan="2">ConnectX-4 Lx</td><td>encap</td><td>NO</td><td>YES</td></tr><tr><td>decap</td><td>NO</td><td>YES</td></tr><tr><td rowspan="2">ConnectX-5</td><td>encap</td><td>YES</td><td>YES</td></tr><tr><td>decap</td><td>YES</td><td>YES</td></tr></tbody></table>			NIC	FDB	ConnectX-4	encap	NO	YES	decap	NO	NO	ConnectX-4 Lx	encap	NO	YES	decap	NO	YES	ConnectX-5	encap	YES	YES	decap	YES	YES
			NIC	FDB																						
	ConnectX-4	encap	NO	YES																						
		decap	NO	NO																						
	ConnectX-4 Lx	encap	NO	YES																						
		decap	NO	YES																						
ConnectX-5	encap	YES	YES																							
	decap	YES	YES																							
	Workaround: N/A																									
	Keywords: Steering Encapsulation / Decapsulation																									
	Discovered in Version: 14.23.1020																									
1424873	Description: Modifying VMQoS rate limiter parameters during traffic might cause transmission failure																									
	Workaround: N/A																									
	Keywords: VMQoS, rate limiter																									
	Discovered in Version: 14.22.1002																									
1398556	Description: Although the default FEC configuration of the MMA2L20-AR module is FC-FEC, different FEC configuration can be set by the PPLM access register / MLXLink.																									
	Workaround: N/A																									
	Keywords: FEC, PPLM access registry, MMA2L20-AR module																									
	Discovered in Version: 14.22.1002																									
1316221	Description: Health counter increases every 50ms instead of 10ms.																									
	Workaround: N/A																									
	Keywords: Health counter																									
	Discovered in Version: 14.22.1002																									
1284452/1282926	Description: mlxconfig tool presents all possible expansion ROM images, instead of presenting only the existing images.																									
	Workaround: N/A																									
	Keywords: mlxconfig																									
	Discovered in Version: 14.22.1002																									

Table 18 - Known Issues (Sheet 3 of 4)

Internal Ref.	Issue
1277762	Description: An Ethernet multicast loopback packet is not counted (even if it is not a local loopback packet) when running the <code>nic_receive_steering_discard</code> command.
	Workaround: N/A
	Keywords: Ethernet multicast loopback packet
	Discovered in Version: 14.22.1002
1190753	Description: When a dual-port VHCA sends a RoCE packet on its non-native port, and the packet arrives to its affiliated vport FDB, a mismatch might happen on the rules that match the packet source vport.
	Workaround: N/A
	Keywords: Dual-port VHCA, RoCE packet, vport FDB
	Discovered in Version: 14.22.1002
1079027/1126921	Description: Occasionally, when adding module info page for Bell 1G BaseT module to the <code>mlxlink</code> data, the information is not updated correctly.
	Workaround: N/A
	Keywords: Bell 1G BaseT module, <code>mlxlink</code>
	Discovered in Version: 14.21.2010
1177500	Description: QoS for SR-IOV between VPorts per Traffic Class and Para-VPort achieves the optimal performance on Port1.
	Workaround: N/A
	Keywords: vPort TC QoS
	Discovered in Version: 14.21.1000
1168594	Description: RoCE Dual Port Mode (a.k.a Multi-Port vHCA: MPV) is not supported in Multi-Host setups.
	Workaround: N/A
	Keywords: Multi-Port vHCA, Multi-Host
	Discovered in Version: 14.21.1000
1176407/1171665	Description: Secure Firmware devices block any "not secure" access to the CR-Space, private ICMDs and Flash random access commands. Thus, tools/capabilities such as <code>ibdump</code> or <code>wqedump</code> or packet sniffing do not function properly.
	Workaround: Load a secure customer token to use any blocked tools.
	Keywords: Secure Firmware
	Discovered in Version: 14.21.1000
1119458	Description: When RoCE Dual Port Mode is enable, if VHCA0 has a different VLAN stripping configuration on its E-SW vport context than the configuration on the affiliated vport (VHCA1), the NIC steering on the VLAN might be wrong for the single-port VHCA traffic.
	Workaround: N/A
	Keywords: VLAN, dual port affiliation
	Discovered in Version: 14.21.1000

Table 18 - Known Issues (Sheet 4 of 4)

Internal Ref.	Issue
1159246	Description: OEM commands must be used with minre field == 0 when setting the host Rate Limiter, meaning the user is prevented from configuring bandwidth share fields of any hosts/PF.
	Workaround: N/A
	Keywords: Rate Limiter
	Discovered in Version: 14.21.1000
1047184	Description: RDMA resq_local_length_error and resp_remote_invalid_request counters do not function properly.
	Workaround: N/A
	Keywords: RDMA counters
	Discovered in Version: 14.21.1000
1132277	Description: An encapsulated packet that its inner packet has a unicast MAC and its outer packet has a multicast MAC (or vice versa) might be counted by the hardware on the wrong vPort counter.
	Workaround: N/A
	Keywords: Packet encapsulation/decapsulation, unicast MAC, multicast MAC
	Discovered in Version: 14.21.1000
1059782	Description: If BMC does not support the recovery flow, no communication will occur after fwreset.
	Workaround: N/A
	Keywords: BMC
	Discovered in Version: 14.20.1010
1048128	Description: Using ECN with RDMA Read backpressure on the NIC side, may cause low percentage of pauses.
	Workaround: N/A
	Keywords: ECN, RDMA
	Discovered in Version: 14.20.1010
1031744	Description: Same flow counter cannot be used on different table types.
	Workaround: N/A
	Keywords: Flow counter
	Discovered in Version: 14.20.1010
1009067	Description: In case of an ip_protocol match (on UDP/TCP) related to fragmented packet, the l4_type match might be missed when the hardware steering does not see the L4 headers.
	Workaround: Add to the driver ip_frag match for all steering rules that use ip_protocol match.
	Keywords: ip_protocol match, L4 headers
	Discovered in Version: 14.20.1010

4 Bug Fixes History

Table 19 lists the bugs fixed in this release. For a list of old firmware Bug Fixes, please see ConnectX-4 Lx Firmware Archived Bug Fixes file

(http://www.mellanox.com/pdf/firmware/ConnectX4Lx-Firmware_Archived_Bug_Fixes.pdf)

Table 19 - Bug Fixes History (Sheet 1 of 8)

Internal Ref.	Issue
1417642	Description: Fixed an issue that showed the link as down even when the link was up due to missing asynchronous event.
	Keywords: Asynchronous event
	Discovered in Version: 14.23.1000 (Internal)
	Fixed in Release: 14.23.1020
1391409	Description: Fixed a rare scenario that resulted in NMI hardware failure.
	Keywords: Virtualization, NMI
	Discovered in Version: 14.18.1000
	Fixed in Release: 14.23.1020
1433172	Description: Fixed a rare scenario that resulted in TX pipe getting hang.
	Keywords: TX pipe
	Discovered in Version: 14.18.1000
	Fixed in Release: 14.23.1020
1288757	Description: Closed the vport as part of the fast teardown flow, to prevent Ack to be sent without been scatter to memory.
	Keywords: Fast Teardown Flow
	Discovered in Version: 14.22.1002
	Fixed in Release: 14.23.1020
1356954	Description: Fixed an issue that occasionally caused the link to drop after running the 'reboot' command due to I2C being stuck.
	Keywords: I2C stuck, link in polling
	Discovered in Version: 14.22.1002
	Fixed in Release: 14.23.1020
1390343	Description: Fixed an issue that caused the tx_pause_storm_warning_events and the tx_pause_storm_error_events counters to increase even when there was TX pause.
	Keywords: TX pause storm events
	Discovered in Version: 14.22.1002
	Fixed in Release: 14.23.1020
1225025	Description: Fixed a congestion control issue in large scale A2A case.
	Keywords: Congestion Control
	Discovered in Version: 14.21.2010
	Fixed in Release: 14.23.1020

Table 19 - Bug Fixes History (Sheet 2 of 8)

Internal Ref.	Issue
1331545	Description: Fixed a congestion control issue that occurred in a system with a large number of QPs.
	Keywords: Congestion Control
	Discovered in Version: 14.21.2010
	Fixed in Release: 14.23.1020
1333168	Description: Fixed an issue that resulted in a wrong sensor state report when running the PLDM sensor reading command and when a negative temperature was set with the hook.
	Keywords: Temperature
	Discovered in Version: 14.22.1002
	Fixed in Release: 14.23.1020
1342760	Description: Fixed a speed issue for link mode Speed25GBBaseSR.
	Keywords: Link speed
	Discovered in Version: 14.21.2010
	Fixed in Release: 14.23.1020
1063148	Description: Pause duration: Physical port counters count in 512bits quantas, instead of microseconds.
	Keywords: Pause duration, Physical port counters
	Discovered in Version: 14.20.1010
	Fixed in Release: 14.23.1020
1187411	Description: Added the option to avoid unintentionally powering off the backplane port cage upon reboot when in standby mode.
	Keywords: Backplane port cage, standby mode
	Discovered in Version: 14.21.1000
	Fixed in Release: 14.22.1002
1231791	Description: Fixed an issue that caused the driver to return a wrong logical OR of the 2 physical ports, when querying the vport state when the LAG was enabled.
	Keywords: LAG, vport
	Discovered in Version: 14.21.2010
	Fixed in Release: 14.22.1002
1252833	Description: Increased the Full Wire Speed (FWS) threshold value to improve EDR link results.
	Keywords: Full Wire Speed (FWS) threshold, EDR
	Discovered in Version: 14.21.1000
	Fixed in Release: 14.22.1002
1281622	Description: Fixed an issue that resulted in “Destroy LAG” command failure if a VFs received an FLR while its affinity QPs were open.
	Keywords: ECMP / SR-IOV LAG
	Discovered in Version: 14.21.1000
	Fixed in Release: 14.22.1002

Table 19 - Bug Fixes History (Sheet 3 of 8)

Internal Ref.	Issue
1172293	Description: When RoCE Dual Port mode is enabled, tcpdump is not functional on the 2nd port.
	Keywords: Dual Port vHCA, Multi-port, RoCE Dual Port
	Discovered in Version: 14.21.1000
	Fixed in Release: 14.22.1002
1124226	Description: Fixed an issue that caused QP connection timeout due to firmware not being able to handle duplicate packets with AckReq bit set. The fix stopped ignoring duplicate AckReq packets to avoid timeout on the sender side.
	Keywords: QP connection timeout
	Discovered in Version: 14.21.1000
	Fixed in Release: 14.21.2010
1179155	Description: MPFS load balance (DUP_MAC_ACTION==LOAD_BALANCE(1)) is not working as long as IB_ROUTING_MODE/SRIOV_IB_ROUTING_MODE is configured to LID.
	Keywords: MPFS load balance, LID, GID
	Discovered in Version: 14.21.1000
	Fixed in Release: 14.21.2010
1155392	Description: Fixed an issue that caused a SX engine deadlock (the SX engine handles software port/priority changes for a specific Send Queue) when more than a single SX engine handled the "prio diff" flow simultaneously, thus caused the hardware to get stuck. The issue happened as the firmware releases the SX engine and waits for various operation to complete. However, due to a race that allows a different transaction to get into the SX engine and cause the lock to be taken by it, the SX engine release is prevented.
	Keywords: SX engine deadlock, system hang
	Discovered in Version: 14.21.1000
	Fixed in Release: 14.21.2010
1168271	Description: Fixed an issue that caused the system to hang while changing QPTS/QPDPM/QPDP parameters during traffic.
	Keywords: QPTS/QPDPM/QPDP, system hang
	Discovered in Version: 14.21.1000
	Fixed in Release: 14.21.2010
1190215	Description: Fixed unfairness between senders in RoCE LAG while ECN is configured.
	Keywords: RoCE LAG, ECN
	Discovered in Version: 14.21.1000
	Fixed in Release: 14.21.2010
1121688	Description: Fixed an issue which displayed diagnostic counters only on the adapter that was initialized first, which is the counters' owner. The owner received correct values, while the other adapter only received zeros.
	Keywords: Diagnostic counters
	Discovered in Version: 14.21.1000
	Fixed in Release: 14.21.2010

Table 19 - Bug Fixes History (Sheet 4 of 8)

Internal Ref.	Issue
1175146	Description: Fixed an issue that caused the rdma_cm traffic to fail on the 2nd port when more than 32 VFs were configured, when the RoCE Dual Port vHCA (a.k.a Multi-Port vHCA: MPV) feature was enabled.
	Keywords: MPV
	Discovered in Version: 14.21.1000
	Fixed in Release: 14.21.2010
1167218/ 1168567	Description: Fixed an issue related to RDMA_CM driver that might have caused the QP Rate Limit to be activate unexpectedly and reduce the bandwidth significantly on this QP.
	Keywords: Performance
	Discovered in Version: 14.20.1010
	Fixed in Release: 14.21.1000
1122718	Description: Fixed an issue that caused low throughput when ECN was enabled in a many-to-one scenario.
	Keywords: ECN
	Discovered in Version: 14.20.1010
	Fixed in Release: 14.21.1000
1090723	Description: Fixed an issue that wrongly reported the maximum temperature in a setup as the current temperature regardless of the actual temperature.
	Keywords: PCI Gen4 receiver
	Discovered in Version: 14.20.1010
	Fixed in Release: 14.21.1000
1091424	Description: Fixed Virtual Addressing Capability incorrect report in NC-SI (Sideband).
	Keywords: NC-SI
	Discovered in Version: 14.20.1010
	Fixed in Release: 14.21.1000
1134021	Description: Improved RX sensitivity per Vcore.
	Keywords: RN, Vcore
	Discovered in Version: 14.20.1010
	Fixed in Release: 14.21.1000
1099880	Description: Disabled the option to write to the protected modules to avoid receiving NACK upon module initialization.
	Keywords: Thermal temperature
	Discovered in Version: 14.20.1010
	Fixed in Release: 14.21.1000

Table 19 - Bug Fixes History (Sheet 5 of 8)

Internal Ref.	Issue
1111202	Description: Enabled cold_flicks reset when sending invalidate_all command to invalidate all the ECs.
	Keywords: CQE, QoS
	Discovered in Version: 14.20.1010
	Fixed in Release: 14.21.1000
743242/ 938019	Description: SR-IOV min & max rate limiter can only support up to 64 VFs per port.
	Keywords: SR-IOV min & max rate limiter
	Discovered in Version: 14.20.1010
	Fixed in Release: 14.21.1000
1046427/ 1047180	Description: Fixed the issue where ECN did not function as expected when the number of QPs > ~500 per host.
	Keywords: ECN
	Discovered in Version: 14.20.1010
	Fixed in Release: 14.21.1000
1054335/ 1054671	Description: Fixed the issue where when using UD RoCE multicast traffic over SR-IOV, packets were scattered to all the attached QPs in the e-sw (PF and its VFs) and not only on the vport that was specified in the e-se FDB.
	Keywords: UD RoCE multicast traffic, SR-IOV
	Discovered in Version: 14.20.1010
	Fixed in Release: 14.21.1000
1060650	Description: Fixed a link issue on Intel 10GbE Optical module PN: R8H2F, Y3KJN.
	Keywords: Intel 10GbE Optical module
	Discovered in Release: 14.18.2000
	Fixed in Release: 14.20.1010
1052064	Description: Fixed an issue that caused the device to hang upon warm reboot.
	Keywords: Warm reboot
	Discovered in Release: 14.18.2000
	Fixed in Release: 14.20.1010
1047533	Description: Fixed an issue that caused the TX traffic not to send packets when using VF index (ARI) bigger than 127.
	Keywords: VFs
	Discovered in Release: 14.18.2000
	Fixed in Release: 14.20.1010
1009614	Description: Fixed a scaling issue with more than 1k QPs for ECN by moving from per QP caching to per IP to allow better scale with number of host in the fabric.
	Keywords: Performance
	Discovered in Release: 14.18.2000
	Fixed in Release: 14.20.1010

Table 19 - Bug Fixes History (Sheet 6 of 8)

Internal Ref.	Issue
1013035	Description: Removed the following false alarm asserts reported in Multi-Host environment: • host blocked due to flow control credits • dead irisc timer indication was increased
	Keywords: False alarm asserts, Multi-Host
	Discovered in Release: 14.18.2000
	Fixed in Release: 14.20.1010
1041108	Description: Enabled firmware resync of the internal clocks after getting out of the standby mode to prevent PTP time sync from getting out of sync after system warm-rebooted due to system getting into a low-power (standby) mode.
	Keywords: PTP time sync, standby mode
	Discovered in Release: 14.18.2000
	Fixed in Release: 14.20.1010
1047693	Description: When running RoCE over VRRP, enabled the device to receive RoCE packet with different source MAC than the original RoCE packet's destination MAC, to allow routing between different subnets.
	Keywords: RoCE over VRRP, Destination MAC
	Discovered in Release: 14.18.2000
	Fixed in Release: 14.20.1010
1050234	Description: Fixed an issued that caused LLDP not to enable PFC configuration currently when DCBX transitioning flow control configurations was set from Global Pause to PFC.
	Keywords: RoCE Lossy & ECN
	Discovered in Release: 14.18.2000
	Fixed in Release: 14.20.1010
1047533	Description: Rephrased and improved external troubleshoot messages in PDDR register.
	Keywords: PDDR register
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.20.1010
999261	Description: Improved SR-IOV performance.
	Keywords: SR-IOV
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.20.1010
954822	Description:: The ipoib_enhanced_offloads indication in the HCA capabilities reports 0 while SRIOV_EN=1.
	Keywords: SR-IOV, IPoIB Offloads
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.20.1010

Table 19 - Bug Fixes History (Sheet 7 of 8)

Internal Ref.	Issue
959464	Description: When the Max Rate Limiter is enabled and a Teardown/FLR is issued upon the last gvmi with <code>max_rate_limiter</code> enabled Teardown/FLR, the hardware remains enabled (<code>rate_limiter_en = 1</code>). ** "max rate limiter enabled" = at least 1 (per chip). <code>create/modify_scheduling_elemnt</code> command has been issued by the driver, with <code>max_average_bw != 0</code> .
	Keywords: Teardown/FLR, Max Rate Limiter
	Discovered in Version: 14.18.1000
	Fixed in Release: 14.20.1010
981598	Description: Fixed an issue on an ETH port with SR-IOV enabled that prevented packets from reaching the BMC (failure in steering loopback resolution) if the BMC addresses were configured after VF initialization, and the VF was trying to send traffic to the BMC (that located on the same phy port).
	Keywords: BMC, SR-IOV, packets
	Discovered in Release: 14.18.2000
	Fixed in Release: 14.20.1010
906144	Description: Fixed an issue which caused the rate limiter not to function when setting a rate to tc 7.
	Keywords: QOS - ETH - rate limit per TC
	Discovered in Release: 14.18.2000
	Fixed in Release: 14.20.1010
893261	Description: Fixed the PCIe TX glitch during Recovery.Speed state of the link training to PCIe Gen3.
	Keywords: PCIe TX glitch
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.20.1010
1002190	Description: Fixed an issue related to the <code>PortRcvDataVLExtended/PortXmitDataVLExtended</code> parameter that caused the counters' value to be reported in octets instead of dwrods.
	Keywords: Counters
	Discovered in Release: 14.18.2000
	Fixed in Release: 14.20.1010
1025741/ 781339/ 1050373	Description: QP ULP modes 0 and 1 cannot be assigned to the same Multicast group.
	Keywords: Multicast Group (MCG), QPs
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.20.1010
913451	Description: Fixed an issue in standby (WoL) modes only that caused the actual current consumption in 1.2V rail to be higher by <33mA than the advertised values although the total IC consumption is as advertised.
	Keywords: Standby (WoL) modes, current consumption
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.20.1010

Table 19 - Bug Fixes History (Sheet 8 of 8)

Internal Ref.	Issue
852744	Description: Mapping an SL to VL 15 is currently not supported. Trying to do so, will cause a health buffer fatal internal error report.
	Keywords: SL to VL mapping
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.20.1010
1019003/ 1019039/ 995878	Description: Fixed an issue causing physical errors observed on Ixia 100GbE receiver.
	Keywords: Ixia 100GbE receiver
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.20.1010

5 Firmware Changes and New Feature History

Table 20 - Firmware Changes and New Feature History (Sheet 1 of 9)

Feature/Change	Description
Disable SL/diff Flow	Added support for disable SL/diff flow to avoid performance degradation for single queue using multiple priorities. This functionality should not be used when DCB (PFC, ETS) is enabled.
Software Reset Flow	Software Reset Flow enables the device to recover from fatal errors. The flow includes software detection of a fatal error, automatic creations of an mstdump file for future debug by the software, and resetting of the device. The feature is enabled using an mlxconfig command. Note: The flow is currently not supported on Multi host devices, Socket Direct devices and devices running management traffic (NCSI, MCTP).
Steering Discard Packet Counters	Any received packet which is dropped by the device is accounted for. To enable this functionality, the following counters were added to count the discard packets (per vport): • <code>nic_receive_steering_discard</code>: Number of packets that completed the NIC Receive Flow Table steering, and were discarded because they did not match any flow in the final Flow Table. • <code>receive_discard_vport_down</code>: Number of packets that were steered to a VPort, and discarded because the VPort was not in a state to receive packets. • <code>transmit_discard_vport_down</code>: Number of packets that were transmitted by a vNIC, and discarded because the VPort was not in a state to transmit packets.
Virtual Functions (VF)	Increased the number of VFs that can work with full VMQoS (8 TC) per PFs as follow: • in dual port devices to 20 VFs • in single port devices to 58 VFs
Pause Frame Duration and XOFF Resend Time	Increased the Pause Frame Duration and the XOFF Resend Time to the maximum value defined by the specification.
PCI Relax Ordering	mlxconfig configuration can now enable or disable forced PCI relaxed ordering in mkey_context. If this feature is enabled, the software per mkey configuration is ignored.
vport Mirroring	Packets are mirrored based on certain mirroring policy. The policy is set using the “set FTE command” that supports forward action in the ACL tables (ingress/egress). The firmware support the following destination list format: 1. new destination vport (analyzer) 2. another Flow Table this way, the driver can forward the SX/RX packet related to the vport once it reaches the ACL table (forward it to the analyzer vport).
Resiliency: Special Error Event	Firmware uses error events to monitor the health of core transport engines, both Rx and Tx, and to detect if a system hang occurred and was not cured by other error mechanisms. Upon such detection, events are sent to the driver to perform any required action (e.g., software reset).

Table 20 - Firmware Changes and New Feature History (Sheet 2 of 9)

Feature/Change	Description
10GBaseT module	Added support for 10GBaseT modules connected to a QSFP cage. Note: This connectivity supported was only tested with eNet's E10GS-FPT-ENC 10GBase-T SFP+ device, and 10Gtek's ASF-10G-T when using firmware v14.22.1002
QP's Creation Time	Accelerated QP's creation time.
SR-IOV LID based Routing Mode	SR-IOV default routing mode is now LID based. The configuration change is available via mlxconfig tool. Note that in such mode, the VF will get its own LID, hence the GRH is not required. Note: LID based routing support for vports is supported using SM v4.8.1
Expansion ROM	Added PXE and UEFI to additional ConnectX-4 Lxadapter cards. ConnectX-4 Lx now holds PXE, x86-UEFI and Arm-UEFI
Bug Fixes	See Section 4, "Bug Fixes History", on page 21
Rev. 14.21.2010	
Query vPort Environments (Debug Counters)	Debug counters are a group of counters that handle traffic performance issue related to firmware overhead in transport flow. The following are the additional counters added to this firmware version: - current_q_under_processor_handle - total_q_under_processor_handle - qp_priority_update_flow
Bug Fixes	See Section 4, "Bug Fixes History", on page 21
Rev. 14.21.1000	
Receiver Signal Integrity Improvements	Raised the network link only with phase greater than 15 ticks to improve signal integrity.
	Extended measurement test between 2 similar RX configurations.
	Moved the data path to use the second input buffer to improve signal integrity.
Dynamic Bandwidth Allocation per Host	Added support for new chassis MC OEM commands which allowed dynamically configuring rate limit and bandwidth share, per host and physical function.
RoCE Dual Port Mode (a.k.a Multi-Port vHCA: MPV)	Enables the usage of a dual port Virtual HCA (vHCA) to share RDMA resources (e.g., MR, CQ, SRQ, PDs) across the two Ethernet (RoCE) NIC network ports and display the NIC as a dual port device. For this feature to function properly, the following requirements must be met: - Either the LAG or the Dual Port mode is enabled by the driver - Dual port device: both ports must be set as ETH - In ConnectX-4/ConnectX-4 Lx adapter cards, the maximum allowed number of VFs per PF is 32. - Function per port is enabled Note: This feature is only supported in single host device

Table 20 - Firmware Changes and New Feature History (Sheet 3 of 9)

Feature/Change	Description
DSCP	Added QPDPM register to support dynamic mapping between DSCP and priority.
	Added trust level for QoS prioritization according to the DSCP or PCP.
	Added ingress buffer management for: - ingress traffic mapping to a buffer according to priority - buffers sizes and lossless parameters
Steering Rules Rate Improvement	Improved steering rules update rate to up to 50K rules per sec.
Windows SR-IOV Enhanced eIPoIB	Enabled Windows SR-IOV Enhanced eIPoIB (without Secure Connection) for Windows-over-Windows setups.
Driver CR Dump	crdump operation takes a snapshot of the device's crspace dword-by-dword. It enables the driver to collect debug information upon firmware failure.
Secured Firmware Update	Secure Firmware Updates provides devices with the ability to verify digital signatures of new firmware binaries, in order to ensure that only officially approved versions are installed on the devices. Note: This feature is only available in adapter cards that support this feature.
Cables	Changed the default FEC mode for cables with attenuation 16 and below from RS to FC.
ECN	Enabled ECN (CongestionControl) by default for all priorities on Ethernet ports.
Bug Fixes	See Section 4, "Bug Fixes History", on page 21
Rev. 14.20.1010	
DSCP	Added trust level for QoS prioritization according to the DSCP or PCP.
	[Beta] Added ingress buffer management for: - ingress traffic mapping to a buffer according to priority - buffers sizes and lossless parameters
Secured Firmware Updates	[Beta] Secure Firmware Updates provides devices with the ability to verify digital signatures of new firmware binaries, in order to ensure that only officially approved versions are installed on the devices. Note: This feature is only available in adapter cards that support this feature.
Relaxed Ordering	[Beta] Added support for relaxed ordering write in memory keys.
RDMA Counters	Enhanced RDMA counter
Outbound Buffer Overflow Mechanism	Added the ability to enable and disable drop of packets in case of PCIe outbound buffer overflow. In addition added a counter for counting these drops.
TLV for PCI class code	Added 2 new per Host TLVs (see Table 30, "Per host Settings," on page 50)

Table 20 - Firmware Changes and New Feature History (Sheet 4 of 9)

Feature/Change	Description
Fast Teardown	Enables fast unloading driver by using Teardown HCA with op_mode=1 (force_close). For further information, refer to the PRM.
SFP Power Flow Improvement (level 2,1)	Added support for SFP power class.
Bug Fixes	See Section 4, “Bug Fixes History”, on page 21
Rev. 14.18.2000	
Bug Fixes	See Section 4, “Bug Fixes History”, on page 21
Rev. 14.18.1000	
RX Loss (BaseT link down indication)	Added logical link indication in SFP to BaseT modules and disabled logical link when peer port is down.
SFP Rate	Added support for 10GbE in 25GbE SFP optical modules
PDDR	Enables mlxlink tool to collect data on the PHY link status and provides link down reasons and additional link related information.
KR Tx Response	Enabled TX configuration response and movement during Link Training in Ethernet.
Phy Test mode	Added support at lane rate of 12.89Gb.
Head of Queue (HoQ) per TC	Limits the amount of time a packet may head a Traffic Class (TC) transmission queue, without being transmitted. Stale packets are discarded. Active by default for TCs adhering to link level flow control
User Access Region (UAR) 4KB Granularity Allocation	UAR page size currently is set to 4KB and not according to what the system page size determines.
No Driver NIC (NODNIC) Performance Improvement	Improved performance of: • Doorbell from User Access Region (UAR) • Clear interrupt from User Access Region (UAR)
Counters	Added support for additional transport counters.
On Demand Paging (ODP) DC	Added ODP support for DC.
Scatter to CQE on Sender for DC	Enabled scatter-to-CQE for sent packets for DC.
CQ modify	Enabled moderation period modification in CQ modify command.
VMQ: Rate limit per Function	[Beta] Added support for minimum/maximum rate limit per vport in SR-IOV.
Network traffic between UEFI-Shell and OS	Enabled network traffic between UEFI-Shell and OS.
non-RDMA capable VFs	Enabled the PF to force disable RoCE for its VFs.
PRM: Access Registers	Added 2 new access registers: • Management Capabilities Mask Register • Ports Capabilities Mask Register Fields For further information, please refer to the PRM.
Loopback Enabled/Disabled	Enabled VNIC the control to enable/disable its local loopback traffic.

Table 20 - Firmware Changes and New Feature History (Sheet 5 of 9)

Feature/Change	Description
RDMA RX Flow Table	Added the option to open a receive RDMA Flow Table and to forward RoCE traffic to some destination QP.
Bug Fixes	See Section 4, “Bug Fixes History”, on page 21
Rev. 14.17.2020	
GENEVE & IP-in-IP Stateless Offload	[Beta] Added support for IP-in-IP and GENEVE network protocols encapsulated into IP frame (L2 tunneling). Encapsulation is suggested as a means to alter the normal IP routing for datagrams, by delivering them to an intermediate destination that would otherwise not be selected based on the (network part of the) IP Destination Address field in the original IP header. Note: For driver support, please see the Release Notes/User Manual of the relevant OS driver.
Bug Fixes	See Section 4, “Bug Fixes History”, on page 21
Rev. 14.17.1010	
Multi-Host LID Base Routing	Added support for Multi-Host LID base routing. This feature requires a new OpenSM (v4.7.1 and above which comes with MLNX_OFED 3.3-2.0.0.0) with the following attributes: • qos TRUE • lmc 2 (if there is no quad host in the fabric, you can set the lmc to 1) • virt_enabled 2 Note: Multi-Host LID base routing can be configured by the INI only. The default is 0
Resilient RoCE	Resilient RoCE is the ability to send RoCE traffic over a lossy network (a network without flow control enabled), without the need to enable flow control on the network. The ability is accomplished by enabling ECN on both the Switch and the Host.
Multi-Host L3/L4 Classification	Enables load balancing in the Multi PF Switch layer (MPFS) based on the L3/L4 headers
Virtual Functions (VF) per Port	Increased the number of VFs from 64 to 95 per Physical Function (PF). Note: When increasing the number of VFs, the following limitations must be taken into consideration: $\text{server_total_bar_size} \geq (\text{num_pfs}) * (2\log_pf_uar_bar_size + 2\log_vf_uar_bar_size * \text{total_vfs})$ $\text{server_total_msix} \geq (\text{num_pfs}) * (\text{num_pf_msix} + \text{num_vfs_msix} * \text{total_vfs})$ Note: For the maximum number of VFs supported by your driver, please refer to your drivers' Release Notes or User Manual.
HCA Port Flap Counter	Added support for Port Flap Counter.
Fixed Buffer Size (KSM)	Limits the buffer size for all entries to improve performance. KSM is used when associating Key Length My Virtual Address (KLMs) with fixed memory size.
NULL Mkey	This entry (null_mkey) is use to indicate non-present KLM/KSM entries. When accessing is, it causes the device to generate page fault event.

Table 20 - Firmware Changes and New Feature History (Sheet 6 of 9)

Feature/Change	Description
Out-of-Band Online Firmware Update: Firmware Update over PLDM	PLDM firmware burning is based on the DMTF spec DSP0267 (draft 9). The feature enables upgrading firmware and expansion ROM images using the PLDM protocol over MCTP (over PCIe). By doing so, a supporting BMC can query and upgrade the firmware without using OS based tools.
New Group in Ports Performance Counters (PPCNT)	Added a new physical layer statistics counters group. The new group includes BER counters, FEC error correction, clear time, and additional physical layer counters. For further information, please refer to the Ethernet Adapters Programming Manual (PRM) .
Permanent Link Up Mode	Enables the user to set a certain link up state for an unlimited period of time. This mode has 3 states: - Aux power (standby) - Reboot/boot/driver unloaded - the server is active and no driver is up - Driver is up - at least one driver is up (the time between init HCA and teardown or FLR)
No Driver NIC (NODNIC) Performance Improvement	Added support for Doorbell from User Access Region (UAR).
SR-IOV: Rate Limit Per Function	[Beta] Added support for maximum rate limit per function in SR-IOV.
Firmware Resiliency: Suppress Pauses	Allows the user to configure the adapter card to stop sending pauses after x when the receive port is unavailable (in a hang state).
Performance Back-pressure Counters	[Beta] Added support for new performance counters.
Data Center Bridging Exchange (DCBX)	DCBX is used by DCB devices to exchange configuration information with directly connected peers. DCBX uses Link Layer Discovery Protocol (LLDP) to exchange parameters between two link peers. For further information, please refer to the PRM.
Access Register: Default Values Revert	Allows network port registers to revert to their default values when the driver is restarted or the host is rebooted.
Link up Modes	Added additional network link up modes. The new modes decide when to keep the network link up. The new modes are: - keep_eth_link_up - keep_ib_link_up - keep_link_up_on_boot - keep_link_up_on_standby
Bug Fixes	See Section 4, “Bug Fixes History”, on page 21
Rev. 14.16.1020	
Bug Fixes	See Section 4, “Bug Fixes History”, on page 21
Rev. 14.16.1006	
Explicit Congestion Notification (ECN)	[Beta] Explicit Congestion Notification (ECN) is an extension to the Internet Protocol and to the Transmission Control Protocol. ECN allows end-to-end notification of network congestion without dropping packets.

Table 20 - Firmware Changes and New Feature History (Sheet 7 of 9)

Feature/Change	Description
64 VFs per port	Increased the number of VFs from 32 to 64 per PF. Note: When increasing the number of VFs, the following limitations must be taken into consideration: <pre>server_total_bar_size >= (num_pfs)*(2log_pf_uar_bar_size + 2log_vf_uar_bar_size*total_vfs) server_total_msix >= (num_pfs)*(num_pf_msix + num_vfs_msix *total_vfs)</pre>
RoCE Link Aggregation (RoCE LAG)	[Beta] RoCE Link Aggregation provides failover and link aggregation capabilities. In this mode, only one IB port, that represents the two physical ports, is exposed to the application layer. For further information, please refer to the PRM.
OVS Offload	Mellanox Accelerated Switching And Packet Processing (ASAP²) Direct technology allows to offload OVS by handling OVS data-plane in Mellanox ConnectX-4 / ConnectX-4 Lx NIC hardware (Mellanox Embedded Switch or eSwitch) while maintaining OVS control-plane unmodified. The current actions supported by ASAP² Direct include packet parsing and matching, forward, drop along with VLAN push/pop or VXLAN encap/decap and HW based packet/byte flow statistics.
Virtual Extensible LAN (VXLAN) encapsulation/decapsulation	Virtual Extensible LAN (VXLAN) is a network virtualization technology that improves scalability problems associated with large cloud computing deployments. It tunnels Ethernet frames within Ethernet + IP + UDP frames. Mellanox implements VXLAN encapsulation and decapsulation in the hardware.
Data Center Bridging Exchange (DCBX)	[Beta] DCBX is used by DCB devices to exchange configuration information with directly connected peers. DCBX uses Link Layer Discovery Protocol (LLDP) to exchange parameters between two link peers. For further information, please refer to the PRM.
FCS no scatter / FCS check	Enables the user to control whether or not to scatter Frame Check Sequence (FCS) or to check FCS functionality.
Packet Pacing	[Beta] Send Queues (SQ/ Send queue of QP) may be individually rate limited, thus, allowing granular rate control over specific SW-defined flows. A rate-limited flow is allowed to transmit a few packets before its transmission rate is evaluated, and the next packet is scheduled for transmission accordingly.
PRBS Patterns Generation and Tuning	A new PHY test mode in which the device can generate different PRBS patterns for SerDes tuning purpose. For further information, please refer to PRM registers: PPAOS, PPTT, PPRT.
Management Controller Transport Protocol (MCTP) over PCI	Added support for MCTP host management over PCI
OCBB / OCSD support after mlxfwreset	Added support for OCBB/OCSD memory pointers restoration after mlxfwreset
MCTP media migration	Added support for MCTP media migration between SMBUS and PCI
Cables	Removed the RX amplitude configuration on some cable types

Table 20 - Firmware Changes and New Feature History (Sheet 8 of 9)

Feature/Change	Description
Bug Fixes	See Section 4, “Bug Fixes History”, on page 21
Rev. 14.14.2036	
Scatter FCS in RQ	Enables software to scatter or strip FCS in RQ.
Bug Fixes	See Section 4, “Bug Fixes History”, on page 21
Rev. 14.14.1100	
CQE Time Stamping	Keeps track of the creation of a packet. A time-stamping service supports assertions of proof that a datum existed before a particular time.
Priority Flow Control (PFC)	Applies pause functionality to specific classes of traffic on the Ethernet link.
RDMA retransmission counters	Custom port counters provide the user a clear indication about RDMA send/receive statistics and errors.
Link Layer Discovery Protocol (LLDP)	The Link Layer Discovery Protocol (LLDP) is a vendor-neutral Link Layer protocol in the Internet Protocol Suite used by network devices for advertising their identity, capabilities, and neighbors on a IEEE 802 LAN. The protocol is formally defined in IEEE 802.1AB.
1GbE Link Speed	ConnectX-4 Lx adapters now support 1Gb/s Ethernet connectivity in addition to 10GigE, 25GigE, 40GigE, 50GigE
Flow Steering Counters	Provides a clear indication of Flow Steering statistics and errors.
WQE Inline Header	The minimal amount of packet headers inlined in the WQE's Eth Segment.
table-miss Flow	A flow table may include a table-miss flow entry, which renders all Match Fields wildcards. If a packet does not match a flow entry in a flow table, this is a table miss. The behavior on a table miss depends on the table configuration. A table-miss flow entry in the flow table may specify how to process unmatched packets.
Strided WQE User Space	Striding RQ is a receive queue comprised by work queue elements (i.e. WQEs), where multiple packets of LRO segments (i.e. message) are written to the same WQE.
SR-IOV (EN eSwitch & RoCE)	Single Root IO Virtualization (SR-IOV) is a technology that allows a physical PCIe device to present itself multiple times through the PCIe bus.
Vector Calculation/Erasure Coding Offload	Uses the HCA for offloading erasure coding calculations.
Firmware Image Time Stamping for Multi-Host Environment	Enables the administrator to add a timestamp to the firmware they want to upgrade to avoid situations where one host tries to upgrade the firmware and another tries to downgrade; which may lead to two or more unnecessary server reboots. For further information, please refer to MFT User Manual .
Link params modification via access registers	The change includes the following: 1. Changed port configuration which required link re-training (such as speed) 2. PAOS down 3. PAOS up This change, will cause the link to toggle and new configurations to take effect.

Table 20 - Firmware Changes and New Feature History (Sheet 9 of 9)

Feature/Change	Description
Checksum Calculation on Image/Device	Flint utility allows performing an MD5 checksum on the non-persistent sections of the firmware image. For further information, please refer to MFT User Manual .
Rev. 14.12.1240	
Bug Fixes	See Section 4, “Bug Fixes History”, on page 21
Rev. 14.12.1100	
Port Link	Reduced the port link-up time when negotiating according to Clause 73 (DME)
Rev. 14.12.0780	
PCI	- PCIe Function Level Reset (FLR) - Power Management L2/L3 flow support
Ethernet Network	- Large Receive Offload (LRO) - Large Send Offload (LSO) - Receive Side Scaling (RSS) - Global Pause - RoCEv1.0/RoCEv2.0 - Flow Steering - Sniffer Ethernet - Rate Limiter (at Beta level) - Multi packet WQE - Minimal Bandwidth Guarantee (ETS) - Explicit Congestion Notification (ECN) - Priority Flow Control (PFC)
PRM	- Strided SRQ - Self Loopback support - Transport Domain support - CQ2EQ remapping - Added support for the following commands: - MODIFY/QUERY_ESW_VPORT_CONTEXT - QUERY/MODIFY_CONG_STATUS - QUERY/MODIFY_CONG_PARAMS - QUERY_CONG_STATISTICS - ADD/DELETE_VXLAN_UDP_DPORT
Virtualization	- VXLAN/NVGRE Stateless offload In this release, this feature is supported through Windows ONLY - SR-IOV EN (at Beta level)
Performance	CQE zipping
Misc	- Wake-on-Lane/Standby - FlexBoot/UEFI support
Non-Volatile Configuration	Non-Volatile Configuration (NVConfig). For the complete list, please refer to Section 9, on page 50.

6 FlexBoot Changes and New Features

For further information, please refer to FlexBoot Release Notes (www.mellanox.com > Software > InfiniBand/VPI Drivers > FlexBoot).

Table 21 - FlexBoot Changes and New Features (Sheet 1 of 3)

Version	Description
Rev. 3.5.504	
FlexBoot UI	Added “PXE boot without fail-over to iSCSI” and “iSCSI boot without fail-over to PXE boot” options for legacy_boot_protocol configuration. For further information, refer to Mellanox PreBoot Drivers User Manual.
	Enabled PXE to expose the current link speed of the system when in ETH mode.
Rev. 3.5.403	
Enable/Disable FlexBoot in EXPROM via mlxconfig	Added PXE support to additional ConnectX-4 Lx adapter cards. Enabling/Disabling FlexBoot in ConnectX-4 Lx in EXPROM is done via mlxconfig. The default value is: FLEXBOOT enable Please note, not all cards are compiled with FlexBoot. For the full list of the OPNs compiled with FlexBoot, please refer to Section 1.1, “Supported Devices”, on page 6
VLAN Priority	Set the default VLAN priority to 0.
Link Aggregation Control Protocol (LACP)	LACP support is disabled by default. It can be enabled via mlxconfig.
Rev. 3.5.305	
PXE Boot	Added ESC option as an abort key during PXE boot process.
FlexBoot Link Aggregation Control Protocol (LACP)	Enabled/disabled FlexBoot LACP support by editing the INI configuration.
Serial Console	Removed Serial Console support in the ConnectX-4 Lx adapter card.
Upstream sync	Synced the source with iPXE (upstream sync)
Rev. 3.5.210	
Promiscuous VLAN mode	Added support for promiscuous VLAN mode.
Expansion ROM version	Enabled expansion ROM (exp_rom) version exposition according to the new specification (e.g. expose ARCH in flint tool).
FlexBoot UI	Added a FlexBoot menu support for NV_POWER_CONF. Now power consumption configuration is supported from the FlexBoot menu.
	Enhanced FlexBoot/firmware debug capability using Flexboot UI. Added the reg_dump option to the panic_behavior configuration in the FlexBoot menu
Bug Fixes	See Section 6.2, “FlexBoot Bug Fixes History”, on page 46
Rev. 3.5.110	

Table 21 - FlexBoot Changes and New Features (Sheet 2 of 3)

Version	Description
Networking	Ethernet only: The MTU value is set to 1500 upon driver's bring up.
Rev. 3.5.109	
Performance	Performance enhancements in Ethernet mode
FlexBoot UI	Added support for "Undi network wait timeout"
	Enhanced FlexBoot/firmware debug capability using Flexboot UI
Upstream sync	Synced the source with iPXE (upstream sync)
Rev. 3.4.903	
iSCSI re-imaging	Enables the user to install a new image on active iSCSI target
FlexBoot UI	Added new configuration for network link type for supported cards (ConnectX-4 VPI cards)
	Enabled boot configuration menu in ConnectX-4 when the physical port is IB
Booting	Enabled booting with non-default Pkey in ConnectX-4 when the physical port is IB
Link Status	Removed link status line printout at boot time
Boot Menu	Changed the Bus:Device:Function format in boot menu, from PCI - Bus:Dev.Func to 0000:Bus:Dev.Func
Upstream sync	Synced the source with iPXE (upstream sync)
Rev. 3.4.812	
FlexBoot UI	Added debug prints option in the FlexBoot boot menu. For further information, please refer to FlexBoot and UEFI User Manual.
System Diagnosis	Added the ability to diagnose problems in released ROMs by enabling the debug log levels for specific modules. Note: This ability should be used only when debug session is needed.
Interrupts	Added support for ConnectX-4/ConnectX-4 Lx interrupts
Upstream sync	Synced the source with iPXE (upstream sync)
Rev. 3.4.719	
IPv6	Added IPv6 support
x64 Architecture	Added x64 architecture support in ConnectX-4 and Connect-IB adapter cards

Table 21 - FlexBoot Changes and New Features (Sheet 3 of 3)

Version	Description
SHELL CLI	Removed support for the following SHELL CLI commands: • Non-volatile option storage commands • SAN boot commands • Menu commands • Login command • Sync command • DNS resolving command • Time commands • Image crypto digest commands • Loopback testing commands • VLAN commands • PXE commands • Reboot command For further information, please refer to: http://ipxe.org/cmd
Upstream sync	Synced the source with iPXE (upstream sync)
Rev. 3.4.650	
Image size	Added support for .mrom images larger than 128kB
Adapter Cards	Added support for ConnectX-4 EN and ConnectX-4 Lx EN
Flat real mode	Moved to flat real mode when calling INT 1a,b101 to avoid BIOSes issues
Spanning Tree Protocol	Added support for detecting Spanning Tree Protocol non-forwarding ports (RSTP/MSTP)
Upstream sync	Synced the source with iPXE (upstream sync)

6.1 FlexBoot Known Issues

Table 22 - FlexBoot Known Issues (Sheet 1 of 5)

Internal Ref.	Description
1422604	Description: The FlexBoot TFTP enters an infinite loop when it tries to resolve the server name by a DNS request, and the DNS server dose not response to the request.
	Workaround: Terminate the TFTP process by pressing Ctrl + C
	Keywords: Boot, TFTP, DNS
	Discovered in Version: 3.5.504
1295727	Description: In Secure Host mode, the Ctrl + B option will be “read only” and changes will not be applied and may cause unknown behavior.
	WA: N/A
	Keywords: Secure Host mode, Ctrl + B option
	Discovered in Version: 3.5.403
1066544	Description: Chain-loading boot-loaders that works with interrupts fails to boot on multi-host adapter cards
	Workaround: N/A
	Keywords: Chain-loading, boot
	Discovered in Version: 3.5.403
1157875	Description: Pressing any of the arrow keys during boot might cause the boot process to be aborted.
	Workaround: N/A
	Keywords: Abort boot, arrows, FlexBoot
	Discovered in Release: 3.5.305
841198	Description: FlexBoot fails to boot when the following occurs: • Boot priority is set to iSCSI • The iSCSI TCP/IP parameters via DHCP is disabled • iSCSI boot fails or iSCSI boot to target configuration is set to disable
	Workaround: N/A
	Keywords: PXE boot, iSCSI
656001	Description: Booting from WDS and Windows DHCP server when only Option 66 is enabled (without Option 67), is not supported.
	Workaround: N/A
	Keywords: DHCP

Table 22 - FlexBoot Known Issues (Sheet 2 of 5)

Internal Ref.	Description
776057	Description: Citrix PVS boot is not supported.
	Workaround: N/A
	Keywords: Citrix PVS boot
689460	Description: FlexBoot uses system UUID to generate the client DUID-UUID as per RFC 6355, the data conveyed with DHCPv6 Code 1 (Option ID).
	Workaround: N/A
	Keywords: DUID-UUID
928217	Description: Installing ESXi 6.5/6.0 on iSCSI target is currently not supported.
	Workaround: N/A
	Keywords: ESXi 6.5/6.0, iSCSI target
-	Description: To use the DHCP server to identify ipxe requests when using <code>undi-only.kpxe</code> or <code>ipxe.pxe</code> when booting over IB requires special configuration. (see the Workaround below).
	Workaround: Add to the DHCP host declaration the MAC identification alongside the option 61 DUID. For example: <pre>host ib-client1 { option dhcp-client-identifier = ff:00:00:00:00:02:00:00:02:c9:00:<Port-GUID> ; hardware ethernet <Port-MAC> ; fixed-address <IPoIB Address> ; filename "ipxe.pxe" ; if exists user-class and option user-class = "iPXE" { filename "pxelinux.0" ; } }</pre>
	Keywords: undionly.kpxe or ipxe.pxe
-	Description: Due to interoperability issue between the ESXi installer and the lpxelinux boot-loader, when trying to install ESXi 6.5 on iSCSI target using lpxelinux.0 as a bootloader, a PSOD occurs.
	Workaround: Use FlexBoot (or iPXE) to load mboot.c32 directly instead of pxelinux.0 using the script below: <pre>#!/ipxe set base /nfs/Esxi-6.5_INBOX chain \${base}/mboot.c32 -c \${base}/boot.cfg BOOTIF=01- \${mac:hexhyp}</pre> where the "set base ..." specifies a suitable absolute path. Note: iPXE does not need an absolute path, however, mboot.c32 requires it.
	Keywords: mboot.c32, PSOD,

Table 22 - FlexBoot Known Issues (Sheet 3 of 5)

Internal Ref.	Description
976878	Description: When using bootloader grub2 to boot WDS, if the WDS boot fails, an RSOD might appear.
	Workaround: N/A
	Keywords: Bootloader grub2, WDS, RSOD
-	Description: Several BIOS vendors have limited boot-vector space and may not display FlexBoot in their boot menu.
	Workaround: Disable the embedded NIC boot agent in BIOS
	Keywords: BIOS
-	Description: In several BIOS, the server might hang during FlexBoot booting due to wrong configuration of the PMM.
	Workaround: N/A
	Keywords: BIOS
-	Description: Only EBX, ESI, DS, ES registers can be saved in Boot Entry.
	Workaround: N/A
	Keywords: BIOS
-	Description: If a client returned control to the BIOS after a successful connection to an iSCSI target (but did not boot from it), then, unexpected behavior may occur.
	Workaround: Follow the instructions described in the FlexBoot UM for the proper iSCSI boot/install
	Keywords: BIOS
673114/821899	Description: FlexBoot banner might not be shown in some BIOSes.
	Workaround: N/A
	Keywords: BIOS
-	Description: In some cases, PXE boot will not work if the client was given only the file-name without next-server (siaddr).
	Workaround: N/A
	Keywords: PXE Boot
-	Description: PXE boot after iSCSI boot with static configuration is currently not supported.
	Workaround: N/A
	Keywords: PXE Boot

Table 22 - FlexBoot Known Issues (Sheet 4 of 5)

Internal Ref.	Description
-	Description: Boot over VLAN with IB port is currently not supported.
	Workaround: N/A
	Keywords: PXE Boot
-	Description: Some faulty boot loaders do not close the underlying UNDI device which may result in unexpected behavior and possible system crash after the OS starts to load.
	Workaround: N/A
	Keywords: PXE Boot
-	Description: Chain-loading gPXE stack is not supported.
	Workaround: N/A
	Keywords: PXE Boot
647143	Description: Executing a partial boot loop while only downloading the NBP and selecting localboot is unsupported and may cause undefined behavior.
	Workaround: N/A
	Keywords: PXE Boot
-	Description: iSCSI over IB is not tested.
	Workaround: N/A
	Keywords: iSCSI
-	Description: iSCSI over DCB is not supported.
	Workaround: N/A
	Keywords: iSCSI
-	Description: FlexBoot supports only a single active iSCSI connection. Thus, when iSCSI-boot via Port 1 succeeds to connect but fails to boot, it will fail to connect via Port 2.
	Workaround: N/A
	Keywords: iSCSI
-	Description: Boot retries is currently not functional when booting from iSCSI.
	Workaround: N/A
	Keywords: iSCSI
655800	Description: iSCSI over IPv6 is not supported.
	Workaround: N/A
	Keywords: iSCSI

Table 22 - FlexBoot Known Issues (Sheet 5 of 5)

Internal Ref.	Description
-	Description: Boot menu is displayed as READ ONLY if the HCA card does not support flash configuration.
	Workaround: N/A
	Keywords: User Interface
-	Description: FlexBoot Boot Menu will not be visible in serial output.
	Workaround: N/A
	Keywords: User Interface
-	Description: Setting the number of Virtual Functions higher than the machine's memory capability may cause memory issues and system instability.
	Workaround: N/A
	Keywords: Virtualization
-	Description: SLAM, FTP, HTTPS and SRP are currently not supported.
	Workaround: N/A
	Keywords: Protocols
-	Description: Occasionally, using the Spanning Tree Protocol (STP) in the switches may cause packet drops and boot failure in the system.
	Workaround: Enable the "edgemode" if disabled on the switch, or use either portfast or edgemode functionality on the switch ports connected to the NICs.
	Keywords: Protocols
655800	Description: IPv6 can only run if a RADVD service is running in the network.
	Workaround: N/A
	Keywords: Protocols
-	Description: IPv6 over IB is not supported.
	Workaround: N/A
	Keywords: Protocols
655800	Description: Enabling IPv6 first and then IPv4 is currently not supported.
	Workaround: N/A
	Keywords: Protocols

6.2 FlexBoot Bug Fixes History

Table 23 - FlexBoot Bug Fixes History (Sheet 1 of 2)

Version	Issue
1072419	Description: The FlexBoot DHCP loops indefinitely when it continuously gets NACK on the DHCP requests. On some setups, it might also cause an RSOD after a continuous looping.
	Keywords: Bootloader grub2, WDS, RSOD
	Discovered in Release: 3.5.305
	Fixed in Release: 3.5.504
1157875	Description: Pressing any of the arrow keys during boot might cause the boot process to be aborted.
	Keywords: Abort boot, arrows, FlexBoot
	Discovered in Release: 3.5.305
	Fixed in Release: 3.5.403
1113560	Description: Fixed an issue that prevented the first iSCSI target parameters to be reset to their default values.
	Keywords: iSCSI target
	Discovered in Release: 3.5.110
	Fixed in Release: 3.5.305
843209	Description: Fixed an issue which caused the link not to raise in the second port which is set as IB when the first port is ETH in PXE.
	Keywords: Link up, Ports
	Discovered in Release: 3.4.903
	Fixed in Release: 3.5.110
847950	Description: Fixed wrong default value of Boot-To-Target in FlexBoot configuration.
	Keywords: Boot-To-Target, FlexBoot configuration
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.903
691148	Description: When connecting a pre-configured port with VLAN to an IB fabric, the port runs as Ethernet port with the VLAN tag.
	Keywords: VLAN, Port Management
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.903
792432	Description: Booting PXE using Grub2.X over HP G9/G8 servers results in system hang.
	Keywords: PXE boot, Grub2.X, HP G9/G8
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.903

Table 23 - FlexBoot Bug Fixes History (Sheet 2 of 2)

Version	Issue
737512	Description: If the client gets "PXE boot menu" when contacting the DHCP, it will PXE boot first regardless of the boot priority.
	Keywords: ISCSI, DHCP
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.812
690792	Description: If the PMM fails to allocate memory, the system hangs since FlexBoot cannot load from the expansion ROM.
	Keywords: PMM, expansion ROM
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.812
689068	Description: In hybrid BIOSes, if the BIOS loads legacy driver without closing the UEFI driver, the legacy driver fails to load.
	Keywords: BIOS, legacy mode
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.812
634794	Description: Enabled 'boot_pci_busdevfn' initialization when booting from UNDI loader.
	Keywords: UNDI loader
	Discovered in Release: 3.4.650
	Fixed in Release: 3.4.719
-	Description: Removed the instruction that enabled write-protected section modifications after POST.
	Keywords: PXE Boot
	Discovered in Release: 3.4.650
	Fixed in Release: 3.4.719

7 UEFI Changes and Major New Features

Table 24 - UEFI Changes and New Features

Category	Description
Rev. 14.16.15	
Hii Menu	Added “PXE boot without fail-over to iSCSI” and “iSCSI boot without fail-over to PXE boot” options for legacy_boot_protocol configuration. For further information, refer to Mellanox PreBoot Drivers User Manual.
	Enabled UEFI to expose the current link speed of the system in the Hii menu.
Rev. 14.15.19	
Enable/Disable UEFI in EXPROM via mlxconfig	Added UEFI support to additional ConnectX-4 Lx adapter cards. - ConnectX-4 Lx adapter cards are compiled with x86-UEFI and Arm-UEFI Enabling/Disabling UEFI in ConnectX-4 Lx in EXPROM is done via mlxconfig. The default values are: - UEFI_X86 disabled - UEFI_AARCH64 disabled For the full list of the OPNs, please refer to Section 1.1, “Supported Devices”, on page 6

7.1 UEFI Known Issues

The following is a list of general limitations and known issues of the various components of this UEFI release.

Table 25 - UEFI Known Issues

Internal Ref.	Description
1295727	Description: In Secure Host mode, the Hii protocol will be “read only”, changes will not be applied and it may cause unknown behavior.
	WA: N/A
	Keywords: Secure Host mode
	Discovered in Version: 14.15.19
798073	Description: UEFI driver is not supported on Supermicro X9DEW (BIOS version 3.0c).
	WA: N/A
	Keywords: BIOS, Supermicro X9DEW
-	Description: Burning the UEFI driver will remove the Flexboot driver (Legacy BIOS driver) from the firmware.
	WA: N/A
	Keywords: UEFI burning, Flexboot

8 Unsupported Features and Commands

8.1 Unsupported Features

The following advanced feature are unsupported in the current firmware version:

Table 26 - Unsupported Features

Feature	ConnectX-4 Lx
The following service types: • SyncUMR • Mellanox transport • PTP • RAW IPv6	No
INT-A not supported for EQs only MSI-X	No
PCI VPD write flow (RO flow supported)	No
Streaming Receive Queue (STRQ) and collapsed CQ	No
Subnet Manager (SM) on VFs	No
DC in Multi-Host, SR-IOV, and Ethernet (RoCE)	No
RoCE LAG for VFs	No
RoCE LAG in Multi-Host/Socket-Direct	No
Mutlihost Ethernet	No

8.2 Unsupported Commands

- QUERY_MAD_DEMUX
- SET_MAD_DEMUX
- CREATE_RQ - MEMORY_RQ_RMP
- MODIFY_LAG_ASYNC_EVENT

9 Supported Non-Volatile Configurations

Table 27 - Per-physical Port Settings

Name	Parameter Index
VPI settings	0x12
RoCE CC	0x107
RoCE CC ECN	0x108
LLDP_NB_DCBX	0x18E
NV_QOS_CONF	0x192
NV_QOS_CAP	0x193
NV_KEEP_LINK_UP	0x190

Table 28 - Global Settings

Name	Parameter Index
PCI settings	0x80
PCI setting capabilities	0x81
TPT settings	0x82
TPT capabilities	0x83
Option ROM ini	0x100
Option ROM capabilities	0x101
NV_SW_OFFLOAD_CONF	0x10A
NV_PACKET_PACING	0x10C

Table 29 - Per host/function Settings

Name	Parameter Index
Wake-on-LAN	0x10
External Port	0x192

Table 30 - Per host Settings

Name	Parameter Index
NV_PCI_CONF	0x80
NV_PCI_CAP	0x81