



Mellanox ConnectX[®]-5 Firmware Release Notes

Rev 16.19.1200

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	7
1.2.2 Validated and Supported 10/40GbE Cables	7
1.2.3 Validated and Supported QDR/FDR10 Cables	9
1.2.4 Validated and Supported 50Gbs Cables	10
1.2.5 Validated and Supported FDR Cables	10
1.2.6 Validated and Supported EDR/100GB/s Cables	11
1.3 Tested Switches	13
1.3.1 Tested QDR Switches	13
1.3.2 Tested 10/40GbE Switches	13
1.3.3 Tested FDR Switches	14
1.3.4 Tested 100GbE/EDR Switches	14
1.4 Tools, Switch Firmware and Driver Software	14
1.5 Supported FlexBoot	15
1.6 Revision Compatibility	15
Chapter 2 Changes and New Features in Rev 16.19.1200	16
Chapter 3 Known Issues	17
Chapter 4 Bug Fixes History	26
Chapter 5 Firmware Changes and New Feature History	29
Chapter 6 FlexBoot Changes and New Features	36
6.1 FlexBoot Known Issues	37
Chapter 7 Unsupported Features and Commands	42
7.1 Unsupported Features	42
7.2 Unsupported Commands	42
Chapter 8 Supported Non-Volatile Configurations	43

List of Tables

Table 1:	Release Update History	5
Table 2:	Supported Devices	6
Table 3:	Validated and Supported 1GbE Cables	7
Table 4:	Validated and Supported 10/40GbE Cables	7
Table 5:	Validated and Supported QDR/FDR10 Cables	9
Table 6:	Validated and Supported 50Gbs Cables	10
Table 7:	Validated and Supported FDR Cables	10
Table 8:	Validated and Supported EDR/100GB/s Cables	11
Table 9:	Tested QDR Switches	13
Table 10:	Tested 10/40GbE Switches	13
Table 11:	Tested FDR Switches	14
Table 12:	Tested 100GbE/EDR Switches	14
Table 13:	Tools, Switch Firmware and Driver Software	14
Table 14:	Supported FlexBoot	15
Table 15:	Changes and New Features in Rev 16.19.1200	16
Table 16:	Known Issues	17
Table 17:	Bug Fixes History	26
Table 18:	Firmware Changes and New Feature History	29
Table 19:	FlexBoot Changes and New Features	36
Table 20:	FlexBoot Known Issues	37
Table 21:	Per-physical Port Settings	43
Table 22:	Global Settings	43
Table 23:	Per host/function Settings	43

Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 16.19.1200	May 9, 2017	Updated Table 8 , “ Validated and Supported EDR/100GB/s Cables ,” on page 11 , added the following OPNs: MMA1L10-CR, MMA1L11-CR
	April 18, 2017	Added Known Issue # 1025741 to Table 16 , “ Known Issues ,” on page 17
	March 28, 2017	Initial version of this firmware release.

1 Overview

These are the release notes for the ConnectX[®]-5adapters firmware Rev 16.19.1200. This firmware supports the following protocols:

- InfiniBand - SDR, QDR, FDR10, FDR, EDR
- Ethernet - 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, and 100GigE
- PCI Express 3.0/4.0, 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	FlexBoot	UEFI
MCX515A-CCAT	MT_0000000011	ConnectX-5 EN network interface card, 100GbE single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	Yes	No
MCX516A-CCAT	MT_0000000012	ConnectX-5 EN network interface card, 100GbE dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	Yes	No
MCX516A-CDAT	MT_0000000013	ConnectX-5 Ex EN network interface card, 100GbE dual-port QSFP28, PCIe4.0 x16, tall bracket, ROHS R6	Yes	No
MCX555A-ECAT	MT_0000000010	ConnectX-5 VPI adapter card, EDR IB (100Gb/s) and 100GbE, single-port QSFP28, PCIe3.0 x16, tall 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	Yes	No
MCX556A-EDAT	MT_0000000009	ConnectX-5 Ex VPI adapter card, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, PCIe4.0 x16, tall bracket, ROHS R6	Yes	No

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 QSFP TO SFP+ ADAPTER
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
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

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
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
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

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
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 QDR/FDR10 Cables

Table 5 - 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
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

1.2.4 Validated and Supported 50Gbs Cables

Table 6 - 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

1.2.5 Validated and Supported FDR Cables

Table 7 - 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
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

1.2.6 Validated and Supported EDR/100GB/s Cables

Table 8 - 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
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

Table 8 - Validated and Supported EDR/100GB/s Cables

Speed	Cable OPN #	Description
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	MMA1L10-CR	Mellanox® optical transceiver, 100GbE, 100Gb/s, QSFP28, LC-LC, 1310nm, LR4 up to 10km
100Gb/s	MMA1L11-CR	Mellanox® optical module, 100GbE, 100Gb/s, QSFP28, LC-LC, 1310nm, 3.5W, LR4 up to 10km
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
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

Table 8 - Validated and Supported EDR/100GB/s Cables

Speed	Cable OPN #	Description
EDR	MMA1B00-E100	Mellanox® Transceiver, IB EDR, up to 100Gb/s, QSFP28, MPO, 850nm, up to 100m
100GB/S	QSFP-100G-AOC-10M	Cisco 100GBase QSFP Active Optical Cable, 10-meter

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

1.3 Tested Switches

1.3.1 Tested QDR Switches

Table 9 - 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 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

1.3.3 Tested FDR Switches

Table 11 - 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 12 - Tested 100GbE/EDR 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
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 16.19.1200 is tested with the following tools, switch firmware, and driver software:

Table 13 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_OFED	4.0-1.0.0.0/3.4-2.0.0.0

Table 13 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_EN (MLNX_OFED based code)	4.0-1.0.0.0/3.4-2.0.0.0
WinOF-2	1.60/1.50
MFT	4.6.0/4.5.0
WMware	<ul style="list-style-type: none"> ESXi 6.5 v4.16.8.8 ESXi 6.0 v4.15.8.8 ESXi 5.5 v4.5.8.8
MLNX-OS	<ul style="list-style-type: none"> SwitchX: 3.6.2938 Switch-IB: 3.6.2938 Switch-IB 2: 3.6.2938 Spectrum: 3.6.2938
SwitchX®/SwitchX®-2 Firmware	9.4.1050/9.4.2140
Spectrum™ Firmware	13.1300.0120
SwitchX-IB™ Firmware	11.1300.0120
SwitchX-IB 2 Firmware	15.1300.0120
InfiniScale® V Firmware	7.4.3000/v7.4.2200

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 supports the following FlexBoot:

Table 14 - Supported FlexBoot

Expansion ROM	Supported Version
FlexBoot	3.5.110

1.6 Revision Compatibility

Firmware Rev 16.19.1200 complies with the following programmer’s reference manual:

- Mellanox Adapters Programmer’s Reference Manual (PRM), Rev 0.41 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 16.19.1200

Table 15 - Changes and New Features in Rev 16.19.1200

Feature/Change	Description
General	This is the first GA version of the ConnectX-5/Ex adapter cards.
Bug Fixes	See section Section 4, “Bug Fixes History” , on page 26

3 Known Issues

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

Table 16 - Known Issues (Sheet 1 of 9)

Internal Ref.	Issue
-	<p>Description: Occasionally, following an NVCONFIG or firmware update process, the IC might become unresponsive.</p> <p>Workaround: Power Cycle</p> <p>Keywords: NVCONFIG or firmware update</p>
-	<p>Description: There is no interoperability between ConnectX-5 and ConnectX-4 adapter cards when using DC transport.</p> <p>Workaround: N/A</p> <p>Keywords: DC transport, interoperability</p>
954822	<p>Description: The <code>ipoib_enhanced_offloads</code> indication in the HCA capabilities reports 0 while <code>SRIOV_EN=1</code>.</p> <p>Workaround: N/A</p> <p>Keywords: SR-IOV, IPoIB Offloads</p>
946569	<p>Description: Retransmission on a QP that has mkeys signature is not supported.</p> <p>Workaround: Set retry limit to zero for such QPs in the software.</p> <p>Keywords: mkeys signature, QP</p>
-	<p>Description: Destroying XRQ / QP during traffic is currently not supported (Killing application while running has unexpected results, and requires fwreset).</p> <p>Workaround: N/A</p> <p>Keywords: XRQ / QP</p>
-	<p>Description: RNDV transport errors, specifically retransmission is currently not supported</p> <p>Workaround: N/A</p> <p>Keywords: RNDV transport errors, retransmission</p>
955061	<p>Description: Occasionally, when the link is up at a speed of 1GbE, data traffic will not go through.</p> <p>Workaround: N/A</p> <p>Keywords: Link speed, 1GbE</p>
-	<p>Description: Setting a greater RNDV DMA LEN value than the TAG length will result in assertion.</p> <p>Workaround: N/A</p> <p>Keywords: RNDV DMA LEN value, TAG length</p>
937318	<p>Description: Setting more than 8K QPs with <code>hca_sq_owner == 1</code> connected to an RNDV XRQ is currently not supported.</p> <p>Workaround: N/A</p> <p>Keywords: QPs, RNDV XRQ</p>

Table 16 - Known Issues (Sheet 2 of 9)

Internal Ref.	Issue
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
-	Description: Bit Error Rate is not optimal on QDR links
	Workaround: N/A
	Keywords: Link Speed
-	Description: To raise links with platforms based on the following ICs, comply with the following firmware version requirements: <ul style="list-style-type: none"> • Connect-IB® - 10.10.4000 • Switch-IB™ - 11.200.120 (or MLNX-OS 3.4.3050) • Spectrum™ MLNX-OS - 3.5.1000 • ConnectX®-3 - 2.32.5100 • SwitchX® - 9.2.7300 (or MLNX-OS 3.3.5006)
	Workaround: N/A
	Keywords: Interoperability
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
-	Description: If QDR is not enabled for the switch's InfiniBand Port Speed while connected to ConnectX-3/ConnectX-3 Pro or Connect-IB® FDR adapters or to SwitchX® /SwitchX-2 FDR switches, links will rise at SDR or DDR (even if FDR is enabled)
	Workaround: Enable QDR (in addition to FDR) when connecting to peer ports running at FDR
	Keywords: Interoperability
-	Description: Qualified EDR cables currently work with EDR networks (EDR devices, Switch®-IB and ConnectX®-4) only.
	Workaround: N/A
	Keywords: Cables
-	Description: PCIe capability “Device S/N” returns false value.
	Workaround: N/A
	Keywords: PCI
-	Description: When the link is Gen2, entering or exiting L1 state may cause bad CRC or DLLP indication.
	Workaround: N/A
	Keywords: PCI

Table 16 - Known Issues (Sheet 3 of 9)

Internal Ref.	Issue
600534	Description: Configuration of space power management capability <code>PME_EN</code> cannot be set, thus preventing the driver from activating the wake signal.
	Workaround: N/A
	Keywords: PCIe
647144	Description: During server reset (not a power-cycle), a non-maskable interrupt (NMI) might occur due to an Option Card Black Box (OCBB) issue causing PCIe access.
	Workaround: N/A
	Keywords: PCI
685062	Description: OpenSM is supported over host0 only and the <code>MAD_IFC</code> usage is limited to host0 only.
	Workaround: Activate OpenSM and the MFT tools via host0
	Keywords: PCI
651991	Description: OCBB compliance with iLO versions: OCBB is not displayed in the latest iLO versions
	Workaround: N/A
	Keywords: PCI
591240	Description: Traffic that is loopbacked due to <code>QP.force_loopback</code> being equaled to 1, is steered to the PF.
	Workaround: N/A
	Keywords: Ethernet Network
601485/ 599810	Description: <code>mlxfwreset</code> 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

Table 16 - Known Issues (Sheet 4 of 9)

Internal Ref.	Issue
-	<p>Description: Older MFT versions (4.0.0 and 3.8.0) may indicate that the latest GA firmware is old or that it cannot be compared with the existing firmware. A message similar to the below will be displayed upon firmware upgrade stage:</p> <pre># flint -d <mst device> -i <image> burn</pre> <p>Current FW version on flash: 12.1100.6630 New FW version: 12.0012.0572</p> <p>Note: The new FW version is not newer than the current FW version on flash.</p> <p>Do you want to continue ? (y/n) [n] : y</p> <p>Workaround: Choose one of the options below to upgrade firmware:</p> <ul style="list-style-type: none"> • Upgrade to the latest MFT version (4.1.0) • Type "y" after the note flint provides <p>Run flint with the "-force" flag</p> <p>Keywords: Firmware Upgrade/MFT</p>
-	<p>[For customers developing custom low level drivers]</p> <p>Description: VFs internal FLR is not supported in PF teardown HCA command.</p> <p>Workaround: Before unloading the PF driver, PF driver must disable all its active VFs by performing the following:</p> <ol style="list-style-type: none"> 1. Run the <code>disable_hca</code> command on all the <code>function_ids</code> 2. Wait until firmware returns all VFs allocated pages. <p>Keywords: Virtualization, FLR</p>
-	<p>Description: Function (PF/VF) TX port counters are not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>
-	<p>Description: PF driver must work with pages event queue.</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>
591240/ 840738	<p>Description: Any local (internal) loopbacked packet is counted by the Vport counters, although Vport counters should count only traffic that crosses the Vport.</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>
597718	<p>Description: Vport number in virtual trap might be reported incorrectly</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>

Table 16 - Known Issues (Sheet 5 of 9)

Internal Ref.	Issue
648914	Description: Some 10GbE cables are not SFF-8472 compliant. "SFP+ Cable Technology" bits are cleared.
	Workaround: N/A
	Keywords: 10GbE cables, SFP+
691387/ 691415	Description: In a Socket Direct setup, when running a single TCP stream, you might experience sub optimal throughput.
	Workaround: Use multiple streams to reach optimal results
	Keywords: Socket Direct setup, Performance, TCP stream
691754	Description: <code>end_padding_mode</code> is required in <code>CREATE_QP</code> and not in <code>INIT_2_RTR</code> command as defined in the PRM
	Workaround: N/A
	Keywords: <code>end_padding_mode</code> , PRM
691490	Description: LR4 cable events are sent although the port is up
	Workaround: N/A
	Keywords: Management
-	Description: QoS must be configured the same for both ports in order for RoCE LAG to function properly.
	Workaround: N/A
	Keywords: RoCE LAG
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 loopback. Additionally, the locally looped back packet has an encap header as well.
	Workaround: N/A
	Keywords: FDB multicast local loopback packet
756872/ 769604	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	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	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

Table 16 - Known Issues (Sheet 6 of 9)

Internal Ref.	Issue
748292	Description: When a steering rule in the e-Switch 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/ 771507	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
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
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
783742	Description: In order to raise 50GbE link when using ConnectX-5 firmware v16.18.1000 or newer, the following conditions must be met: <ul style="list-style-type: none"> • The minimum ConnectX-5 firmware version should be 16.18.1000 • The minimum MLNX-OS version should be 3.6.1000 (firmware v13.1100.0026)
	Workaround: N/A
	Keywords: MLNX-OS, 50G link
776830	Description: Performing warm reboot during firmware image burning for VPI/IB devices configured with IB port protocol, might cause the device to disappear from the PCIe.
	Workaround: Cold reboot the device instead
	Keywords: Warm/cold reboot
770824	Description: Pressing the Power Down button resets the server and does not initiate the Standby flow (as <code>init 0</code> 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

Table 16 - Known Issues (Sheet 7 of 9)

Internal Ref.	Issue
693832	<p>Description: In an InfiniBand Multihost and SR-IOV setups, traffic should contain GRH (GID index) if the <code>grh_required</code> bit is set in the <code>query_hca_vport_context</code> command. OpenSM should be configured as follow (<code>opensm.conf</code>):</p> <ul style="list-style-type: none"> • <code>virt_enable</code> should be 2 • Enable Qos: • <code>qos TRUE</code> <p>Note: In this case, traffic without GRH will be forwarded to <code>vport0</code> ("Host0")</p>
	<p>Workaround: N/A</p>
	<p>Keywords: Multihost/SR-IOV setups</p>
778257	<p>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.</p>
	<p>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.</p>
	<p>Keywords: Warn reboot, firmware image burning, VPI/IB devices</p>
758803	<p>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 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..."</i></p>
	<p>Workaround: In order to re-train the physical link, <code>sendbug PortInfo.physical_port_state = POLLING</code> is required.</p>
	<p>Keywords: Physical link, CPortState=down</p>
852744	<p>Description: Mapping an SL to VL 15 is currently not supported. Trying to do so, will cause a health buffer fatal internal error report.</p>
	<p>Workaround: N/A</p>
	<p>Keywords: SL to VL mapping</p>
854805	<p>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.</p>
	<p>Workaround: Set the <code>max_avergae_bw</code> in <code>scheduling_context</code> commands to equal or less than the supported wire speed.</p>
	<p>Keywords: Bandwidth rate, speed</p>
827444	<p>Description: FDR link can raise with symbol errors on optic EDR cable longer than 30M.</p>
	<p>Workaround: N/A</p>
	<p>Keywords: FDR link, EDR cable</p>
-	<p>Description: PDDR access register reports incorrect FEC request in the Phy Info page.</p>
	<p>Workaround: N/A</p>
	<p>Keywords: PDDR access register</p>

Table 16 - Known Issues (Sheet 8 of 9)

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
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
967336/ 967397	Description: FDB Root Flow Table must be created for eSwitch vport VLAN insertion (VST) to be supported.
	Workaround: N/A
	Keywords: VST, Root Flow Table
937998/ 926739/ 979505	Description: Static Rate Limiter per QP cannot activate 2 rates simultaneously.
	Workaround: N/A
	Keywords: Static Rate Limiter, QP
979173/ 983749	Description: DC transport is currently at beta level and does not function properly.
	Workaround: N/A
	Keywords: DC transport
984178	Description: NVMeoF staging buffer must be contiguous and at least 128KB.
	Workaround: N/A
	Keywords: NVMeoF
968266	Description: NVMeoF IO transactions size supported is up to 20KB.
	Workaround: N/A
	Keywords: NVMeoF
998069/ 752194	Description: When using an Operating Systems with kernels 3.18 – 4.9 (inclusive), HCAs' passthrough can fail with the following line in trace: genirq: Flags mismatch irq 16. 00000000 (vfio-intx(0000:84:00.0)) vs. 00000000 (vfio-intx(0000:0b:00.0))
	Workaround: Use either older kernels that do not have this issue, or kernels 4.10 and up.
	Keywords: passthrough, kernels 3.18 – 4.9
983813	Description: Power cycle is required when updating from firmware v16.18.1000.
	Workaround: N/A
	Keywords: Power cycle, firmware update

Table 16 - Known Issues (Sheet 9 of 9)

Internal Ref.	Issue
1007508	<p>Description: The maximum supported Virtual Functions (VF) per port is 64 in firmware version 16.19.1200. Using more than 64 VFs, can cause lack of the DMA pages depending on the server capability.</p>
	<p>Workaround: N/A</p>
	<p>Keywords: Virtual Functions (VF)</p>
1025741	<p>Description: A Multicast Group (MCG) with QPs from mixed ULP types is not supported in firmware 16.19.1200 and onwards (QPC.ulp_stateless_offloads_mode = 0/1/2).</p>
	<p>Workaround: N/A.</p>
	<p>Keywords: Multicast Group (MCG), QPs</p>

4 Bug Fixes History

Table 17 lists the bugs fixed in this release.

Table 17 - Bug Fixes History (Sheet 1 of 3)

Internal Ref.	Issue
979293/ 992180	Description: RDMA READ Message Rate (MR) of 8bytes is lower than expected in ConnectX-5 compare to ConnectX-4 adapter cards.
	Keywords: RDMA READ Message Rate
	Discovered in Release: 16.19.1000
	Fixed in Release: 16.19.1200
1007303	Description: Fixed a rare case of long copper cables reinsertion which caused a long link-up time.
	Keywords: Cable, port, phy, linkup time
	Discovered in Release: 16.19.1000
	Fixed in Release: 16.19.1200
975514	Description: Fixed an issue causing the mkey command query for mkey type KSM to mask some of the KLM bits in the command output.
	Keywords: mkey, KLM
	Discovered in Release: 16.18.1000
	Fixed in Release: 16.19.1200
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: 16.18.1000
	Fixed in Release: 16.19.1200
988184/ 974569	Description: On rare occasions, VF FLR/reboot can cause TX to hang.
	Keywords: VF FLR/reboot, TX
	Discovered in Release: 16.19.1000
	Fixed in Release: 16.19.1200
963653/ 961833	Description: Diagnostic counters are not reset when enabled with on_demand mode.
	Keywords: on_demand mode, Diagnostic counters
	Discovered in Release: 16.18.1000
	Fixed in Release: 16.19.1000

Table 17 - Bug Fixes History (Sheet 2 of 3)

Internal Ref.	Issue
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: 16.18.1000
	Fixed in Release: 16.19.1000
945987	Description: When performing soft reboot after burning (after ~10 iteration), the card disappears from the LSPCI output and the mst status.
	Keywords: Soft reboot, LSPCI output, mst status
	Discovered in Release: 16.18.1000
	Fixed in Release: 16.19.1000
959527	Description: Missing invalidation upon Set () .pkey leads to bad Pkey checks.
	Keywords: Pkeys, PortInfo.LID
	Discovered in Release: 16.18.1000
	Fixed in Release: 16.19.1000
928872	Description: When performing Pkey check for IPoIB enhanced traffic, the Pkey membership bit is ignored.
	Keywords: Pkeys
	Discovered in Release: 16.18.1000
	Fixed in Release: 16.19.1000
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: 16.18.1000
	Fixed in Release: 16.19.1000
919526	Description: HCA mad response contains the incoming packet Pkey and not the matched Pkey.
	Keywords: Pkey
	Discovered in Release: 16.18.1000
	Fixed in Release: 16.19.1000

Table 17 - Bug Fixes History (Sheet 3 of 3)

Internal Ref.	Issue
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: 16.18.1000
	Fixed in Release: 16.19.1000
687113	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: 16.18.1000
	Fixed in Release: 16.19.1000
655688	Description: When arming SRQ for limit event, the device might issue an event with <code>context_index=0</code> .
	Keywords: RoCE
	Discovered in Release: 16.18.1000
	Fixed in Release: 16.19.1000
759571/ 759655	Description: Modifying the <code>encap_id</code> of FTE is not supported.
	Keywords: <code>encap_id</code> , FTE
	Discovered in Release: 16.18.1000
	Fixed in Release: 16.19.1000

5 Firmware Changes and New Feature History

Table 18 - Firmware Changes and New Feature History (Sheet 1 of 7)

Feature/Change	Description
Rev. 16.18.2000 (Beta)	
Virtual Extensible LAN (VXLAN) encapsulation/decapsulation	[Beta Level] 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.
NVMeoF Target Offload	[Beta Level] Added support for NVMe over Fabrics (NVMeoF) offload, an implementation of the new NVMeoF standard target (server) side in hardware.
Tag Matching	[Beta Level] Added support for offloading MPI tag matching to HCA. It matches the source send operations to the destination receivers.
VLAN Switch Tagging (VST)	Enables the virtual machine interface to have no VLAN tag over it, thus allowing VLAN tagging to be handled by the Hypervisor.
On Demand Paging (ODP)	Added supported for Demand Paging (ODP).
Rev. 16.18.1000 (Beta)	
Power MGMT	<p>Added support for PCIe Express standard "Slot capability register" message (PCIe base rev 3.1, section 6.9 – "Slot Power Limit Control")</p> <p>When ConnectX-5 Ex based adapter is inserted to a PCIe slot that supports the reporting of the slot power limit control, the ConnectX-5 Ex may disable the 2nd port if PCIe slot message, showing that the power in this slot is insufficient. If not, both ports will stay enabled.</p> <p>In cases where ConnectX-5 Ex based adapter is inserted to a PCI slot that DOES NOT support the "Slot capability register" message, and the adapters' configuration is 2 active optic cables/ transceivers, only one port will be enabled (the first inserted optic).</p> <p>Custom and OEM branded card based on ConnectX-5 Ex may be configured by INI to support/not-support the Power management feature.</p> <p>In hosts which do not support the "Slot capability register" Message and have enough power to support 2 active optical cable, the user will have the option to override the configuration resulted from "Slot capability register" by running the following NVconfig command:</p> <ul style="list-style-type: none"> • <code>echo "MLNX_RAW_TLV_FILE" > /tmp/power_conf_tlv.cfg;</code> <code>echo "0x00000004 0x00000088 0x00000000 0xc0000000"</code> <code>>> /tmp/power_conf_tlv.cfg</code> • <code>mlxconfig -d <device> -f /tmp/power_conf_tlv.cfg</code> <code>set_raw</code> • <code>mlxfwreset -d <device> reset</code> <p>For details on ConnectX-5 Ex power, please refer to ConnectX-5 Ex Datasheet</p>

Table 18 - Firmware Changes and New Feature History (Sheet 2 of 7)

Feature/Change	Description
NVM Express over Fabrics (NVMeF)	NVMeF is a protocol for communicating block storage IO requests over RDMA. For further information, please refer to the PRM section <i>“NVMe over Fabric Target Application Offload (NVMeF)”</i> .
Tag Matching	In Tag Matching, the software holds a list of matching entries called matching list. Each matching entry contains a tag and a pointer to an application buffer. The matching list is used to steer arriving messages to a specific buffer according to the message tag. The action of traversing the matching list and finding the matching entry is called tag matching. For further information, refer to the PRM section <i>“Tag Matching and Rendezvous Offload”</i>
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 collection of data on the PHY link status and provides reasons for any link down cases as well as 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.
Performance	Improved performance for Send Queues (SQs) transmitting multiple priorities in a single Traffic Class (TC) configuration.
Dropless TCP	Added the ability to avoid packet drops due to temporary lack of posted Receive buffers (WQEs), for trusted Receive Queues (RQs).
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.
Counters	Added support for additional transport counters.
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.
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.
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 18 - Firmware Changes and New Feature History (Sheet 3 of 7)

Feature/Change	Description
GENEVE & IP-in-IP Stateless Offload	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.
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.
Virtual Functions (VF) per Port	The maximum Virtual Functions (VF) per port is 95. 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> Note: For the maximum number of VFs supported by your driver, please refer to your drivers' Release Notes or User Manual.
QoS per VFs	[InfiniBand Only] Added support for multiple VLs in SR-IOV/multi-host environments. Note: The number of VLs can be configured by the NVCONFIG. The default VL number is 4 VLs.
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) .

Table 18 - Firmware Changes and New Feature History (Sheet 4 of 7)

Feature/Change	Description
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: <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: <ul style="list-style-type: none"> • Doorbell from User Access Region (UAR) • Clear interrupt from User Access Region (UAR)
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	Added support for new performance counters.
Data Center Bridging Exchange (DCBX)	DCBX is used by DCB devices to exchange configuration information with directly connected peers. DCBX uses Link Layer Discovery Protocol (LLDP) to exchange parameters between two link peers. For further information, please refer to the PRM.
Access Register: Default Values Revert	Allows network port registers to revert to their default values when the driver is restarted or the host is rebooted.
Link up Modes	Added additional network link up modes. The new modes decide when to keep the network link up. The new modes are: <ul style="list-style-type: none"> • keep_eth_link_up • keep_ib_link_up • keep_link_up_on_boot • keep_link_up_on_standby
Explicit Congestion Notification (ECN)	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.
RoCE Link Aggregation (RoCE LAG)	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 NIC hardware (Mellanox Embedded Switch or eSwitch) while maintaining OVS control-plane unmodified.
FCS no scatter / FCS check	Enables the user to control whether or not to scatter Frame Check Sequence (FCS) or to check FCS functionality.

Table 18 - Firmware Changes and New Feature History (Sheet 5 of 7)

Feature/Change	Description
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
IPoIB checksum and LSO off-load	Added IPoIB checksum and LSO offload support
Scatter FCS in RQ	Enables software to scatter or strip FCS in RQ.
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.
Flow Steering Counters	Provides a clear indication of Flow Steering statistics and errors.
WQE Inline Header	The minimal amount of packet headers inlined in the WQE's Eth Segment.
table-miss Flow	A flow table may include a table-miss flow entry, which renders all Match Fields wildcards. If a packet does not match a flow entry in a flow table, this is a table miss. The behavior on a table miss depends on the table configuration. A table-miss flow entry in the flow table may specify how to process unmatched packets.
Strided WQE User Space	Striding RQ is a receive queue comprised by work queue elements (i.e. WQEs), where multiple packets of LRO segments (i.e. message) are written to the same WQE.
SR-IOV (EN eSwitch & RoCE)	Single Root IO Virtualization (SR-IOV) is a technology that allows a physical PCIe device to present itself multiple times through the PCIe bus.
Vector Calculation/ Erasure Coding Offload	Uses the HCA for offloading erasure coding calculations.

Table 18 - Firmware Changes and New Feature History (Sheet 6 of 7)

Feature/Change	Description
Link params modification via access registers	<p>The change includes the following:</p> <ol style="list-style-type: none"> 1. Changed port configuration which required link re-training (such as speed) 2. PAOS down 3. PAOS up <p>This change, will cause the link to toggle and new configurations to take effect.</p>
Checksum Calculation on Image/Device	<p>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.</p>
Port Link	<p>Reduced the port link-up time when negotiating according to Clause 73 (DME)</p>
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 • Multi packet WQE • Minimal Bandwidth Guarantee (ETS) • Explicit Congestion Notification (ECN) • Priority Flow Control (PFC)
PCI	<ul style="list-style-type: none"> • PCIe Function Level Reset (FLR) • Power Management L2/L3 flow support
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, NVGRE is supported through Windows ONLY • SR-IOV EN
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 43.
Port management	<ul style="list-style-type: none"> • Enabled port management. Now one port can be set as Ethernet and one as InfiniBand.

Table 18 - Firmware Changes and New Feature History (Sheet 7 of 7)

Feature/Change	Description
InfiniBand Network	<ul style="list-style-type: none"> • Dynamically Connected (DC) transport Note: There is no interoperability between ConnectX-5 and ConnectX-4 adapter cards when using DC. • Unreliable Datagram Connection transport • Atomic Operation • CORE-Direct® <ul style="list-style-type: none"> • 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
General	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • NC-SI over RMII support
MAD	<ul style="list-style-type: none"> • 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 19 - FlexBoot Changes and New Features

Version	Description
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	Enhanced FlexBoot/firmware debug capability using Flexboot UI
Upstream sync	Synced the source with iPXE (upstream sync)

6.1 FlexBoot Known Issues

Table 20 - FlexBoot Known Issues (Sheet 1 of 5)

Internal Ref.	Issue
-	<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	<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 filename without next-server (siaddr).</p> <p>Workaround: N/A</p> <p>Keywords: PXE Boot</p>
-	<p>Description: In ConnectX-4, the PXE boot time measurement over TFTP Ethernet is 3:40 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.</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>

Table 20 - FlexBoot Known Issues (Sheet 2 of 5)

Internal Ref.	Issue
-	Description: Boot over VLAN with IB port is currently not supported.
	Workaround: N/A
	Keywords: PXE Boot
-	Description: Some faulty boot loaders do not close the underlying UNDI device which may result in unexpected behavior and possible system crash after the OS starts to load.
	Workaround: N/A
	Keywords: PXE Boot
-	Description: Chain-loading gPXE stack may result in undesirable behavior.
	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 20 - FlexBoot Known Issues (Sheet 3 of 5)

Internal Ref.	Issue
655800	Description: 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: Flexboot supports only 2K MTU.
	Workaround: N/A
	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 20 - FlexBoot Known Issues (Sheet 4 of 5)

Internal Ref.	Issue
-	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: IPv6 over WDS 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
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
735159	Description: LACP packets handling is not supported during boot.
	Workaround: N/A
	Keywords: PXE boot
735159	Description: Interrupts are not supported.
	Workaround: N/A
	Keywords: Interrupts
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

Table 20 - FlexBoot Known Issues (Sheet 5 of 5)

Internal Ref.	Issue
776057	Description: Citrix PVS boot is not supported.
	Workaround: N/A
	Keywords: Citrix PVS boot

7 Unsupported Features and Commands

7.1 Unsupported Features

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

- Rate Limit (static rate)
- SR-IOV: Rate Limit per Function
- Packet Pacing
- 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: SR-IOV, and RoCE
- RoCE LAG for VFs and Multihost are not supported in RoCE LAG

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