Mellanox ConnectX®-4 Firmware Release Notes

Rev 12.20.1010

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

Table of Contents

List of Tables

Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 12.20.1010	August 24, 2017	Updated section Section 4, "Bug Fixes History", on page 28: Removed old Bug Fixes and added reference to the ConnectX4-Firmware_Archived_Bug_Fixes_v1.0.pdf file.
	August 14, 2017	Updated section Section 4, "Bug Fixes History", on page 28: added issue # 1086254
	July 30, 2017	Updated section Section 4, "Bug Fixes History", on page 28: added issue # 1060650
	July 17, 2017	Updated the following sections: • Section 2, "Changes and New Features in Rev 12.20.1010", on page 19 • Table 1.2.7, "Validated and Supported EDR/100GB/s Cables," on page 13
	July 03, 2017	Initial version of this firmware release.

1   Overview

These are the release notes for the ConnectX®-4 adapters firmware Rev 12.20.1010. This firmware supports the following protocols:

- InfiniBand - SDR, QDR, FDR10, FDR, EDR

- Ethernet - 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE[1] and 100GigE

- 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 - Supported Devices (Sheet 1 of 2)

Device Part Number	PSID	Device Name	Compiled with FlexBoot	Compiled with UEFI[a]
MCX413A-BCAT	MT_2120110027	ConnectX®-4 EN network interface card, 40GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	Yes	No
MCX413A-GCAT	MT_2600110035	ConnectX®-4 EN network interface card, 50GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	Yes	No
MCX414A-BCAT	MT_2130110027	ConnectX®-4 EN network interface card, 40GbE dual-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	Yes	No
MCX414A-GCAT	MT_2610110035	ConnectX®-4 EN network interface card, 50GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	Yes	No
MCX415A-BCAT	MT_2120111027	ConnectX®-4 EN network interface card, 40GbE single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	Yes	No
MCX415A-CCAT	MT_2140110033	ConnectX®-4 EN network interface card, 100GbE single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	Yes	No
MCX415A-GCAT	MT_2120110035	ConnectX®-4 EN network interface card; 50GbE single-port QSFP28; PCIe3.0 x16; ROHS R6	Yes	No
MCX416A-BCAT	MT_2130111027	ConnectX®-4 EN network interface card, 40GbE dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	Yes	No

---

1. 56 GbE is a Mellanox propriety link speed and can be achieved while connecting a Mellanox adapter cards to Mellanox SX10XX switch series or connecting a Mellanox adapter card to another Mellanox adapter card.

Table 2 - Supported Devices (Sheet 2 of 2)

Device Part Number	PSID	Device Name	Compiled with FlexBoot	Compiled with UEFI[a]
MCX416A-CCAT	MT_2150110033	ConnectX®-4 EN network interface card, 100GbE dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R	Yes	No
MCX416A-GCAT	MT_2130110035	ConnectX®-4 EN network interface card; 50GbE dual-port QSFP28; PCIe3.0 x16; ROHS R6	Yes	No
MCX445B-CCAN	MT_0000000016	ConnectX®-4 EN network interface card for OCP, 100GbE single-port QSFP28, PCIe3.0 x16, no bracket, ROHS R6	Yes	No
MCX445B-ECAN	MT_0000000018	ConnectX®-4 VPI network interface card for OCP, EDR IB (100Gb/s) and 100GbE single-port QSFP28, PCIe3.0 x16, no bracket, ROHS R6	Yes	No
MCX556A-ECAT	MT_0000000008	ConnectX®-5 VPI adapter card, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	No	No
MCX453A-FCAT	MT_2160110021	ConnectX®-4 VPI adapter card, FDR IB 40GbE, single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	Yes	No
MCX454A-FCAT	MT_2170110021	ConnectX®-4 VPI adapter card, FDR IB and 40GbE, dual-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	Yes	No
MCX455A-ECAT	MT_2180110032	ConnectX®-4 VPI adapter card, EDR IB (100Gb/s) and 100GbE, single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	Yes	No
MCX455A-FCAT	MT_2160111021	ConnectX®-4 VPI adapter card, FDR IB and 40GbE, single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	Yes	No
MCX456A-FCAT	MT_2170111021	ConnectX®-4 VPI adapter card, FDR IB and 40GbE, dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	Yes	No
MCX456A-ECAT	MT_2190110032	ConnectX®-4 VPI adapter card, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	Yes	No

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® module, ETH 1GbE, 1Gb/s, SFP, Base-T, up to 100m

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

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
40GB/S	MC2210511-LR4	Mellanox® optical module, IB FDR10, 40Gb/s, QSFP, LC-LC, 1310nm, LR4 up to 10km
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
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

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
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
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

Table 5 - Validated and Supported 25GbE Cables

Speed	Cable OPN #	Description
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

Table 8 - Validated and Supported FDR Cables

Speed	Cable OPN #	Description
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

1.2.7 Validated and Supported EDR/100GB/s Cables

Table 9 - Validated and Supported EDR/100GB/s Cables

Speed	Cable OPN #	Description
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
EDR	MCP1600-E001[a]	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 1M
EDR	MCP1600-E002[a]	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 2M
EDR	MCP1600-E003	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 3M
EDR	MCP1600-E00A[a]	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 0.5M
EDR	MCP1600-E01A[a]	Mellanox® Passive Copper cable, VPI, up to 100Gb/s, QSFP, LSZH, 1.5m

Table 9 - Validated and Supported EDR/100GB/s Cables

Speed	Cable OPN #	Description
EDR	MCP1600-E02A	Mellanox® Passive Copper cable, VPI, up to 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
EDR	MFA1A00-E005[a]	MELLANOX Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 5m
EDR	MFA1A00-E010[a]	MELLANOX Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 10m
EDR	MFA1A00-E015[a]	MELLANOX Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 15m
EDR	MFA1A00-E020	MELLANOX Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 20m
EDR	MFA1A00-E030	MELLANOX Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 30m
EDR	MFA1A00-E050	MELLANOX Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 50m
EDR	MFA1A00-E100	MELLANOX Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 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

Table 9 - Validated and Supported EDR/100GB/s Cables

Speed	Cable OPN #	Description
100GB/S	MMS1C00-C500	Mellanox® transceiver, 100GbE, QSFP28, MPO, 1550nm PSM4, up to 2km
EDR	MFS1200-E005	Mellanox® Active Fiber Cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 5m
EDR	MFS1200-E010	Mellanox® Active Fiber Cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 10m
EDR	MFS1200-E015	Mellanox® Active Fiber Cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 15m
EDR	MFS1200-E020	Mellanox® Active Fiber Cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 20m
EDR	MFS1200-E030	Mellanox® Active Fiber Cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 30m
EDR	MFS1200-E050	Mellanox® Active Fiber Cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 50m
EDR	MFS1200-E100	Mellanox® Active Fiber Cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 100m
100GB/S	MMA1B00-C100D	Mellanox® Transceiver, 100GbE, QSFP28, MPO, 850nm, up to 100m
EDR	MMA1B00-E100	Mellanox® Transceiver, IB EDR, up to 100Gb/s, QSFP28, MPO, 850nm, up to 100m
100GB/S	QSFP-100G-AOC10M	Cisco 100GBase QSFP Active Optical Cable, 10-meter
100GB/S	MMS1C00-C500	Mellanox® transceiver, 100GbE, QSFP28, MPO, 1550nm PSM4, up to 2km

a. Forward Error Correction (FEC) is deactivated on this cable.

1.3 Tested Switches

1.3.1 Tested QDR Switches

Table 10 - Tested QDR Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
QDR	N/A	12300	36-Port 40Gb QDR Infiniband Switch, Management Module, Dual Power	QLogic
QDR	InfiniScale® IV	IS5025Q-1SFC	36-port 40Gb/s InfiniBand Switch Systems	Mellanox
QDR	InfiniScale® IV	Switch 4036	Grid Director™ 4036E	Mellanox

1.3.2 Tested 10/40GbE Switches

Table 11 - 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

1.3.3 Tested FDR Switches

Table 12 - Tested FDR Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
FDR	SwitchX®	SX6018F-1SFR	18-port 56Gb/s InfiniBand/VPI Switch Systems	Mellanox
FDR	SwitchX®	SX6036F-1BFR	36-port 56Gb/s InfiniBand/VPI Switch Systems	Mellanox
FDR	SwitchX®	SX6506	108-Port 56Gb/s InfiniBand Director Switch	Mellanox
FDR	SwitchX®-2	SX6710-FB2F2	36-port 56Gb/s InfiniBand/VPI Switch Systems	Mellanox

1.3.4 Tested 100GbE/EDR Switches

Table 13 - Tested 100GbE/EDR Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
100Gb/s	N/A	7060CX	32-port 100Gb Switch	Arista

Table 13 - Tested 100GbE/EDR Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
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
EDR	Switch-IB	SB7790-EB2F	36-port EDR 100Gb/s InfiniBand Switch Systems	Mellanox
EDR	Switch-IB 2	SB7800-ES2R	36-port Non-blocking Managed EDR 100Gb/s InfiniBand Smart Switch	Mellanox
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 12.20.1010 is tested with the following tools, Switch firmware, and driver software:

Table 14 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_OFED	4.1-1.0.0.0/4.0-1.0.1.0
MLNX_EN (MLNX_OFED based code)	4.1-1.0.0.0/4.0-1.0.1.0
WinOF-2	1.70/1.60
MFT	4.7.0/4.6.0
WMware	• ESXi 6.5 v4.16.10.3 • ESXi 6.0 v4.15.10.3 • ESXi 5.5 v4.5.10.3
MLNX-OS	• SwitchX: 3.6.3004 • Switch-IB: 3.6.3004 • Switch-IB 2: 3.6.3004 • Spectrum: 3.6.3004
SwitchX®/SwitchX®-2 Firmware	9.4.2160
Spectrum™ Firmware	13.1130.0130
SwitchX-IB™ Firmware	11.1300.0126
SwitchX-IB 2 Firmware	15.1300.0126
InfiniScale® V Firmware	7.4.3000/v7.4.2200

Table 14 - Tools, Switch Firmware and Driver Software

	Supported Version
Linux Inbox Drivers	• Ubuntu 14.04.3 • Ubuntu 14.04.4 • Ubuntu 15.04 • Ubuntu 15.10 • Ubuntu 16.04 • Ubuntu 16.04.1 • Ubuntu 16.10 • SLES12 • SLES12.1 • SLES12.2 • RHEL6.6 • RHEL6.7 • RHEL6.8 • RHEL7.1 • RHEL7.2 • RHEL7.3
Windows Inbox Drivers	Windows Server 2016

1.5 Supported FlexBoot

> 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 12.20.1010 supports the following FlexBoot:

Table 15 - Supported FlexBoot

Expansion ROM	Supported Version
FlexBoot	3.5.210

1.6 Revision Compatibility

Firmware Rev 12.20.1010 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 12.20.1010

Table 16 - Changes and New Features in Rev 12.20.1010

Feature/Change	Description
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.
Multi-Host/Socket Direct Routing to be LID based	**[InfiniBand only]** Changed the Multi-Host/Socket Direct routing to be LID based instead of GID based. Thus, GRH/GID index is not required. **Note:** This feature requires SM 4.8.1 and above.
Relaxed Ordering	**[Beta]** Added support for relaxed ordering write in memory keys.
RDMA Counters	Enhanced RDMA counter
TLV for PCI class code	Added 2 new per Host TLVs (see Table 26, "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.
IPoIB Virtualization	Added support for enhanced IPoIB (QP.ulp == 2) in virtualized system (SR-IOV / Multi-Host / Socket Direct)
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 28

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 ConnectX4 Firmware Archived Known Issues file (http://www.mellanox.com/pdf/firmware/ConnectX4-Firmware_Archived_Known_Is-sues_v1.0.pdf)

Table 17 - Known Issues (Sheet 1 of 8)

Internal Ref.	Issue
1063904	**Description:** Messages with mkey signature on offset > 4GB are not supported.
	Workaround: N/A
	Keywords: Signature retransmission
	Discovered in Version: 12.20.1010
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: 12.20.1010
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: 12.20.1010
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: 12.20.1010
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: 12.20.1010
1031744	**Description:** Same flow counter cannot be used on different table types.
	Workaround: N/A
	Keywords: Flow counter
	Discovered in Version: 12.20.1010

Table 17 - Known Issues (Sheet 2 of 8)

Internal Ref.	Issue
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: 12.20.1010
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: 12.20.1010
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: 12.18.1000
966364	**Description: [Ethernet]** TX queue rate limit may sometimes exceed the rate that was set by the user by up to 10%.
	Workaround: N/A
	Keywords: TX queue rate limit
	Discovered in Version: 12.20.1010
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): `mlnx_qos -i <network interface> -p 0,1,2,3,4,5,6,6`
	Keywords: Performance

Table 17 - Known Issues (Sheet 3 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 `max_rate_limiter` enabled Teardown/FLR, the hardware remains enabled (rate_limiter_en = 1). ** "`max rate limiter enabled`" = at least 1 (per chip). `create/modify_sched-uling_elemnt` command has been issued by the driver, with `max_average_bw != 0`.
	Workaround: Set a default rate (`modify_schduling_element.max_aver-age_bw=0`), 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 `DUP_MAC_ACTION==LAST_CFG` 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 `min_avg_bw` and `max_avg_bw` together. The issue starts when configuring high proportion for `min_avg_bw` 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

Table 17 - Known Issues (Sheet 4 of 8)

Internal Ref.	Issue
929267	**Description:** Copper cables 3m and above are not supported vs. SX1024 switch system.
	Workaround: N/A
	Keywords: Cables
911628	**Description:** Host rate limiter values are statically configured and do not change when changing the port speed.
	Workaround: N/A
	Keywords: Rate limiter
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
877646	**Description:** The e-sw uplink state can affect the traffic only if the driver determines a root Flow Table for the e-sw FDB.
	Workaround: N/A
	Keywords: FDB, e-sw uplink state
864200	**Description:** Running the `modify_scheduling_context` 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 `max_average_bw` 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 `max_avergae_bw` in `scheduling_context` commands to equal or less than the supported wire speed.
	Keywords: Bandwidth rate, speed
854206/856355	**Description:** If the vport state is DOWN and a packet is sent in the local loopback, the sx_sniffer tool will not function.
	Workaround: N/A
	Keywords: sx_sniffer, vport
827444	**Description:** FDR link can raise with symbol errors on optic EDR cable longer than 30M.
	Workaround: N/A
	Keywords: FDR link, EDR cable

Table 17 - Known Issues (Sheet 5 of 8)

Internal Ref.	Issue
824525	**Description:** The first duplicated MAC address in the MPFS is prioritized (instead of the last address) under the DUP_MAC_ACTION==LAST_CFG configuration (default).
	Workaround: N/A
	Keywords: Duplicated MAC address
783742	**Description:** In order to raise 50GbE link when using ConnectX-4 firmware v12.16.1006 or newer, the following conditions must be met: • 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
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: Warn 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 `keep_link_up`, which opens the port when the firmware is loaded.
	Workaround: Use init 0 to start the Standby flow.
	Keywords: Warm/cold reboot
758803	**Description:** The firmware and the hardware do not reset the physical link upon `CPort-State=down`. According to the IB Specification, MANAGEMENT STATE CHANGE COMMANDS: *"CPortState… when phy_link=up and CPortState=down, the state machine will transition to the LinkDown 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…"*
	Workaround: In order to re-train the physical link, sendbug PortInfo.physical_port_state = POLLING is required.
	Keywords: Physical link, CPortState=down
756872/769604/850198	**Description:** Flow Counter is supported only for FTE that does not include a flow_tag 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

Table 17 - Known Issues (Sheet 6 of 8)

Internal Ref.	Issue
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 send 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 encap 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 encap header.
	Workaround: N/A
	Keywords: FDB, multicast 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
747961	**Description:** Force loopback in the QPC in virtualized environment (Multi-Host or SR-IOV) is not supported.
	Workaround: N/A
	Keywords: Force loopback, Multi-Host, SR-IOV
693832	**Description:** In an InfiniBand SR-IOV setups, traffic should contain GRH (GID index) if the grh_required bit is set in the query_hca_vport_context command. **Note:** In this case, traffic without GRH will be forwarded to vport0 ("Host0")
	Workaround: N/A
	Keywords: SR-IOV setups, GRH
693832	**Description:** In virtualized (SR-IOV/Multi-Host/Socket Direct) setups OpenSM should be configured as follow (opensm.conf): • 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

Table 17 - Known Issues (Sheet 7 of 8)

Internal Ref.	Issue
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-voltile configuration of port type TLV more than 50 times might cause system to hang.
	Workaround: Run `mlxconfig reset` after every 50 consecutive updates of port type TLV.
	Keywords: Non-voltile configuration, TLV
686032	**Description:** While transmitting a packet from a NODNIC interface to BMC (on the same phy port) the packet will be duplicated and will be transmitted to the wire too (in addition to the packet that arrived to BMC).
	Workaround: N/A
	Keywords: NODNIC interface, BMC
685062	**Description:** Multi-Host InfiniBand: OpenSM is supported over host0 only and the `MAD_IFC` usage is limited to host0 only.
	Workaround: Activate OpenSM and the MFT tools via host0
	Keywords: Multi-Host InfiniBand
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+
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

Table 17 - Known Issues (Sheet 8 of 8)

Internal Ref.	Issue
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 `reboot/power cycle` after upgrading firmware.
	Keywords: Firmware Tool
572150	**Description:** A low link speed issue occurs when connecting a ConnectX®-4 EDR adapter card with a QDR InfiniScale® IV based switch using a copper cable longer than 3M. The link is raised as DDR.
	Workaround: N/A
	Keywords: Link Speed, QDR, DDR

4 Bug Fixes History

Table 18 lists the bugs fixed in this release. For a list of old firmware Bug Fixes, please see ConnectX4 Firmware Archived Bug Fixes file (http://www.mellanox.com/pdf/firmware/ConnectX4-Firmware_Archived_Bug_Fixes_v1.0.pdf)

Table 18 - Bug Fixes History (Sheet 1 of 7)

Internal Ref.	Issue
1060650	**Description:** Fixed a link issue on Intel 10GbE Optical module PN: R8H2F, Y3KJN.
	Keywords: Intel 10GbE Optical module
	Discovered in Release: 12.18.2000
	Fixed in Release: 12.20.1010
1040099	**Description:** Fixed an issue that caused the link to raise as DDR instead of QDR after firmware reset when connected to switch 4036.
	Keywords: QDR, DDR
	Discovered in Release: 12.18.1000
	Fixed in Release: 12.20.1010
1052064	**Description:** Fixed an issue that caused the device to hang upon warm reboot.
	Keywords: Warm reboot
	Discovered in Release: 12.18.2000
	Fixed in Release: 12.20.1010
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: 12.18.2000
	Fixed in Release: 12.20.1010
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: 12.18.2000
	Fixed in Release: 12.20.1010
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: 12.18.2000
	Fixed in Release: 12.20.1010

Table 18 - Bug Fixes History (Sheet 2 of 7)

Internal Ref.	Issue
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: 12.18.2000
	Fixed in Release: 12.20.1010
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: 12.18.2000
	Fixed in Release: 12.20.1010
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: 12.18.2000
	Fixed in Release: 12.20.1010
1047533	**Description:** Rephrased and improved external troubleshoot messages in PDDR register.
	Keywords: PDDR register
	Discovered in Release: 12.18.1000
	Fixed in Release: 12.20.1010
999261	**Description:** Improved SR-IOV performance.
	Keywords: SR-IOV
	Discovered in Release: 12.18.1000
	Fixed in Release: 12.20.1010
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: 12.18.1000
	Fixed in Release: 12.20.1010
1002884	**Description:** Fixed an issue that prevented ibdump from functioning properly on Connect-X-4 second port.
	Keywords: ibdump
	Discovered in Release: 12.18.1000
	Fixed in Release: 12.20.1010

Table 18 - Bug Fixes History (Sheet 3 of 7)

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: 12.18.2000
	Fixed in Release: 12.20.1010
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: 12.18.2000
	Fixed in Release: 12.20.1010
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: 12.18.1000
	Fixed in Release: 12.20.1010
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: 12.18.2000
	Fixed in Release: 12.20.1010
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: 12.18.1000
	Fixed in Release: 12.20.1010
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: 12.18.1000
	Fixed in Release: 12.20.1010
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: 12.18.1000
	Fixed in Release: 12.20.1010

Table 18 - Bug Fixes History (Sheet 4 of 7)

Internal Ref.	Issue
902828/ 915047	**Description:** When using a firmware based LLDP/DCBX software based, LLDP tools (such as lldptool 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: 12.18.1000
	Fixed in Release: 12.18.2000
759571/ 759655	**Description:** Modifying the encap_id of FTE is not supported.
	Keywords: encap_id, FTE
	Discovered in Release: 12.16.1020
	Fixed in Release: 12.18.2000
966472	**Description:** Fixed an issue which caused bi-directional traffic 10% BW degradation in Multi-Host.
	Keywords: Performance
	Discovered in Release: 12.18.1000
	Fixed in Release: 12.18.2000
959369	**Description:** Increased the CQE zipping aggressive mode timer to 9000.
	Keywords: Performance
	Discovered in Release: 12.18.1000
	Fixed in Release: 12.18.2000
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: 12.18.1000
	Fixed in Release: 12.18.2000
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: 12.18.1000
	Fixed in Release: 12.18.2000

Table 18 - Bug Fixes History (Sheet 5 of 7)

Internal Ref.	Issue
655688	**Description:** When arming SRQ for limit event, the device might issue an event with `context_index=0`.
	Keywords: RoCE
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.18.2000
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: 12.18.1000
	Fixed in Release: 12.18.2000
928872	**Description:** When performing Pkey check for IPoIB enhanced traffic, the Pkey membership bit is ignored.
	Keywords: Pkeys
	Discovered in Release: 12.18.1000
	Fixed in Release: 12.18.2000
862480	**Description:** Stopping the Rate Limiter while traffic is being transmitted might cause the adapter card to hang.
	Keywords: Rate Limiter
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.18.2000
597718	**Description:** Privileged Vport egress traffic is not blocked when Vport is not active
	Keywords: Virtualization
	Discovered in Release: 12.12.1100
	Fixed in Release: 12.18.2000
-	**Description:** PF direct pass-through is not supported in InfiniBand (since PF FLR is not supported)
	Keywords: PF direct pass-through, InfiniBand
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.18.2000
959527	**Description:** Missing invalidation upon `Set().pkey` leads to bad Pkey checks.
	Keywords: Pkeys, PortInfo.LID
	Discovered in Release: 12.18.1000
	Fixed in Release: 12.18.2000

Table 18 - Bug Fixes History (Sheet 6 of 7)

Internal Ref.	Issue
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: 12.17.2020
	Fixed in Release: 12.18.1000
963653/ 961833	**Description:** Diagnostic counters are not reset when enabled with on_demand mode.
	Keywords: on_demand mode, Diagnostic counters
	Discovered in Release: 12.18.1000
	Fixed in Release: 12.1000
920552	**Description:** Modified PCIe Tx configuration.
	Keywords: PCIe TX
	Discovered in Release: 12.17.2020
	Fixed in Release: 12.18.1000
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: 12.17.2020
	Fixed in Release: 12.18.1000
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: 12.17.2020
	Fixed in Release: 12.18.1000
886357	**Description:** Fixed an issue which prevented physical counters from reseting. Now the physical counters are reset on first driver start.
	Keywords: Physical counters
	Discovered in Release: 12.17.2020
	Fixed in Release: 12.18.1000
	Description: Fixed possible negotiation issues with 3rd parties.
	Keywords: Link negotiation
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.18.1000

Table 18 - Bug Fixes History (Sheet 7 of 7)

Internal Ref.	Issue
827444	**Description:** Fixed a rare issue which caused FDR/56GbE link to raise with errors.
	Keywords: Link speed
	Discovered in Release: 12.16.1020
	Fixed in Release: 12.18.1000
867367/ 867787	**Description:** Fixed an issue which caused `scheduling_context.element_type` to be taken into consideration with performing verifications, when running the `modify_scheduling_context` command, although the field is reserved.
	Keywords: SR-IOV Rate Limiter
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.18.1000
865373/ 865820	**Description:** Fixed an issue which caused the eSwitch `max_average_bw` ref counter to decrement in `TEARDOWN_HCA/ FLR VF` regardless of the `max_average_bw` value set, although the ref counter design was to increment on every max_average_bw != 0 (limited).
	Keywords: Bandwidth rate, VFs, TEARDOWN_HCA/ FLR VF
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.18.1000

5 Firmware Changes and New Feature History

Table 19 - Firmware Changes and New Feature History (Sheet 1 of 8)

Feature/Change	Description
Rev. 12.18.2000	
Bug Fixes	See Section 4, "Bug Fixes History", on page 28
Rev. 12.18.1000	
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.
RDMA RX Flow Table	Added the option to open a receive RDMA Flow Table and to forward RoCE traffic to some destination QP.

Table 19 - Firmware Changes and New Feature History (Sheet 2 of 8)

Feature/Change	Description
Bug Fixes	See Section 4, "Bug Fixes History", on page 28
Rev. 12.17.2020	
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 28
Rev. 12.17.1010	
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
InfiniBand Multi-Host Isolation	Enabled isolation between separate Hosts using the same HCA. All the Hosts can be rebooted, the driver can be stopped and the FLR signal can be sent independently.

95 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:
	`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)`
	Note: For the maximum number of VFs supported by your driver, please refer to your drivers' Release Notes or User Manual.

Table 19 - Firmware Changes and New Feature History (Sheet 3 of 8)

Feature/Change	Description
QoS per VFs	**[InfiniBand Only]** Added support for multiple VLs in SR-IOV/mutli-host environments. **Note:** The number of VLs can be configured by the NVCONFIG. The default VL number is 4 VLs.
InfiniBand Rate Limit per QP (static rate)	Added support for QP Rate Limit in InfiniBand.
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)	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.

Table 19 - Firmware Changes and New Feature History (Sheet 4 of 8)

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: • 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 28
Rev. 12.16.1020	
Bug Fixes	See Section 4, "Bug Fixes History", on page 28
Rev. 12.16.1006	
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:
	
	`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)`

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-plain in Mellanox ConnectX-4 / ConnectX-4 Lx NIC hardware (Mellanox Embedded Switch or eSwitch) while maintaining OVS control-plain 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 19 - Firmware Changes and New Feature History (Sheet 5 of 8)

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 28
Rev. 12.14.2036	
IPoIB checksum and LSO offload	Added IPoIB checksum and LSO offload support
Scatter FCS in RQ	Enables software to scatter or strip FCS in RQ.
Bug Fixes	See Section 4, "Bug Fixes History", on page 28
Rev. 12.14.1100	
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 and 56GbE Link Speed	ConnectX-4adapters now support 1Gb/s and 56GbE Ethernet connectivity in addition to 10GigE, 25GigE, 40GigE, 50GigE, and 100GigE

Table 19 - Firmware Changes and New Feature History (Sheet 6 of 8)

Feature/Change	Description
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.
Multi-Host InfiniBand	Enables connecting multiple compute or storage hosts into a single interconnect adapter by separating the adapter PCIe interface into multiple and independent PCIe interfaces.
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. 12.12.1240	
Bug Fixes	See Section 4, "Bug Fixes History", on page 28
Rev. 12.12.1100	
Port Link	Reduced the port link-up time when negotiating according to Clause 73 (DME)
Rev. 12.12.0780	
PCI	• PCIe Function Level Reset (FLR) • Power Management L2/L3 flow support

Table 19 - Firmware Changes and New Feature History (Sheet 7 of 8)

Feature/Change	Description
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	- 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
InfiniBand Network	• Dynamically Connected (DC) transport
Misc	- Wake-on-Lane/Standby - FlexBoot/UEFI support
Non-Volatile Configuration	• Non-Volatile Configuration (NVConfig). For the complete list, lease refer to Section 8, on page 54.
Port management	• Enabled port management. Now one port can be set as Ethernet and one as InfiniBand.
Rev. 12.1100.6630	
Virtualization	- Added support for SR-IOV - Added support for MADs Virtualization Attributes according to ib_-virt_annex_v17
PRM	• Updated virtualization command set according to PRM 0.26
Configuration tools	• Enabled SR-IOV, NUM_VFS and INT_LOG_MAX_PAYLOAD_-SIZE configuration via the mlxconfig tool
Rev. 12.0100.6440	
All	• Initial Release of ConnectX®-4 adapter cards

Table 19 - Firmware Changes and New Feature History (Sheet 8 of 8)

Feature/Change	Description
Port Speed	• InfiniBand port speed up to EDR • Ethernet port speed up to 100GigE
Virtualization	• Function per port

InfiniBand Network	• Dynamically Connected transport
	• Unreliable Datagram Connection transport
	• Atomic Operation
	• CORE-Direct®
	  • Provides Collective Off-loading in HCA
	  • Frees CPU to perform computation in parallel with collective operations
	• T10 DIF pipeline Data Integrity Signature off-loading (at beta level)
	• User Memory Registration (UMR)
	• Automatic Path Migration
	• On Demand Paging (ODP) - Memory can now be used without pinning memory beforehand.
	• Congestion Control
	• Shrink Address Vectors for RC and UD
	• Programmable Port/Node GUID

Ethernet Network	**Note:** All the Ethernet features listed below are at Beta level. • Large Receive Offload (LRO) • Large Send Offload (LSO) • Receive Side Scaling (RSS) • Global Pause • RoCEv1/RoCEv2. RoCE is supported only in Reliable Connection (RC) transport • Flow Steering
General	• Thermal monitoring and protection • Port LEDs indications • NVConfig Tool • Suspend to RAM (S3) support • Diagnostic counters vendor-specific MAD support, as defined by VS-MAD spec version 1.2 • Physical Port Counter - Beta level • Q Counter - Beta level • Firmware burning (using mstflint) when the driver is down • CPLD field upgrade • V Port commands
Host management	• NC-SI over RMII support
MAD	• Config space address in MAD management class 0x09

6   FlexBoot Changes and New Features

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

Table 20 - 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 (`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
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 `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.

Table 20 - FlexBoot Changes and New Features (Sheet 2 of 2)

Version	Description
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
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 21 - FlexBoot Known Issues

Internal Ref.	Description
-	**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
-	**Description:** Boot over VLAN with IB port is currently not supported.
	Workaround: N/A
	Keywords: PXE Boot

Table 21 - 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
670421	**Description:** Using filename for PXE boot with rootpath for hooking an iSCSI target (to install) is not supported when the PXE boot loader uses UNDI API, since all traffic must get to the boot loader.
	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

Table 21 - FlexBoot Known Issues

Internal Ref.	Description
655800	**Description:** iSCSI over IPv6 is not supported.
	Workaround: N/A
	Keywords: iSCSI
-	**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:** Large Receive Offload (LRO) and iSCSI may not interoperate due to a bug in current Linux kernel distributions.
	Workaround: Disable LRO in the IPoIB module when using iSCSI. See the Mellanox FlexBoot user's manual for details under the Diskless Machines chapter (InfiniBand Ports).
	Keywords: Networking
-	**Description:** 56Gb/s is currently not supported.
	Workaround: N/A
	Keywords: Link Speed
-	**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

Table 21 - FlexBoot Known Issues

Internal Ref.	Description
-	**Description:** FCoE, BCV are not supported.
	Workaround: N/A
	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
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
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

Table 21 - FlexBoot Known Issues

Internal Ref.	Description
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
689460	**Description:** To use the DHCP server to identify ipxe requests when using `undi-only.kpxe` or `ipxe.pxe` 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:00:00:00:00:00:02:00:00:02: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 bootloader, 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 21 - FlexBoot Known Issues

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
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: Bootloader grub2, WDS, RSOD

6.2   FlexBoot Bug Fixes History

Table 22 - 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 22 - 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 feature 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 23 - 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 24 - 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 25 - Per host/function Settings

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

Table 26 - Per host Settings

Name	Parameter Index
NV_PCI_CONF	0x80
NV_PCI_CAP	0x81