



Mellanox ConnectX[®]-4 Lx Firmware Release Notes

Rev 14.20.1030

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 NON-INFRINGEMENT 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 2017. Mellanox Technologies Ltd. All Rights Reserved.

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

All other trademarks are property of their respective owners.

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

Table of Contents

Release Update History	5
Chapter 1 Overview	6
1.1 Supported Devices	6
1.2 Supported Cables and Modules	6
1.2.1 Validated and Supported 1GbE Cables	6
1.2.2 Validated and Supported 10/40GbE Cables	7
1.2.3 Validated and Supported 25GbE Cables	8
1.2.4 Validated and Supported QDR/FDR10 Cables	9
1.2.5 Validated and Supported 50Gbs Cables	10
1.2.6 Validated and Supported FDR Cables	10
1.2.7 Validated and Supported 100GB/s Cables	11
1.3 Tested Switches	12
1.3.1 Tested 10/40GbE Switches	12
1.3.2 Tested 100GbE Switches	12
1.4 Tools, Switch Firmware and Driver Software	13
1.5 Supported FlexBoot, UEFI	13
1.6 Revision Compatibility	14
Chapter 2 Changes and New Features in Rev 14.20.1030	15
Chapter 3 Known Issues	16
Chapter 4 Bug Fixes History	23
Chapter 5 Firmware Changes and New Feature History	37
Chapter 6 FlexBoot Changes and New Features	44
6.1 FlexBoot Known Issues	46
6.2 FlexBoot Bug Fixes History	51
Chapter 7 Unsupported Features and Commands	53
7.1 Unsupported Features	53
7.2 Unsupported Commands	53
Chapter 8 Supported Non-Volatile Configurations	54

List of Tables

Table 1:	Release Update History	5
Table 2:	Supported Devices	6
Table 3:	Validated and Supported 1GbE Cables	6
Table 4:	Validated and Supported 10/40GbE Cables	7
Table 5:	Validated and Supported 25GbE Cables	8
Table 6:	Validated and Supported QDR/FDR10 Cables	9
Table 7:	Validated and Supported 50Gbs Cables	10
Table 8:	Validated and Supported FDR Cables	10
Table 9:	Validated and Supported 100GB/s Cables	11
Table 10:	Tested 10/40GbE Switches	12
Table 11:	Tested 100GbE Switches	12
Table 12:	Tools, Switch Firmware and Driver Software	13
Table 13:	Supported FlexBoot, UEFI	13
Table 14:	Changes and New Features in Rev 14.20.1030	15
Table 15:	Known Issues	16
Table 16:	Bug Fixes History	23
Table 17:	Firmware Changes and New Feature History	37
Table 18:	FlexBoot Changes and New Features	44
Table 19:	FlexBoot Known Issues	46
Table 20:	FlexBoot Bug Fixes History	51
Table 21:	Per-physical Port Settings	54
Table 22:	Global Settings	54
Table 23:	Per host/function Settings	54
Table 24:	Per host Settings	54

Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 14.20.1030	August 15, 2017	Initial version of this firmware release.

1 Overview

These are the release notes for the ConnectX®-4 Lx adapters firmware Rev 14.20.1030. 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

Device Part Number	PSID	Device Name	Compiled with FlexBoot	Compiled with UEFI ^a
SN30L11422	LNV0000000003	Lenovo ThinkSystem ConnectX®-4 Lx PCIe 25Gb 2-Port SFP28 Ethernet Adapter by Mellanox	Yes	No
01GR252	LNV2420110034	Lenovo ThinkSystem ConnectX®-4 Lx PCIe 25Gb 2-Port SFP28 Ethernet Adapter by Mellanox	Yes	Yes
00MM952	LNV2430110027	Mellanox ConnectX®-4 Lx 1x40GbE QSFP+ Adapter	Yes	Yes
00MN992 / 00MN990	LNV2950110034	Mellanox ConnectX®-4 Lx ML2 1x25GbE SFP+ Adapter	Yes	Yes
SN30L21983	LNV2960110034	Lenovo ThinkSystem ConnectX®-4 Lx ML2 25Gb 2-Port SFP28 Ethernet Adapter by Mellanox	Yes	Yes

a. If you need to compile your adapter card with an UEFI expansion ROM, please contact Mellanox Support (support@mellanox.com)

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 optical module, Base-T

1.2.2 Validated and Supported 10/40GbE Cables

Table 4 - Validated and Supported 10/40GbE 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
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
40GB/S	MC2210126-004	Mellanox® Passive Copper Cable, ETH 40GbE, 40Gb/s, QSFP, 4m
40GB/S	MC2210126-005	Mellanox® Passive Copper Cable, ETH 40GbE, 40Gb/s, QSFP, 5m
40GB/S	MC2210128-003	Mellanox Passive Copper Cable ETH 40GBE 40GB/S QSFP 3M
40GB/S	MC2210130-001	Mellanox Passive Copper Cable ETH 40GBE 40GB/S QSFP 1M
40GB/S	MC2210130-002	Mellanox Passive Copper Cable ETH 40GBE 40GB/S QSFP 2M
40GB/S	MC2210130-00A	Mellanox® Passive Copper Cable, ETH 40GbE, 40Gb/s, QSFP, 0.5m
40GB/S	MC2210130-00B	Mellanox® Passive Copper Cable, ETH 40GbE, 40Gb/s, QSFP, 0.75m
40GB/S	MC2210310-XXX	Mellanox Active Fiber Cable ETH 40GBE 40GB/S QSFP from 3M up to 100M
40GB/S	MC2210411-SR4L	Mellanox Optical Module 40GB/S QSFP MPO 850NM UP TO 30M
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
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

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
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
40GB/S	QSFP-40G-SR-BD	Cisco 40GBASE-SR-BiDi, duplex MMF
40GB/S	QSFP-40G-SR4	Cisco 40GBASE-SR4, 4 lanes, 850 nm MMF
40GB/S	QSFP-H40G-ACU10M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 10-meter, active
40GB/S	QSFP-H40G-AOC10M	Cisco 40GBase-AOC QSFP direct-attach Active Optical Cable, 10-meter
40GB/S	QSFP-H40G-CU1M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 1-meter, passive
40GB/S	QSFP-H40G-CU3M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 3-meter, passive
40GB/S	QSFP-H40G-CU5M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 5-meter, passive
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

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
25GB/S	MCP2M00-A001	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1m

Table 5 - Validated and Supported 25GbE Cables

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

1.2.4 Validated and Supported QDR/FDR10 Cables

Table 6 - Validated and Supported QDR/FDR10 Cables

Speed	Cable OPN #	Description
QDR	MC2206125-007	Mellanox Passive Copper Cable IB QDR 40GB/S QSFP 7M
QDR	MC2206126-006	Mellanox Passive Copper Cable IB QDR 40GB/S QSFP 6M
FDR10	MC2206128-004	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 4M
FDR10	MC2206128-005	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 5M

Table 6 - Validated and Supported QDR/FDR10 Cables

Speed	Cable OPN #	Description
FDR10	MC2206130-001	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 1M
FDR10	MC2206130-002	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 2M
FDR10	MC2206130-003	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 3M
FDR10	MC2206130-00A	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 0.5M
FDR10	MC2206310-XXX	Mellanox Active Fiber Cable IB QDR/FDR10 40GB/S QSFP from 3M up to 100M
FDR10	MC2210411-SR4	Mellanox Optical Module 40GB/S QSFP MPO 850NM UP TO 100M
FDR10	MC2210411-SR4E	Mellanox Optical Module 40GB/S QSFP MPO 850NM UP TO 300M
FDR10	MFS4R12CB-XXX	Mellanox Active Fiber Cable VPI UP TO 40GB/S QSFP from 3M up to 100M

1.2.5 Validated and Supported 50Gbs Cables

Table 7 - Validated and Supported 50Gbs Cables

Speed	Cable OPN #	Description
50GE	MCP7H00-G001	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 2X50GBS QSFP28 TO 2XQSFP28 1M
50GE	MCP7H00-G002	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 2X50GBS QSFP28 TO 2XQSFP28 2M
50GE	MCP7H00-G003	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 2X50GBS QSFP28 TO 2XQSFP28 3M
50GE	MCP7H00-G01A	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 2X50GBS QSFP28 TO 2XQSFP28 1.5M
50GE	MCP7H00-G02A	Mellanox Passive Copper Hybrid Cable ETH 100GBE TO 2X50GBS QSFP28 TO 2XQSFP28 2.5M

1.2.6 Validated and Supported FDR Cables

Table 8 - Validated and Supported FDR Cables

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

Table 8 - Validated and Supported FDR Cables

Speed	Cable OPN #	Description
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

1.2.7 Validated and Supported 100GB/s Cables

Table 9 - Validated and Supported 100GB/s Cables

Speed	Cable OPN #	Description
100GbE	CBL-00195-02	100GbE QSFP28 to QSFP28 copper cable 3M
100GB/S	MCP1600-C001	Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 1M
100GB/S	MCP1600-C002	Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 2M
100GB/S	MCP1600-C003	Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 3M
100GB/S	MCP1600-C00A	Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 0.5M
100GE	MCP1600-C01A	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 1.5m
100GE	MCP1600-C02A	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 2.5m
100GB/S	MFA1A00-C005	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 5m
100GB/S	MFA1A00-C010	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 10m
100GB/S	MFA1A00-C015	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 15m
100GB/S	MFA1A00-C020	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 20m
100GB/S	MFA1A00-C030	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 30m
100GB/S	MFA1A00-C050	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 50m
100GB/S	MFA1A00-C100	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 100m
100GB/S	MFS1200-C005	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 5m
100GB/S	MFS1200-C010	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 10m
100GB/S	MFS1200-C015	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 15m
100GB/S	MFS1200-C020	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 20m
100GB/S	MFS1200-C030	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 30m
100GB/S	MFS1200-C050	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 50m
100GB/S	MFS1200-C100	Mellanox® Active Fiber Cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 100m
100GB/S	MMS1C00-C500	Mellanox® transceiver, 100GbE, QSFP28, MPO, 1550nm PSM4, up to 2km
100GB/S	MMA1B00-C100D	Mellanox® Transceiver, 100GbE, QSFP28, MPO, 850nm, up to 100m
100GB/S	MMS1C00-C500	Mellanox® transceiver, 100GbE, QSFP28, MPO, 1550nm PSM4, up to 2km

1.3 Tested Switches

1.3.1 Tested 10/40GbE Switches

Table 10 - Tested 10/40GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
10/40GbE	N/A	3064	48-port 10Gb/40Gb Switch	Cisco
10GbE	N/A	5548	Cisco 10GB ETH switch	Cisco
40GbE	N/A	3132Q	Cisco 40GB ETH switch	Cisco
10/40GbE	N/A	7050Q	16-port 40Gb Switch	Arista
40GbE	N/A	7050QX	32-port 40Gb Switch	Arista
10/40GbE	N/A	7050S	48-port 10Gb/40Gb Switch	Arista
10GbE	N/A	G8264	BNT 10/40GB ETH switch	BNT
40GbE	N/A	G8316	BNT 40GB RackSwitch G8316	BNT
10GbE	N/A	QFX3500	Juniper 10/40GB ETH switch	Juniper
10GbE	N/A	S4810P-AC	48-port 10Gb/40Gb Switch	Force10
40GbE	N/A	S6000	32-port 40Gb Switch	Dell
10GbE	SwitchX®	SX1016X-1BFR	64-Port 10GbE Switch System	Mellanox
40GbE	SwitchX®	SX1036B-1BFR	36-Port 40/56GbE Switch System	Mellanox

1.3.2 Tested 100GbE Switches

Table 11 - Tested 100GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
100Gb/s	N/A	7060CX	32-port 100Gb Switch	Arista
100Gb/s	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
100Gb/s	N/A	C3232C	High-Density, 100 Gigabit Ethernet Switch	Cisco
100Gb/s	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

1.4 Tools, Switch Firmware and Driver Software

Firmware Rev 14.20.1030 is tested with the following tools, Switch firmware, and driver software:

Table 12 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_OFED	4.1-1.0.2.0/4.0-1.0.1.0
MLNX_EN (MLNX_OFED based code)	4.1-1.0.2.0/4.0-1.0.1.0
WinOF-2	1.70/1.60
MFT	4.7.0/4.6.0
WMIware	<ul style="list-style-type: none"> ESXi 6.5 v4.16.10.3 ESXi 6.0 v4.15.10.3 ESXi 5.5 v4.5.10.3
MLNX-OS	<ul style="list-style-type: none"> SwitchX: 3.6.3004 Spectrum: 3.6.3004
SwitchX®/SwitchX®-2 Firmware	9.4.2160
Spectrum™ Firmware	13.1130.0130
Linux Inbox Drivers	<ul style="list-style-type: none"> 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 (support may vary between cards, see [Section 1.1, “Supported Devices”, on page 6](#)).

Firmware Rev 14.20.1030 supports the following FlexBoot:

Table 13 - Supported FlexBoot, UEFI

Expansion ROM	Supported Version
FlexBoot	3.5.210
UEFI	14.12.24

1.6 Revision Compatibility

Firmware Rev 14.20.1030 complies with the following programmer's reference manual:

- *Mellanox Adapters Programmer's Reference Manual (PRM), Rev 0.44 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 14.20.1030

Table 14 - Changes and New Features in Rev 14.20.1030

Feature/Change	Description
Virtualization	[Beta] Set the number of VF's value to 8 and enabled SR-IOV for VPI adapter cards.
Bug Fixes	See Section 4, “Bug Fixes History” , on page 23

3 Known Issues

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

For a list of old firmware Know Issues, please see ConnectX4Lx Firmware Archived Known Issues file (http://www.mellanox.com/pdf/firmware/ConnectX4Lx-Firmware_Archived_Known_Issues_v1.0.pdf)

Table 15 - Known Issues (Sheet 1 of 7)

Internal Ref.	Issue
1063904	Description: Messages with mkey signature on offset > 4GB are not supported.
	Workaround: N/A
	Keywords: Signature retransmission
	Discovered in Version: 14.20.1010
1063148	Description: Pause duration: Physical port counters count in 512bits quantas, instead of microseconds.
	Workaround: To normalize the counter, do not change the speed: $counter_value_in_microsec = current_counter_value * 512 / port_speed$
	Keywords: Pause duration, Physical port counters
	Discovered in Version: 14.20.1010
1059782	Description: If BMC does not support the recovery flow, no communication will occur after fwre-set.
	Workaround: N/A
	Keywords: BMC
	Discovered in Version: 14.20.1010
1054335/1054671	Description: When using UD RoCE multicast traffic over SR-IOV, packets are scattered to all the attached QPs in the e-sw (PF and its VFs) and not only on the vport that is specified in the e-se FDB.
	Workaround: N/A
	Keywords: UD RoCE multicast traffic, SR-IOV
	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
1046427/1047180	Description: ECN does not function as expected when the number of QPs > ~500 per host.
	Workaround: N/A
	Keywords: ECN
	Discovered in Version: 14.20.1010

Table 15 - Known Issues (Sheet 2 of 7)

Internal Ref.	Issue
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
743242	Description: SR-IOV min & max rate limiter can only support up to 64 VFs per port.
	Workaround: N/A
	Keywords: SR-IOV min & max rate limiter
	Discovered in Version: 14.20.1010
979364	Description: Changing SL2VL (QTCT commands in ETH or SL2VL mad in IB) during traffic may cause the chip to hang.
	Workaround: Run SL2VL commands before running traffic.
	Keywords: SL2VL, traffic
	Discovered in Version: 14.18.1000
964783	Description: Querying Vport/eSwitch that are not set to FOLLOW using the max_tx_speed command, returns information as if the FOLLOW mode is enabled.
	Workaround: N/A
	Keywords: max_tx_speed, Vport/eSwitch
963540	Description: Enabling the s-vlan strip on a vport for which the user configured an s-vlan match on its Flow Steering tables, results in the corruption of the steering on that specific vport.
	Workaround: N/A
	Keywords: s-vlan strip, vport, Flow Steering
963473	Description: When running min_avg_bw and max_avg_bw together, and when configuring a high proportion for min_avg_bw between vports (for example: 1:40, 1:100), the vport with the lowest proportion will get high deviation.
	Workaround: Set one TC not to be mapped by any user priority. For example (TC7 is not mapped): <code>mlnx_qos -i <network interface> -p 0,1,2,3,4,5,6,6</code>
	Keywords: Performance

Table 15 - Known Issues (Sheet 3 of 7)

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_element</code> command has been issued by the driver, with <code>max_average_bw != 0</code> .
	Workaround: Set a default rate (<code>modify_scheduling_element.max_average_bw=0</code>), or destroy all the scheduling elements on the chip prior to issuing a Teardown/FLR
	Keywords: Teardown/FLR, Max Rate Limiter
955595	Description: Under the <code>DUP_MAC_ACTION==LAST_CFG</code> configuration (default), the first duplicated MAC address in the MPFS is prioritized instead of the last address.
	Workaround: N/A
	Keywords: MAC address, MPFS
955061	Description: Occasionally, when the link is up at a speed of 1GbE, data traffic will not go through.
	Workaround: N/A
	Keywords: Link speed, 1GbE
949485	Description: TX doorbell via UAR and CQ doorbell via UAR are currently not supported in Multi-Host devices.
	Workaround: N/A
	Keywords: TX doorbell, CQ doorbell, Multi-Host, UAR
946800	Description: PXE booting in RedHat 7.3 is currently not supported.
	Workaround: N/A
	Keywords: PXE, RedHat 7.3
941203	Description: Occasionally, mapping 2 SLs to a single VL results in bad results in BW allocation for both SLs.
	Workaround: N/A
	Keywords: SLs to VL mapping, BW allocation
938322	Description: Performance issues occur when running <code>min_avg_bw</code> and <code>max_avg_bw</code> together. The issue starts when configuring high proportion for <code>min_avg_bw</code> between vports. For example: 1:40, 1:100: the vport with the low proportion will get high deviation.
	Workaround: N/A
	Keywords: Performance
935581	Description: When SR-IOV is enabled, some multicast traffic might be lost if another vport that a listening on the same multicast GID is down.
	Workaround: N/A
	Keywords: Multicast traffic, vport
911628	Description: Host rate limiter values are statically configured and do not change when changing the port speed.
	Workaround: N/A
	Keywords: Rate limiter

Table 15 - Known Issues (Sheet 4 of 7)

Internal Ref.	Issue
898603	Description: If multiple processes in RX RDMA Flow Table are used, vport counters may be counted more than once.
	Workaround: N/A
	Keywords: vport counters
877646	Description: In IB virtualization, transmitted vport counter cannot count traffic between functions on the same phy port.
	Workaround: N/A
	Keywords: vport counter
864200	Description: Running the <code>modify_scheduling_context</code> command does not include checking whether the scheduling element was created or not.
	Workaround: Do not modify non-existing elements
	Keywords: SR-IOV Rate Limiter
854805/864202	Description: Setting/modifying the <code>max_average_bw</code> rate for a function, or setting speeds over the maximum supported speed (as indicated in INI) may result in inaccurate rates, and in an assert.
	Workaround: Set the <code>max_avergae_bw</code> in <code>scheduling_context</code> commands to equal or less than the supported wire speed.
	Keywords: Bandwidth rate, speed
854206/856355/ 954512	Description: Host Shaper rate limiter values (<code>static_host_rate_limiter</code> , <code>port_speed</code> , <code>pci_speed</code>) are reported according to their minimum values. Thus, it is not recommended to change the <code>static_host_rate_limiter</code> value when modifying the port speed.
	Workaround: N/A
	Keywords: Host Shaper rate limiter
827444	Description: FDR link can raise with symbol errors on optic EDR cable longer than 30M.
	Workaround: N/A
	Keywords: FDR link, EDR cable
824525	Description: The first duplicated MAC address in the MPFS is prioritized (instead of the last address) under the <code>DUP_MAC_ACTION==LAST_CFG</code> configuration (default).
	Workaround: N/A
	Keywords: Duplicated MAC address
783742	Description: In order to raise 50GbE link when using ConnectX-4 Lx firmware v14.16.1006 or newer, the following conditions must be met: <ul style="list-style-type: none"> • The minimum ConnectX-4 firmware version should be 12.16.1006 • The minimum ConnectX-4 Lx firmware version should be 14.16.1006 • The minimum MLNX-OS version should be 3.6.1000 (firmware v13.1100.0026)
	Workaround: N/A
	Keywords: MLNX-OS, 50G link

Table 15 - Known Issues (Sheet 5 of 7)

Internal Ref.	Issue
776830/778257	Description: Performing warm reboot during firmware image burning in VPI/IB devices configured with IB port protocol, might cause the device to disappear from the PCIe.
	Workaround: Power Cycle the server (cold reboot). Once a cold reboot is performed, the device will reboot with the previous image that was already burned.
	Keywords: Warm reboot, firmware image burning, VPI/IB devices
770824	Description: Pressing the Power Down button resets the server and does not initiate the Standby flow (as init 0 does). As a result, both ports are up due to <code>keep_link_up</code> , which opens the port when the firmware is loaded.
	Workaround: Use <code>init 0</code> to start the Standby flow.
	Keywords: Warm/cold reboot
758803	Description: The firmware and the hardware do not reset the physical link upon <code>CPortState=down</code> . According to the IB Specification, MANAGEMENT STATE CHANGE COMMANDS: “ <i>CPortState... when phy_link=up and CPortState=down, the state machine will transition to the Link-Down state which will reset other link state machines. Since phy_link=up, this will be followed by a transition to the LinkInitialize state. Thus a command to change link port state to down provides a way to re-initialize the link layer...</i> ”
	Workaround: In order to re-train the physical link, <code>sendbug PortInfo.physical_port_state = POLLING</code> is required.
	Keywords: Physical link, <code>CPortState=down</code>
756872/769604/ 850198	Description: Flow Counter is supported only for FTE that does not include a <code>flow_tag</code> or for FTE that have TIR as a destination.
	Workaround: N/A
	Keywords: Flow Counter, FTE
756871/770208/ 850199	Description: Using Flow Counters in the FDB Flow Table causes the transmitted IB traffic vport counters not to function properly.
	Workaround: N/A
	Keywords: Flow Counter, FDB Flow Table, vport counters
756870/769605/ 850199/850208	Description: Using Flow Counters in the FDB Flow Table may harm vport counters’ clearing functionality.
	Workaround: N/A
	Keywords: Flow Counter, FDB Flow Table, vport counters
754914	Description: When e-switch FDB is not created, the VF functional loopback traffic is sent to vport 0 (PF).
	Workaround: N/A
	Keywords: e-switch FDB, vport, SR-IOV
748292	Description: When a steering rule in the e-sw FDB includes an <code>encap</code> action and an external port as destination, a transmitted multicast packet that matches the rule is sent to the wire and the loopback and the locally looped back packet will also have an <code>encap</code> header.
	Workaround: N/A
	Keywords: FDB, multicast packet

Table 15 - Known Issues (Sheet 6 of 7)

Internal Ref.	Issue
748199	Description: In case of a steering rule in the e-sw FDB with encap action and an external port as destination, the transmitted multicast packet that matches this rule is sent to the wire and the loop-back. Additionally, the locally looped back packet has an encap header as well.
	Workaround: N/A
	Keywords: FDB multicast local loopback packet
747967	Description: Burning firmware on the same device in parallel from multiple interfaces (e.g. PCIe and MTUSB) is not supported.
	Workaround: N/A
	Keywords: PCIe, MTUSB, burning in parallel
693832	Description: In virtualized (SR-IOV/Multi-Host/Socket Direct) setups OpenSM should be configured as follow (opensm.conf): <ul style="list-style-type: none"> • virt_enable should be 2 • [Recommended] Enable Qos: qos TRUE
	Workaround: N/A
	Keywords: SR-IOV/Multi-Host/Socket Direct, OpenSM
691754	Description: end_padding_mode is required in CREATE_QP and not in INIT_2_RTR command as defined in the PRM
	Workaround: N/A
	Keywords: end_padding_mode, PRM
691490	Description: LR4 cable events are sent although the port is up
	Workaround: N/A
	Keywords: Management
691387/691415	Description: In a Multi-Host/Socket Direct setup, when running a single TCP stream, you might experience sub optimal throughput.
	Workaround: Use multiple streams to reach optimal results
	Keywords: Multi-Host/Socket Direct setup, Performance, TCP stream
690890	Description: Updating a non-volatile configuration of port type TLV more than 50 times might cause system to hang.
	Workaround: Run <code>mlxconfig reset</code> after every 50 consecutive updates of port type TLV.
	Keywords: Non-volatile configuration, TLV
682518	Description: Interoperability issue between ConnectX-4 or ConnectX-4 Lx adapter cards and ConnectX-2 adapter card when trying to raise a 10GbE link.
	Workaround: N/A
	Keywords: Interoperability
648914	/651063/1066193 Description: Some 10GbE cables are not SFF-8472 compliant. "SFP+ Cable Technology" bits are cleared.
	Workaround: N/A
	Keywords: 10GbE cables, SFP+

Table 15 - Known Issues (Sheet 7 of 7)

Internal Ref.	Issue
600534	Description: Configuration of space power management capability PME_EN cannot be set, thus preventing the driver from activating the wake signal.
	Workaround: N/A
	Keywords: PCIe
599810/601485	Description: mlxfwreset does not function properly in old MFT versions after upgrading the firmware image.
	Workaround: Upgrade MFT to the latest release or use <code>reboot/power cycle</code> after upgrading firmware.
	Keywords: Firmware Tool
594964	Description: A minimum of 200 LFM is required in order to cool the MCX4411A-ACAN adapter card.
	Workaround: N/A
	Keywords: Temperature

4 Bug Fixes History

Table 16 lists the bugs fixed in this release.

Table 16 - Bug Fixes History (Sheet 1 of 14)

Internal Ref.	Issue
1090723	Description: Fixed an issue that wrongly reported the maximum temperature in a setup as the current temperature regardless of the actual temperature.
	Keywords: Temperature
	Discovered in Release: 14.20.1010
	Fixed in Release: 14.20.1030
1090566	Description: Fixed an issue that caused the adapter cards not to send an event after a temperature recovery incident after running PLDM thermal.
	Keywords: PLDM thermal, temperature
	Discovered in Release: 14.20.1010
	Fixed in Release: 14.20.1030
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
1013035	Description: Removed the following false alarm asserts reported in Multi-Host environment: <ul style="list-style-type: none"> • 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

Table 16 - Bug Fixes History (Sheet 2 of 14)

Internal Ref.	Issue
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
1063449	Description: Fixed an issue that caused TX to get stuck when a link fail-over occurred in LAG and the firmware switched between the two ports. Additional credits reset flow were added when the firmware moved between different port,vl.
	Keywords: TX, LAG
	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 16 - Bug Fixes History (Sheet 3 of 14)

Internal Ref.	Issue
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 PortRcvDataVLExtended/PortXmitDataVLExtended 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
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

Table 16 - Bug Fixes History (Sheet 4 of 14)

Internal Ref.	Issue
902828/ 915047	Description: When using a firmware based LLDP/DCBX software based, LLDP tools (such as lldp-tool in Linux) should be disabled. When intending to use software based LLDP, firmware LLDP must be disabled by using mlxconfig. Using both the LLDP software and the firmware based LLDP will result in an unexpected results. This applies to both Physical Functions (Bare Metal OS) and Virtual Functions.
	Keywords: LLDP/DCBX
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.18.2000
966472	Description: Fixed an issue which caused bi-directional traffic 10% BW degradation in Multi-Host.
	Keywords: Performance
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.18.2000
959369	Description: Increased the CQE zipping aggressive mode timer to 9000.
	Keywords: Performance
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.18.2000
962901	Description: Moving IPoIB enhanced QP to ERR or RST state results in the corruption of the service_type and pm_state in the QPC.
	Keywords: IPoIB enhanced QP
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.18.2000
961194	Description: Attaching RoCE IPv4 QPs to MCG when the vport state is set to toggle (DOWN/UP), prevents the QPs that are listed on that MCG from receiving any traffic.
	Keywords: RoCE IPv4 QPs
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.18.2000
655688	Description: When arming SRQ for limit event, the device might issue an event with context_index=0 .
	Keywords: RoCE
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.18.2000
949458	Description: Occasionally, when moving UD QP from error state to RTS, the QP re-enters the error state.
	Keywords: UD QP, Error state, RTS
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.18.2000

Table 16 - Bug Fixes History (Sheet 5 of 14)

Internal Ref.	Issue
928872	Description: When performing Pkey check for IPoIB enhanced traffic, the Pkey membership bit is ignored.
	Keywords: Pkeys
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.18.2000
862480	Description: Stopping the Rate Limiter while traffic is being transmitted might cause the adapter card to hang.
	Keywords: Rate Limiter
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.18.2000
597718	Description: Privileged Vport egress traffic is not blocked when Vport is not active
	Keywords: Virtualization
	Discovered in Release: 14.12.1100
	Fixed in Release: 14.18.2000
959527	Description: Missing invalidation upon Set () .pkey leads to bad Pkey checks.
	Keywords: Pkeys, PortInfo.LID
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.18.2000
919526	Description: Fixed an issue which caused the HCA mad response to contain the incoming packet Pkey and not the matched Pkey.
	Keywords: Pkey
	Discovered in Release: 14.17.2020
	Fixed in Release: 14.18.1000
963653/ 961833	Description: Diagnostic counters are not reset when enabled with on_demand mode.
	Keywords: on_demand mode, Diagnostic counters
	Discovered in Release: 14.18.1000
	Fixed in Release: 14.1000
920552	Description: Modified PCIe Tx configuration.
	Keywords: PCIe TX
	Discovered in Release: 14.17.2020
	Fixed in Release: 14.18.1000
943484	Description: Fixed an issue that prevented the software to set ECN parameters (min_rate, max_rate, rate_to_set_on_first_cnp) to values >32768.
	Keywords: RoCE Lossy, ECN
	Discovered in Release: 14.17.2020
	Fixed in Release: 14.18.1000

Table 16 - Bug Fixes History (Sheet 6 of 14)

Internal Ref.	Issue
876275	Description: Fixed an issue which caused the link speed to raise as DDR when connected with certain copper cables to devices supporting up to QDR speed.
	Keywords: DDR, QDR
	Discovered in Release: 14.17.2020
	Fixed in Release: 14.18.1000
886357	Description: Fixed an issue which prevented physical counters from resetting. Now the physical counters are reset on first driver start.
	Keywords: Physical counters
	Discovered in Release: 14.17.2020
	Fixed in Release: 14.18.1000
	Description: Fixed possible negotiation issues with 3rd parties.
	Keywords: Link negotiation
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.18.1000
827444	Description: Fixed a rare issue which caused 56GbE link to raise with errors.
	Keywords: Link speed
	Discovered in Release: 14.16.1020
	Fixed in Release: 14.18.1000
867367/ 867787	Description: Fixed an issue which caused <code>scheduling_context.element_type</code> to be taken into consideration with performing verifications, when running the <code>modify_scheduling_context</code> command, although the field is reserved.
	Keywords: SR-IOV Rate Limiter
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.18.1000
865373/ 865820	Description: Fixed an issue which caused the eSwitch <code>max_average_bw</code> ref counter to decrement in <code>TEARDOWN_HCA/ FLR VF</code> regardless of the <code>max_average_bw</code> value set, although the ref counter design was to increment on every <code>max_average_bw != 0</code> (limited).
	Keywords: Bandwidth rate, VFs, <code>TEARDOWN_HCA/ FLR VF</code>
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.18.1000
866181	Description: Fixed an issue which caused system fail when enabled SR-IOV.
	Keywords: SR-IOV
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.17.2020

Table 16 - Bug Fixes History (Sheet 7 of 14)

Internal Ref.	Issue
883830	Description: Fixed the NC-SI "set" Flow Control so it could always report "unsupported command" when operating the smbus.
	Keywords: NC-SI "set" Flow Control
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.17.2020
855533	Description: Fixed ODP flow issues that caused occasional fatal error reporting and RX hanging.
	Keywords: ODP flow
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.17.2020
860574/ 860716	Description: Fixed performance issues to improve Packet Pacing performance.
	Keywords: Performance, Packet Pacing
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.17.2020
883834	Description: Information is now received per channel port instead of the lowest port which received LLDP.
	Keywords: LLDP
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.17.2020
883834	Description: Changed the maximum TLV size per each TLV to 28B instead of 32B.
	Keywords: TLV size
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.17.2020
883834	Description: Fixed an issue which added a header to the packet which already contained header and data.
	Keywords: Packet headers
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.17.2020
883834	Description: Enabled LLDP by default in the ini after adding the MCTP command.
	Keywords: LLDP, MCTP command
	Discovered in Release: 14.17.1010
	Fixed in Release: 14.17.2020
826702	Description: Fixed an issue which caused RX to hang when a UDP packet with destination port of RoCE v2 arrived and the data matched the DC transport service.
	Keywords: RoCE v2, DC Transport, UDP, RX
	Discovered in Release: 14.16.1010
	Fixed in Release: 14.17.1010

Table 16 - Bug Fixes History (Sheet 8 of 14)

Internal Ref.	Issue
802148	Description: Fixed an issue which caused the link not to come up when the port was toggled in a rapid frequency.
	Keywords: Link up
	Discovered in Release: 14.16.1010
	Fixed in Release: 14.17.1010
736528	Description: On rare occasions during UEFI boot cycles system got stuck while WinPE is loaded. (OS WinPE, system DL160).
	Keywords: WinPE OS, UEFI boot cycles
	Discovered in Release: 14.16.1010
	Fixed in Release: 14.17.1010
689471	Description: Single FTE that catches both untagged and prio-tagged packets (by giving an FTE with match_value.vlan_tag = 0 and match_value.vid = 0) is currently not supported.
	Keywords: Ethernet Network
	Discovered in Release: 14.14.1006
	Fixed in Release: 14.17.1010
-	Description: Flashing the firmware requires server reboot. Firmware cannot be flashed twice without server reboot after first flashing
	Keywords: Upgrading/Downgrading
	Discovered in Release: 14.14.1006
	Fixed in Release: 14.17.1010
828608	Description: Fixed an issue causing bubbles to appear as symbol errors when link raised FDR 1x.
	Keywords: FDR 1x
	Discovered in Release: 14.16.1010
	Fixed in Release: 14.17.1010
677359	Description: When Clause 74 Fire-Code FEC is active, and there are FC corrected errors, both the FC_correctable counter and the FC_uncorrectable counter are increment.
	Keywords: Clause 74 Fire-Code FEC, Counters
	Discovered in Release: 14.16.1010
	Fixed in Release: 14.17.1010
827579/ 826702	Description: Fixed an issue which caused RX to hang when the UDP packet had a reserved UDP destination port.
	Key Words: UDP packet, RX
	Discovered in Release: 14.14.1006
	Fixed in Release: 14.17.1010
835735	Description: Fixed DMAC reporting mapping per host.
	Key Words: DMAC reporting
	Discovered in Release: 14.14.1006
	Fixed in Release: 14.17.1010

Table 16 - Bug Fixes History (Sheet 9 of 14)

Internal Ref.	Issue
846520	Description: Fixed an EEH error from PCI which caused firmware to hang.
	Key Words: EEH error
	Discovered in Release: 14.14.1006
	Fixed in Release: 14.17.1010
857644	Description: Fixed the default value of the PCIe target_link_speed to Gen3 in link control2.
	Key Words: PCIe, Link Speed
	Discovered in Release: 14.14.1006
	Fixed in Release: 14.17.1010
783733/ 774373	Description: Fixed an issue which prevented LEDs from blinking when the traffic was less than 0.1% of the link speed.
	Key Words: LEDs, Blink, Link Speed
	Discovered in Release: 14.14.1006
	Fixed in Release: 14.17.1010
781039	Description: A server getting into a Standby mode while Packet-Pacing is enabled might cause firmware to hang and driver call-trace.
	Keywords: Packet-Pacing
	Discovered in Release: 14.14.1020
	Fixed in Release: 1417.1010
753349	Description: Fixed an issue which caused unexpected QoS functionality in case of multiple sources to single destination traffic transmission.
	Key Words: QoS
	Discovered in Release: 14.14.1020
	Fixed in Release: 1417.1010
801374	Description: Fixed an issue which occasionally caused the RX traffic to hang in DC when received a PCI error on WQE fetch.
	Key Words: RX traffic, DC
	Discovered in Release: 14.14.1020
	Fixed in Release: 1417.1010
773110	Description: Fixed OOB connection issue during Intel's ITP inject errors test.
	Keywords: OOB, ITP inject errors test
	Discovered in Release: 14.14.2036
	Fixed in Release: 1417.1010
773724	Description: Modified PCIe settings.
	Key Words: PCIe
	Discovered in Release: 14.14.1006
	Fixed in Release: 14.16.1020

Table 16 - Bug Fixes History (Sheet 10 of 14)

Internal Ref.	Issue
670185	Description: Added protection from IOPX thermal diode destabilization to prevent UEFI IPv6 PXE boot failure on ConnectX-4 Lx 25GE OCP card.
	Key Words: UEFI, OCP card
	Discovered in Release: 14.14.2036
	Fixed in Release: 14.16.1006
780651	Description: Fixed an issue causing single port devices to query and write Physical Port TLVs to Port 2.
	Keywords: Physical Port TLVs, single port device
	Discovered in Release: 14.14.2036
	Fixed in Release: 14.16.1006
752392	Description: Enabled mlxfwreset to work using the PCIe Secondary Bus Reset.
	Keywords: mlxfwreset
	Discovered in Release: 14.14.2036
	Fixed in Release: 14.16.1006
775393	Description: Fixed an issue causing link flapping as a result, incorrect link settings.
	Keywords: Link flapping, link settings
	Discovered in Release: 14.14.2036
	Fixed in Release: 14.16.1006
756570	Description: Fixed an issue causing wrong alignment markers to be used when running 50G with Clause91 FEC enabled.
	Keywords: Clause91 FEC, 50G
	Discovered in Release: 14.14.2036
	Fixed in Release: 14.16.1006
767281	Description: Reduced the default BAR size for VF (SR-IOV) from 5 (32 MB) to 1 (2MB).
	Keywords: BAR size for VF (SR-IOV)
	Discovered in Release: 14.14.2036
	Fixed in Release: 14.16.1006
735159/ 747595/ 752533	Description: Added legacy interrupts support in FlexBoot.
	Keywords: Interrupts, PXE
	Discovered in Release: 14.14.2036
	Fixed in Release: 14.16.1006
752343	Description: Modified the TX configuration to support EMI crossing margins in 16Ghz
	Keywords: TX configuration, EMI
	Discovered in Release: 14.14.2036
	Fixed in Release: 14.16.1006

Table 16 - Bug Fixes History (Sheet 11 of 14)

Internal Ref.	Issue
691194	Description: In some cases, a Bit Error Rate is not optimal on 10G/40G links.
	Keywords: 10G/40G links, Bit Error Rate
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036
689788	Description: Instability of Link Training Flow occurs during 100G Auto-Negotiation.
	Keywords: Link Training Flow, Auto-Negotiation
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036
684496	Description: Fixed a rare issue which caused the command to hang when moved the QP to RESET and back to RTS.
	Keywords: QP, RTS
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036
665089	Description: Improved RDMA READ bandwidth under packet lost scenario.
	Keywords: RDMA READ bandwidth
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036
736195	Description: Added support for pnat = 1 in HCA <code>access_reg</code> command as required by the <code>ibdiag-net</code> tool.
	Keywords: <code>access_reg</code> command, <code>ibdiagnet</code>
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036
696486	Description: Increased the steering hash tables static size from 128 to a maximum of 32K entries.
	Keywords: Steering hash tables static size
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036
691649	Description: Prevented miscalculation of module temperature when using 100Gb/s cables (OPN: MFA1A00-Cxxx for 100GbE).
	Keywords: 100Gb/s cables
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036
693446	Description: Reduced one hop for Unicast RX steering, steering pipes balancing.
	Keywords: Ethernet Steering performance
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036

Table 16 - Bug Fixes History (Sheet 12 of 14)

Internal Ref.	Issue
690614	Description: Non-volatile configuration of Port Type TLV more than 50 times might cause system hang.
	Keywords: Non-volatile configuration, Port Type TLV
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036
691043	Description: Enabled RoCE IPv4 Multicast prevents MCG command from failing when an IPv4 is mapped to an IPv6 address.
	Keywords: RoCE IPv4 Multicast
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036
649696/ 690681	Description: If the PF driver or the tool (e.g. ethtool) use PAOS DOWN command (e.g. by ifconfig down or ip link set down), loopback traffic is blocked for all functions on this port (PF->VFs / VF->VF) In Multi-Host loopback, the traffic will be blocked once the firmware receives the PAOS down command from all PFs. However, the loopback traffic will not be blocked when the port is down due to the physical link (for example: cable plugged out, switch port down).
	Keywords: Multi-Host loopback
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036
659307	Description: Fixed a 25G and 50G link issue when Clause 91 RS FEC was active.
	Keywords: 25G and 50G link, Clause 91 RS FEC
	Discovered in Release: 14.12.1100
	Fixed in Release: 14.14.1100
676877	Description: Added a missing invalidation of eSwitch cache upon FLR which caused the upcoming driver load to either fail or not to be able to transmit.
	Keywords: Packet Transmit, FLR
	Discovered in Release: 14.12.1240
	Fixed in Release: 14.14.1100
670185	Description: Fixed a UEFI IPv6 PXE boot failure on ConnectX-4 Lx 25GE OCP card.
	Keywords: UEFI, OCP card
	Discovered in Release: 14.12.1240
	Fixed in Release: 14.14.1100
668221	Description: Fixed an issue which prevented Vport counters from counting local loopback packets. Packets now are filter by the self-loopback prevention.
	Keywords: Vport, local loopback packets
	Discovered in Release: 14.12.1240
	Fixed in Release: 14.14.1100

Table 16 - Bug Fixes History (Sheet 13 of 14)

Internal Ref.	Issue
667288	Description: Reported INTx as unsupported to allow PFs Passthrough on PowerKVM.
	Keywords: Passthrough, PowerKVM
	Discovered in Release: 14.12.1100
	Fixed in Release: 14.14.1100
596637	Description: SR-IOV Ethernet supports up to 18 VFs per port only.
	Keywords: Virtualization
	Discovered in Release: 14.12.1240
	Fixed in Release: 14.14.1100
591240	Description: Fixed an incident what allowed local (internal) loopbacked packets to be counted by the Vport counters, although Vport counters should count only traffic that crosses the Vport.
	Keywords: Virtualization
	Discovered in Release: 14.12.1240
	Fixed in Release: 14.14.1100
664558	Description: Fixed an issue preventing driver loading or TX traffic sending upon reboot, after ungraceful driver unload.
	Keywords: Driver Load
	Discovered in Release: 14.12.1240
	Fixed in Release: 14.14.1100
657680	Description: Fixed casting of BMC MAC before steering API.
	Keywords: BMC, Steering API
	Discovered in Release: 14.12.1240
	Fixed in Release: 14.14.1100
614403	Description: Fixed the PCI write flow to take into consideration the PCI MTU. This fix eliminates the need for NOPs in the flow, which resulted from PPC larger PCI MTU. The single queue limitation for READ is due to a hardware limitation of the number of READ request in a given time.
	Keywords: PCI MTU
	Discovered in Release: 14.12.1100
	Fixed in Release: 14.14.1100
630327	Description: Fixed a case that caused FlexBoot to not work as expected with systems that run with "large bar" enabled (Above 4G Decoding) over Connect-IB or ConnectX-4 HCAs.
	Keywords: FlexBoot, 4G Decoding
	Discovered in Release: 14.12.1100
	Fixed in Release: 14.14.1100

Table 16 - Bug Fixes History (Sheet 14 of 14)

Internal Ref.	Issue
629563	Description: Fixed an over-subscription on the RX buffer when Flow Control was not enabled which caused the RX pipe to hang.
	Keywords: Flow Control
	Discovered in Release: 14.12.0780
	Fixed in Release: 14.12.1240

5 Firmware Changes and New Feature History

Table 17 - Firmware Changes and New Feature History (Sheet 1 of 7)

Feature/Change	Description
DSCP	Added trust level for QoS prioritization according to the DSCP or PCP.
	[Beta] Added ingress buffer management for: <ul style="list-style-type: none"> 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.
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 24, “Per host Settings,” on page 54)
Fast Teardown	Enables fast unloading driver by using Teardown HCA with op_mode=1 (panic mode). For further information, refer to the PRM.
SFP Power Flow Improvement (level 2,1)	Added support for SFP power class.
10GBaseT module	Added support for 10GBaseT modules.
Bug Fixes	See Section 4, “Bug Fixes History,” on page 23
Rev. 14.18.2000	
Bug Fixes	See Section 4, “Bug Fixes History,” on page 23
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.

Table 17 - Firmware Changes and New Feature History (Sheet 2 of 7)

Feature/Change	Description
No Driver NIC (NODNIC) Performance Improvement	Improved performance of: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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.
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 23
Rev. 14.17.2020	
GENEVE & IP-in-IP Stateless Off-load	[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 23
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: <ul style="list-style-type: none"> • 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.

Table 17 - Firmware Changes and New Feature History (Sheet 3 of 7)

Feature/Change	Description
Multi-Host L3/L4 Classification	Enables load balancing in the Multi PF Switch layer (MPFS) based on the L3/L4 headers
95 Virtual Functions (VF) per Port	<p>Increased the number of VFs from 64 to 95 per Physical Function (PF).</p> <p>Note: When increasing the number of VFs, the following limitations must be taken into consideration:</p> <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> <p>Note: For the maximum number of VFs supported by your driver, please refer to your drivers' Release Notes or User Manual.</p>
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.
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)	<p>Added a new physical layer statistics counters group. The new group includes BER counters, FEC error correction, clear time, and additional physical layer counters.</p> <p>For further information, please refer to the Ethernet Adapters Programming Manual (PRM).</p>
Permanent Link Up Mode	<p>Enables the user to set a certain link up state for an unlimited period of time. This mode has 3 states:</p> <ul style="list-style-type: none"> 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)	<p>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.</p> <p>For further information, please refer to the PRM.</p>

Table 17 - Firmware Changes and New Feature History (Sheet 4 of 7)

Feature/Change	Description
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: <ul style="list-style-type: none"> • 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 23
Rev. 14.16.1020	
Bug Fixes	See Section 4, “Bug Fixes History” , on page 23
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.
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 (ASAP ²) 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 ASAP ² 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.

Table 17 - Firmware Changes and New Feature History (Sheet 5 of 7)

Feature/Change	Description
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
Bug Fixes	See Section 4, “Bug Fixes History”, on page 23
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 23
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 17 - Firmware Changes and New Feature History (Sheet 6 of 7)

Feature/Change	Description
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.
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 23
Rev. 14.12.1100	
Port Link	Reduced the port link-up time when negotiating according to Clause 73 (DME)
Rev. 14.12.0780	
PCI	<ul style="list-style-type: none"> • PCIe Function Level Reset (FLR) • Power Management L2/L3 flow support

Table 17 - Firmware Changes and New Feature History (Sheet 7 of 7)

Feature/Change	Description
Ethernet Network	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • Strided SRQ • Self Loopback support • Transport Domain support • CQ2EQ remapping • Added support for the following commands: <ul style="list-style-type: none"> • MODIFY/QUERY_ESW_VPORT_CONTEXT • QUERY/MODIFY_CONG_STATUS • QUERY/MODIFY_CONG_PARAMS • QUERY_CONG_STATISTICS • ADD/DELETE_VXLAN_UDP_DPORT
Virtualization	<ul style="list-style-type: none"> • VXLAN/NVGRE Stateless offload In this release, this feature is supported through Windows ONLY • SR-IOV EN (at Beta level)
Performance	<ul style="list-style-type: none"> • CQE zipping
Misc	<ul style="list-style-type: none"> • Wake-on-Lane/Standby • FlexBoot/UEFI support
Non-Volatile Configuration	<ul style="list-style-type: none"> • Non-Volatile Configuration (NVConfig). For the complete list, please refer to Section 8, on page 54.

6 FlexBoot Changes and New Features

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

Table 18 - FlexBoot Changes and New Features (Sheet 1 of 2)

Version	Description
Rev. 3.5.210	
Promiscuous VLAN mode	Added support for promiscuous VLAN mode.
MTU	[InfiniBand] Added support for configurable MTU.
Expansion ROM version	Enabled expansion ROM (<code>exp_rom</code>) version exposition according to the new specification (e.g. expose ARCH in flint tool).
FlexBoot UI	Added a FlexBoot menu support for <code>NV_POWER_CONF</code> . Now power consumption configuration is supported from the FlexBoot menu.
	Enhanced FlexBoot/firmware debug capability using Flexboot UI. Added the <code>reg_dump</code> option to the <code>panic_behavior</code> configuration in the FlexBoot menu
Upstream sync	Synced the source with iPXE (upstream sync)
Rev. 3.5.110	
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 <code>PCI - Bus:Dev.Func</code> to <code>0000:Bus:Dev.Func</code>
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

Table 18 - FlexBoot Changes and New Features (Sheet 2 of 2)

Version	Description
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
SHELL CLI	Removed support for the following SHELL CLI commands: <ul style="list-style-type: none"> • 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 19 - FlexBoot Known Issues

Internal Ref.	Description
-	<p>Description: Several BIOS vendors have limited boot-vector space and may not display FlexBoot in their boot menu.</p> <p>Workaround: Disable the embedded NIC boot agent in BIOS</p> <p>Keywords: BIOS</p>
-	<p>Description: In several BIOS, the server might hang during FlexBoot booting due to wrong configuration of the PMM.</p> <p>Workaround: N/A</p> <p>Keywords: BIOS</p>
-	<p>Description: Only EBX, ESI, DS, ES registers can be saved in Boot Entry.</p> <p>Workaround: N/A</p> <p>Keywords: BIOS</p>
-	<p>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.</p> <p>Workaround: Follow the instructions described in the FlexBoot UM for the proper iSCSI boot/install</p> <p>Keywords: BIOS</p>
673114/821899	<p>Description: FlexBoot banner might not be shown in some BIOSes.</p> <p>Workaround: N/A</p> <p>Keywords: BIOS</p>
-	<p>Description: In some cases, PXE boot will not work if the client was given only the file-name without next-server (siaddr).</p> <p>Workaround: N/A</p> <p>Keywords: PXE Boot</p>
-	<p>Description: PXE boot after iSCSI boot with static configuration is currently not supported.</p> <p>Workaround: N/A</p> <p>Keywords: PXE Boot</p>
-	<p>Description: Boot over VLAN with IB port is currently not supported.</p> <p>Workaround: N/A</p> <p>Keywords: PXE Boot</p>

Table 19 - FlexBoot Known Issues

Internal Ref.	Description
-	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 19 - FlexBoot Known Issues

Internal Ref.	Description
-	<p>Description: Boot menu is displayed as READ ONLY if the HCA card does not support flash configuration.</p> <p>Workaround: N/A</p> <p>Keywords: User Interface</p>
-	<p>Description: FlexBoot Boot Menu will not be visible in serial output.</p> <p>Workaround: N/A</p> <p>Keywords: User Interface</p>
-	<p>Description: Large Receive Offload (LRO) and iSCSI may not interoperate due to a bug in current Linux kernel distributions.</p> <p>Workaround: Disable LRO in the IPoIB module when using iSCSI. See the Mellanox FlexBoot user's manual for details under the Diskless Machines chapter (Infini-Band Ports).</p> <p>Keywords: Networking</p>
-	<p>Description: 56Gb/s is currently not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Link Speed</p>
-	<p>Description: Setting the number of Virtual Functions higher than the machine's memory capability may cause memory issues and system instability.</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>
-	<p>Description: SLAM, FTP, HTTPS and SRP are currently not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Protocols</p>
-	<p>Description: Occasionally, using the Spanning Tree Protocol (STP) in the switches may cause packet drops and boot failure in the system.</p> <p>Workaround: Enable the "edgemode" if disabled on the switch, or use either portfast or edgemode functionality on the switch ports connected to the NICs.</p> <p>Keywords: Protocols</p>
-	<p>Description: FCoE, BCV are not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Protocols</p>

Table 19 - FlexBoot Known Issues

Internal Ref.	Description
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
841198	Description: FlexBoot fails to boot when the following occurs: <ul style="list-style-type: none"> • 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
843377/849223	Description: The physical MAC assigned via the boot menu is displayed as zeroes instead of the set MAC when ConnectX-4 VPI adapter card is configured as InfiniBand.
	Workaround: N/A
	Keywords: Physical MAC, Boot menu
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
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

Table 19 - FlexBoot Known Issues

Internal Ref.	Description
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
689460	Description: To use the DHCP server to identify ipxe requests when using <code>undionly.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
928217	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 boot-loader, 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,
976878	Description: When using boot-loader grub2 to boot WDS, if the WDS boot fails, an RSOD might appear.
	Workaround: N/A
	Keywords: Boot-loader grub2, WDS, RSOD
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 a continues looping.
	Workaround: N/A
	Keywords: Boot-loader grub2, WDS, RSOD

6.2 FlexBoot Bug Fixes History

Table 20 - FlexBoot Bug Fixes History (Sheet 1 of 2)

Version	Issue
843209	Description: Fixed and issue which cause 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: 3.5.109
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
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
697291	Description: In ConnectX-4, the PXE boot time measurement over TFTP Ethernet is 1:30 min for image size of 1GB, TFTP InfiniBand is 1:20 min, and iSCSI boot time measurement is 8 seconds for image size of 25 MB.
	Keywords: PXE Boot
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.812

Table 20 - FlexBoot Bug Fixes History (Sheet 2 of 2)

Version	Issue
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 Unsupported Features and Commands

7.1 Unsupported Features

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

- Service types not supported:
 - SyncUMR
 - Mellanox transport
 - PTP
 - RAW IPv6
 - PTP (IEEE 1588)
- INT-A not supported for EQs only MSI-X
- PCI VPD write flow (RO flow supported)
- Streaming Receive Queue (STRQ) and collapsed CQ
- Precise clock synchronization over the network (IEEE 1588)
- SM is not supported on VFs
- DC is not supported in: Multi-Host, SR-IOV, and Ethernet (RoCE)
- RoCE LAG for VFs and Multi-Host/Socket-Direct are not supported in RoCE LAG
- QoS per VFs feature is supported up to 14 VFs per PF in dual port device with 8 VLs.
- Mutlihost Ethernet

7.2 Unsupported Commands

- QUERY_MAD_DEMUX
- SET_MAD_DEMUX
- PAGE_FAULT_RESUME
- ACTIVATE_TRACER
- DEACTIVATE_TRACER
- ACCESS_REG_SPACE
- ACCESS_REG_SPACE_DWORD
- ACTIVATE/DEACTIVATE_TRACER
- QUERY/MODIFY_SCHED_QUEUE
- CREATE_RQ - MEMORY_RQ_RMP
- MODIFY_LAG_ASYNC_EVENT

8 Supported Non-Volatile Configurations

Table 21 - 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 22 - 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 23 - Per host/function Settings

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

Table 24 - Per host Settings

Name	Parameter Index
NV_PCI_CONF	0x80
NV_PCI_CAP	0x81