



Connect. Accelerate. Outperform.™

Mellanox ConnectX®-4 Firmware Release Notes

Rev 12.14.1100

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (“PRODUCT(S)”) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES “AS-IS” WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER'S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies
 350 Oakmead Parkway Suite 100
 Sunnyvale, CA 94085
 U.S.A.
www.mellanox.com
 Tel: (408) 970-3400
 Fax: (408) 970-3403

© Copyright 2016. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, CloudX logo, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, GPUDirect®, InfiniHost®, InfiniScale®, Kotura®, Kotura logo, Mellanox Federal Systems®, Mellanox Open Ethernet®, Mellanox ScalableHPC®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, Open Ethernet logo, PhyX®, SwitchX®, TestX®, The Generation of Open Ethernet logo, UFM®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

Accelio™, CyPU™, FPGADirect™, HPC-X™, InfiniBridge™, LinkX™, Mellanox Care™, Mellanox CloudX™, Mellanox Multi-Host™, Mellanox NEO™, Mellanox PeerDirect™, Mellanox Socket Direct™, Mellanox Spectrum™, NVMeDirect™, StPU™, Spectrum logo, Switch-IB™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Chapter 1	Overview	6
1.1	Supported Devices	6
1.2	Supported Cables and Modules	7
1.2.1	Validated and Supported 1GbE Cables	7
1.2.2	Validated and Supported 10/40GbE Cables	7
1.2.3	Validated and Supported 25GbE Cables	9
1.2.4	Validated and Supported QDR/FDR10 Cables	9
1.2.5	Validated and Supported FDR Cables	10
1.2.6	Validated and Supported EDR 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	13
1.3.4	Tested 100Gb/s/EDR Switches	14
1.4	Tools, Switch Firmware and Driver Software	15
1.5	Supported FlexBoot	16
1.6	Revision Compatibility	16
Chapter 2	Changes and New Features in Rev 12.14.1100	17
Chapter 3	Known Issues	19
Chapter 4	Bug Fixes History	25
Chapter 5	Firmware Changes and New Feature History	29
Chapter 6	FlexBoot Changes and New Features	31
6.1	FlexBoot Known Issues	32
Chapter 7	Unsupported Features and Commands	37
7.1	Unsupported Features	37
7.2	Unsupported Commands	37
Chapter 8	Supported Non-Volatile Configurations	38

List of Tables

Table 1:	Release Update History	5
Table 2:	Supported PSIDs	6
Table 3:	Validated and Supported 1GbE Cables	7
Table 4:	Validated and Supported 10/40GbE Cables	7
Table 5:	Validated and Supported 25GbE Cables	9
Table 6:	Validated and Supported QDR/FDR10 Cable	9
Table 7:	Validated and Supported FDR Cables	10
Table 8:	Validated and Supported EDR Cables	11
Table 9:	Tested QDR Switches	13
Table 10:	Tested 10/40GbE Switches	13
Table 11:	Tested FDR Switches	13
Table 12:	Tested 10Gb/s/EDR Switches	14
Table 13:	Tools, Switch Firmware and Driver Software	15
Table 14:	Supported FlexBoot, UEFI and CLP	16
Table 15:	Firmware Rev 12.14.1100 Changes and New Feature	17
Table 16:	Known Issues	19
Table 17:	Fixed Bugs List	25
Table 18:	Firmware Changes and New Feature History	29
Table 19:	FlexBoot Changes and New Feature	31
Table 20:	FlexBoot Known Issues	32
Table 21:	Per-physical Port Settings	38
Table 22:	Global Settings	38
Table 23:	Per host/function Settings	38

Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 12.14.1100	November, 2015	Initial version of this firmware release.

1 Overview

These are the release notes for the ConnectX®-4 adapters firmware Rev 12.14.1100. This firmware supports the following protocols:

- InfiniBand - SDR, QDR, FDR10, FDR, EDR
- Ethernet - 1GigE, 10GigE, 25GigE, 40GigE, 50GigE, 56GigE¹ 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](#).

Table 2 - Supported PSIDs (Sheet 1 of 2)

Device Part Number	PSID	Device Name
MCX413A-BCAT	MT_2120110027	ConnectX®-4 EN network interface card, 40GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6
MCX413A-GCAT	MT_2600110035	ConnectX®-4 EN network interface card, 50GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6
MCX414A-BCAT	MT_2130110027	ConnectX®-4 EN network interface card, 40GbE dual-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6
MCX414A-GCAT	MT_2610110035	ConnectX®-4 EN network interface card, 50GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6
MCX415A-BCAT	MT_2120111027	ConnectX®-4 EN network interface card, 40GbE single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6
MCX415A-CCAT	MT_2140110033	ConnectX®-4 EN network interface card, 100GbE single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6
MCX415A-GCAT	MT_2120110035	ConnectX®-4 EN network interface card; 50GbE single-port QSFP28; PCIe3.0 x16; ROHS R6
MCX416A-BCAT	MT_2130111027	ConnectX®-4 EN network interface card, 40GbE dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6
MCX416A-CCAT	MT_2150110033	ConnectX®-4 EN network interface card, 100GbE dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R
MCX416A-GCAT	MT_2130110035	ConnectX®-4 EN network interface card; 50GbE dual-port QSFP28; PCIe3.0 x16; ROHS R6
MCX453A-FCAT	MT_2160110021	ConnectX®-4 VPI adapter card, FDR IB 40GbE, single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6
MCX454A-FCAT	MT_2170110021	ConnectX®-4 VPI adapter card, FDR IB and 40GbE, dual-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6

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 PSIDs (Sheet 2 of 2)

Device Part Number	PSID	Device Name
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
MCX455A-FCAT	MT_2160111021	ConnectX®-4 VPI adapter card, FDR IB and 40GbE, single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6
MCX456A-FCAT	MT_2170111021	ConnectX®-4 VPI adapter card, FDR IB and 40GbE, dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6
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

1.2 Supported Cables and Modules

Please refer to the LinkX™ Cables and Transceivers web page (www.mellanox.com -> Products -> Cables and Transceivers) 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-F	Mellanox Optical module, SX, 850nm
1GB/S	MC3208411-T-F	Mellanox ptical module, Base-T

1.2.2 Validated and Supported 10/40GbE Cables

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
10GB/S	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

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
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	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-SR	MELLANOX OPTICAL MODULE ETH 10GBE 10GB/S SFP+ LC-LC 850NM SR UP TO 300M
10GB/S	MFM1T02A-LR	MELLANOX OPTICAL MODULE ETH 10GBE 10GB/S SFP+ LC-LC 1310NM LR UP TO 10KM
10GB/S	Cisco SFP-H10GB-CU1M	Cisco SFP+ cable
10GB/S	Cisco SFP-H10GB-CU3M	Cisco SFP+ cable
10GB/S	Cisco SFP-H10GB-CU5M	Cisco SFP+ cable

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
10GB/S	MC2309124-007	QSFP-4SFP10G
10GB/S	SFP-10G-SR	CISCO 10GBASE-SR SFP Module
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	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	QSFP-H40G-CU1M	Cisco QSFP 40GbE cable
40GB/S	QSFP-H40G-CU3M	Cisco QSFP 40GbE cable
40GB/S	QSFP-H40G-CU5M	Cisco QSFP 40GbE cable
40GB/S	QSFP-40G-SR4	CISCO 40G QSFP Module
NA	MAM1Q00A-QSA	MELLANOX QSFP TO SFP+ ADAPTER

1.2.3 Validated and Supported 25GbE Cables

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-A01A	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1.5m
25GB/S	MCP2M00-A002	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2m
25GB/S	MCP2M00-A02A	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2.5m
25GB/S	MCP2M00-A003	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 3m

1.2.4 Validated and Supported QDR/FDR10 Cables

Table 6 - Validated and Supported QDR/FDR10 Cable

Speed	Cable OPN #	Description
QDR	MC2206125-007	MELLANOX PASSIVE COPPER CABLE IB QDR 40GB/S QSFP 7M

Table 6 - Validated and Supported QDR/FDR10 Cable

Speed	Cable OPN #	Description
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-E	MELLANOX ACTIVE FIBER CABLE IB QDR/FDR10 40GB/S QSFP from 3M up to 100M
FDR10	MC2206310-XXX-F	MELLANOX ACTIVE FIBER CABLE IB QDR/FDR10 40GB/S QSFP from 3M up to 100M
FDR10	MC2206310-300-L	MELLANOX ACTIVE FIBER CABLE IB QDR/FDR10 40GB/S QSFP 300M
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 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

Table 7 - Validated and Supported FDR Cables

Speed	Cable OPN #	Description
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-XXX-E	MELLANOX ACTIVE FIBER CABLE VPI UP TO 56GB/S QSFP from 3M up to 100M
FDR	MC2207310-XXX-T	MELLANOX ACTIVE FIBER CABLE VPI UP TO 56GB/S QSFP from 3M up to 100M
FDR	MC2207312-050	MELLANOX ACTIVE FIBER CABLE VPI UP TO 56GB/S QSFP from 3M up to 300M
FDR	MC2207310-100	MELLANOX ACTIVE FIBER CABLE VPI UP TO 56GB/S QSFP from 3M up to 100M
FDR	MC2207411-SR4L	MELLANOX OPTICAL MODULE IB FDR 56GB/S QSFP MPO 850NM UP TO 30M

1.2.6 Validated and Supported EDR Cables

Table 8 - Validated and Supported EDR Cables

Speed	Cable OPN #	Description
100GB/S	MMA1B00-C100	Mellanox® transceiver, 100GbE, QSFP28, MPO, 850nm, up to 100m
100GB/S	MCP1600-C00A	MELLANOX PASSIVE COPPER CABLE ETH 100GBE 100GBS QSFP LSZH 0.5M
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	MFA1A00-CXXX-M	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, up to 100m
EDR	MMA1B00-E100	Mellanox® transceiver, IB EDR, up to 100Gb/s, QSFP28, MPO, 850nm, up to 100m
EDR	MCP1600-E001	MELLANOX PASSIVE COPPER CABLE VPI 100GB/S QSFP LSZH 1M
EDR	MCP1600-E002	MELLANOX PASSIVE COPPER CABLE VPI 100GB/S QSFP LSZH 2M

Table 8 - Validated and Supported EDR Cables

Speed	Cable OPN #	Description
EDR	MCP1600-E003	MELLANOX PASSIVE COPPER CABLE VPI 100GB/S QSFP LSZH 3M
EDR	MCP1600-E00A	MELLANOX PASSIVE COPPER CABLE VPI 100GB/S QSFP LSZH 0.5M
EDR	MFA1A00-EXXX	MELLANOX active fiber cable, VPI, up to 100Gb/s, QSFP, up to 100m

1.3 Tested Switches

1.3.1 Tested QDR Switches

Table 9 - Tested QDR Switches

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

1.3.2 Tested 10/40GbE Switches

Table 10 - Tested 10/40GbE Switches

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

1.3.3 Tested FDR Switches

Table 11 - Tested FDR Switches

Speed	Switch Family	OPN # / Name	Description
FDR	SwitchX®-2	SX6710-FB2F2	36-port 56Gb/s InfiniBand/VPI Switch Systems
FDR	SwitchX®	SX6018F-1SFR	18-port 56Gb/s InfiniBand/VPI Switch Systems

1.3.4 Tested 100Gb/s/EDR Switches

Table 12 - Tested 10Gb/s/EDR Switches

Speed	Switch Family	OPN # / Name	Description
EDR	Switch-IB	SB7790-EB2F	36-port EDR 100Gb/s InfiniBand Switch Systems
100Gb/s	Spectrum	SN2700-CS2R	32-port Non-blocking 100GbE Open Ethernet Spine Switch System
100Gb/s	N/A	C3232C	High-Density, 100 Gigabit Ethernet Switch

1.4 Tools, Switch Firmware and Driver Software

Firmware Rev 12.14.1100 is tested with the following tools, SwitchX® firmware, and driver software:

Table 13 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_OFED	3.2-1.0.0/3.1-1.0.3
MLNX_EN (MLNX_OFED based code)	3.2-1.0.0/3.1-1.0.3
WinOF-2	1.35/1.30
MFT	4.2.0/4.1.0
MLNX-OS	<ul style="list-style-type: none"> • SwitchX: 3.4.3004 • Switch-IB: 3.4.3206 • Spectrum: 3.5.0530
SwitchX®/SwitchX®-2 Firmware	9.3.5080
Spectrum™ Firmware	13.300.328
SwitchX-IB™ Firmware	11.0202.0126
InfiniScale® V Firmware	7.4.3000/v7.4.2200
Linux Inbox Drivers	<ul style="list-style-type: none"> • RHEL5.9 • RHEL5.10 • RHEL5.11 • RHEL5.12 • RHEL6.0 • RHEL6.1 • RHEL6.2 • RHEL6.3 • RHEL6.4 • RHEL6.5 • RHEL6.6 • RHEL7.0 • RHEL7.1 • Ubuntu 12.04 • Ubuntu 14.04 • SLES11.2 • SLES11.3 • SLES12.0
Windows Inbox Drivers	<ul style="list-style-type: none"> • Windows Server 2016 (Beta)

1.5 Supported FlexBoot

Firmware Rev 12.14.1100 supports the following FlexBoot:

Table 14 - Supported FlexBoot, UEFI and CLP

	Supported Version
FlexBoot	3.4.719

1.6 Revision Compatibility

Firmware Rev 12.14.1100 complies with the following programmer's reference manual:

- *Mellanox Adapters Programmer's Reference Manual (PRM), Rev 0.31 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 12.14.1100

Table 15 - Firmware Rev 12.14.1100 Changes and New Feature

Feature/Change	Description
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
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.
Multihost 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 Multihost 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 .

Table 15 - Firmware Rev 12.14.1100 Changes and New Feature

Feature/Change	Description
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 .

3 Known Issues

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

Table 16 - Known Issues (Sheet 1 of 6)

Internal Ref.	Description
-	<p>Description: Bit error rate is not optimal on QDR links</p> <p>WA: N/A</p> <p>Key Words: Link Speed</p>
572150	<p>Description: A low link speed issue occurs when connecting a ConnectX®-4 EDR adapter card with a QDR InfiniScale® IV based switch. The link is raised as DDR.</p> <p>WA: N/A</p> <p>Key Words: Link Speed</p>
-	<p>Description: To raise links with platforms based on the following ICs, comply with the following firmware version requirements:</p> <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) <p>WA: N/A</p> <p>Key Words: Interoperability</p>
682518	<p>Description: Interoperability issue between ConnectX-4 or ConnectX-4 Lx adapter cards and ConnectX-2 adapter card when trying to raise a 10GbE link.</p> <p>WA: N/A</p> <p>Key Words: Interoperability</p>
-	<p>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)</p> <p>WA: Enable QDR (in addition to FDR) when connecting to peer ports running at FDR</p> <p>Key Words: Interoperability</p>
-	<p>Description: Qualified EDR cables currently work with EDR networks (EDR devices, Switch®-IB and ConnectX®-4) only.</p> <p>WA: N/A</p> <p>Key Words: Cables</p>
-	<p>Description: PCIe capability “Device S/N” returns false value.</p> <p>WA: N/A</p> <p>Key Words: PCI</p>

Table 16 - Known Issues (Sheet 2 of 6)

Internal Ref.	Description
-	Description: When the link is Gen2, entering or exiting L1 state may cause bad CRC or DLLP indication.
	WA: N/A
	Key Words: PCI
600534	Description: Configuration space power management capability PME_EN cannot be set.
	WA: N/A
	Key Words: PCI
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.
	WA: N/A
	Key Words: PCI
685062	Description: MultiHost InfiniBand: OpenSM is supported over host0 only and the MAD_IFC usage is limited to host0 only.
	WA: Activate OpenSM and the MFT tools via host0
	Key Words: PCI
-	Description: PF direct pass-through is not supported (since PF FLR is not supported)
	WA: N/A
	Key Words: PF direct pass-through
687113	Description: Some Port Control Register do not return to the default value after the last port owner host restarts the driver.
	WA: Reboot or reset the driver. reboot / mlxfwreset
	Key Words: PRM

Table 16 - Known Issues (Sheet 3 of 6)

Internal Ref.	Description
-	<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: # flint -d <mst device> -i <image> burn</p> <pre>Current FW version on flash: 12.1100.6630 New FW version: 12.0012.0572</pre> <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>WA: 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>Key Words: Firmware Upgrade/MFT</p>
591240	<p>Description: Traffic that is loopbacked due to QP.force_loopback being equaled to 1, is steered to the PF.</p> <p>WA: N/A</p> <p>Key Words: Ethernet Network</p>
594964	<p>Description: A minimum of 200 LFM is required in order to cool the MCX4411A-ACAN adapter card.</p> <p>WA: N/A</p> <p>Key Words: Temperature</p>
601485/ 599810	<p>Description: mlxfwreset does not function properly in old MFT versions after upgrading the firmware image.</p> <p>WA: Upgrade MFT to the latest release or use <code>reboot/power cycle</code> after upgrading firmware.</p> <p>Key Words: Firmware Tool</p>
-	<p>Description: Windows Server 2016 Inbox driver cannot work with firmware v12.12.0780</p> <p>WA: Use WinOF-2 v1.20 out-of-box driver</p> <p>Key Words: Windows Inbox Drivers</p>
-	<p>Description: Flashing the firmware requires server reboot. Firmware cannot be flashed twice without server reboot after first flashing</p> <p>WA: Reboot the server after firmware flashing</p> <p>Key Words: Upgrading/Downgrading</p>

Table 16 - Known Issues (Sheet 4 of 6)

Internal Ref.	Description
655688	<p>Description: When arming SRQ for limit event, the device might issue an event with <code>context_index=0</code>.</p> <p>WA: N/A</p> <p>Key Words: RoCE</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>WA: 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>Key Words: Virtualization, FLR</p>
-	<p>[For customers developing custom low level drivers]</p> <p>Description: <code>VNodeInfo</code> and <code>VPortGuidInfo</code> virtualization Attributes MADs are not supported.</p> <p>WA: N/A</p> <p>Key Words: Virtualization</p>
-	<p>[For customers developing custom low level drivers]</p> <p>Description: The value of <code>log_max_ra_res_qp</code> in <code>set_hca_cap</code> command should be the same in all functions.</p> <p>WA: N/A</p> <p>Key Words: Virtualization</p>
-	<p>Description: Function (PF/VF) TX port counters are not supported.</p> <p>WA: N/A</p> <p>Key Words: Virtualization</p>
-	<p>Description: PF driver must work with pages event queue.</p> <p>WA: N/A</p> <p>Key Words: Virtualization</p>
596637	<p>Description: SR-IOV Ethernet supports up to 18 VFs per port only.</p> <p>WA: N/A</p> <p>Key Words: Virtualization</p>
597718	<p>Description: Privileged Vport egress traffic is not blocked when Vport is not active</p> <p>WA: N/A</p> <p>Key Words: Virtualization</p>

Table 16 - Known Issues (Sheet 5 of 6)

Internal Ref.	Description
591240	Description: Any local (internal) loopbacked packet is counted by the Vport counters, although Vport counters should count only traffic that crosses the Vport.
	WA: N/A
	Key Words: Virtualization
597718	Description: Vport number in virtual trap might be reported incorrectly
	WA: N/A
	Key Words: Virtualization
689471	Description: Single FTE that catches both untagged and prio-tagged packets (by giving an FTE with match_value.vlan_tag = 0 and match_value.vid = 0) is currently not supported.
	WA: N/A
	Key Words: Ethernet Network
688670	Description: Configuring the SM with VL weight 0 on some VL, and running traffic on it, causes the driver to hang during unload.
	WA: Cold power-cycle the server
	Key Words: QoS
651991	Description: OCBB is not displayed in the latest iLO versions.
	WA: N/A
	Key Words: OCBB compliance, iLO
649696	Description: If the PF driver or the tool (e.g. ethtool) use PAOS DOWN command (e.g. by <code>ifconfig down</code> or <code>ip link set down</code>), loopback traffic is blocked for all functions on this port (PF<->VFs / VF<->VF) In Multihost loopback, the traffic will be blocked once the firmware receives the PAOS down command from all PFs. However, the loopback traffic will not be blocked when the port is down due to the physical link (for example: cable plugged out, switch port down).
	WA: Do not drop the physical link down.
	Key Words: OCBB compliance, iLO
691194	Description: In some cases, a Bit Error Rate is not optimal on 10G/40G links.
	WA: N/A
	Key Words: 10G/40G links, Bit Error Rate
690614	Description: Non-volatile configuration of Port Type TLV more than 50 times might cause system hang.
	WA: Run <code>mlxconfig reset</code> after every 50 consecutive updates of port type TLV.
	Key Words: Non-volatile configuration, Port Type TLV

Table 16 - Known Issues (Sheet 6 of 6)

Internal Ref.	Description
689788	Description: Instability of Link Training Flow occurs during 100G Auto-Negotiation.
	WA: Configure the link partner to Forced Speed mode.
	Key Words: Link Training Flow, Auto-Negotiation
677359	Description: When Clause 74 Fire-Code FEC is active, and there are FC corrected errors, both the <code>FC_correctable</code> counter and the <code>FC_uncorrectable</code> counter are increment.
	WA: N/A
	Key Words: Clause 74 Fire-Code FEC, Counters
691387/ 691415	Description: In a Multihost setup, when running a single TCP stream, you might experience sub optimal throughput.
	WA: Use multiple streams to reach optimal results
	Key Words: Multihost setup, Performance, TCP stream
693832	Description: In a Multihost InfiniBand setup, Host#0 should be loaded first and unloaded last. Currently the device does not support <code>reboot/stop_driver/flr</code> on Host#0 while other hosts are active. Note: Host0 can be selected by changing the INI
	WA: N/A
	Key Words: Multihost InfiniBand
693832	Description: In an InfiniBand Multihost and SR-IOV setups, traffic should contain GRH (GID index), traffic without GRH will be forwarded to <code>vport0</code> (“Host0”). OpenSM should be configured as follow (<code>opensm.conf</code>): <ul style="list-style-type: none"> • <code>virt_enable</code> should be 2 • Enable Qos: <code>qos TRUE</code>
	WA: N/A
	Key Words: Multihost setup, Performance, TCP stream
693832	Description: In a Multihost InfiniBand setup, only chassis reboot is supported
	WA: N/A
	Key Words: Multihost setup, chassis

4 Bug Fixes History

Table 17 lists the bugs fixed in this release.

Table 17 - Fixed Bugs List

Internal Ref.	Description
659307	Description: Fixed a 25G and 50G link issue when Clause 91 RS FEC was active.
	Key Words: 25G and 50G link, Clause 91 RS FEC
	Discovered in Release: 12.12.1100
	Fixed in Release: 12.14.1100
676877	Description: Added a missing invalidation of eSwitch cache upon FLR which caused the upcoming driver load to either fail or not to be able to transmit.
	Key Words: Packet Transmit, FLR
	Discovered in Release: 12.12.1240
	Fixed in Release: 12.14.1100
668221	Description: Fixed an issue which prevented Vport counters from counting local loopback packets. Packets now are filter by the self-loopback prevention.
	Key Words: Vport, local loopback packets
	Discovered in Release: 12.12.1240
	Fixed in Release: 12.14.1100
667288	Description: Reported INTx as unsupported to allow PFs Passthrough on PowerKVM.
	Key Words: Passthrough, PowerKVM
	Discovered in Release: 12.12.1100
	Fixed in Release: 12.14.1100
596637	Description: SR-IOV Ethernet supports up to 18 VFs per port only.
	Key Words: Virtualization
	Discovered in Release: 12.12.1240
	Fixed in Release: 12.14.1100
591240	Description: Fixed and incident what allowed local (internal) loopbacked packets to be counted by the Vport counters, although Vport counters should count only traffic that crosses the Vport.
	Key Words: Virtualization
	Discovered in Release: 12.12.1240
	Fixed in Release: 12.14.1100

Table 17 - Fixed Bugs List

Internal Ref.	Description
664558	Description: Fixed an issue preventing driver loading or TX traffic sending upon reboot, after ungraceful driver unload.
	Key Words: Driver Load
	Discovered in Release: 12.12.1240
	Fixed in Release: 12.14.1100
657680	Description: Fixed casting of BMC MAC before steering API.
	Key Words: BMC, Steering API
	Discovered in Release: 12.12.1240
	Fixed in Release: 12.14.1100
614403	Description: Fixed the PCI write flow to take into consideration the PCI MTU. This fix eliminates the need for NOPs in the flow, which resulted from PPC larger PCI MTU. The single queue limitation for READ is due to a hardware limitation of the number of READ request in a given time.
	Key Words: PCI MTU
	Discovered in Release: 12.12.1100
	Fixed in Release: 12.14.1100
630327	Description: Fixed a case that caused FlexBoot to not work as expected with systems that run with "large bar" enabled (Above 4G Decoding) over Connect-IB or ConnectX-4 HCAs.
	Key Words: FlexBoot, 4G Decoding
	Discovered in Release: 12.12.1100
	Fixed in Release: 12.14.1100
629563	Description: Fixed an over-subscription on the RX buffer when Flow Control was not enabled which caused the RX pipe to hang.
	Key Words: Flow Control
	Discovered in Release: 12.12.0780
	Fixed in Release: 12.12.1240

Table 17 - Fixed Bugs List

Internal Ref.	Description
631225	<p>Description: Fixed an issue causing the firmware to hang when running ibdiagnet. The received DiagData MAD included the following values:</p> <ul style="list-style-type: none"> • Clear_all = 1 • PageNum = 0 • Port_select = 0 <p>To prevent the firmware from hanging, a port check was added to Set() as well.</p>
	<p>Key Words: ibdiagnet</p>
	<p>Discovered in Release: 12.12.1100</p>
	<p>Fixed in Release: 12.12.1240</p>
638024	<p>Description: Fixed an issue causing Port 2's GPIO present to be mapped to the wrong value in ConnectX-4 EN 50GbE single-port adapter card OPN MCX414A-GCAT</p>
	<p>Key Words: Port Link</p>
	<p>Discovered in Release: 12.12.1100</p>
	<p>Fixed in Release: 12.12.1240</p>
-	<p>Description: Removed request for Forward Error Correction (FEC) on copper cables of 2m or shorter. In order to work with Switch-IB without FEC, a minimum firmware version of 11.0200.0120 is required</p>
	<p>Key Words: Cables</p>
	<p>Discovered in Release: 12.1100.6630</p>
	<p>Fixed in Release: 12.12.0780</p>
565905	<p>Description: Fixed a LED issue in adapter cards with bi-color LEDs. The LEDs were activated simultaneously due to a firmware issue.</p>
	<p>Key Words: LEDs</p>
	<p>Discovered in Release: 12.1100.6630</p>
	<p>Fixed in Release: 12.12.0780</p>
-	<p>Description: Fixed an FDR10 incorrect speed indication reported due to the usage of a translation function from the hardware speed to the PRM speed twice.</p>
	<p>Key Words: Port Link</p>
	<p>Discovered in Release: 12.1100.6440</p>
	<p>Fixed in Release: 12.12.0780</p>

Table 17 - Fixed Bugs List

Internal Ref.	Description
592712	Description: Fixed a Phy manager PCS event handling when the port's next state was disable.
	Key Words: Phy Management
	Discovered in Release: 12.1100.6630
	Fixed in Release: 12.12.0780
561387	Description: Fixed an issue that caused invalid data returned by EyeOpening MAD.
	Key Words: MADs
	Discovered in Release: 12.1100.6440
	Fixed in Release: 12.12.0780
552595	Description: Fixed a system call trace event on ConnectX-4 OCP mezz card
	Key Words: Ethernet Network
	Discovered in Release: 12.1100.6440
	Fixed in Release: 12.12.0780
552462	Description: Fixed an issue which caused hardware fatal error when running ibdump.
	Key Words: Diagnostic Tools
	Discovered in Release: 12.1100.6440
	Fixed in Release: 12.12.0780
552227	Description: Reduced the VF ICM footprint for VFs.
	Key Words: Virtualization
	Discovered in Release: 12.1100.6440
	Fixed in Release: 12.1100.6630

5 Firmware Changes and New Feature History

Table 18 - Firmware Changes and New Feature History

Firmware Version	Description
12.12.1240	<ul style="list-style-type: none"> See Section 4, “Bug Fixes History”, on page 25
12.12.1100	<ul style="list-style-type: none"> Reduced the port link-up time when negotiating according to Clause 73 (DME)
12.12.0780	<ul style="list-style-type: none"> Large Receive Offload (LRO) Large Send Offload (LSO) Receive Side Scaling (RSS) Global Pause RoCEv1.0/RoCEv2.0 Flow Steering Sniffer Ethernet Rate Limiter (at Beta level) Multi packet WQE Enhanced Transmission Selection standard (ETS) Explicit Congestion Notification (ECN) Priority Flow Control (PFC) PCIe Function Level Reset (FLR) Power Management L2/L3 flow support 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 VXLAN/NVGRE Stateless offload In this release, this feature is supported through Windows ONLY SR-IOV EN (at Beta level) CQE zipping Dynamically Connected (DC) transport Wake-on-Lane/Standby FlexBoot/UEFI support Non-Volatile Configuration (NVConfig). For the complete list, please refer to Section 8, on page 38.
12.12.0780	<ul style="list-style-type: none"> Enabled port management. Now one port can be set as Ethernet and one as InfiniBand.
12.1100.6630	<ul style="list-style-type: none"> Added support for SR-IOV Added support for MADs Virtualization Attributes according to <code>ib_virt_annex_v17</code> Updated virtualization command set according to PRM 0.26 Enabled SR-IOV, NUM_VFS and INT_LOG_MAX_PAYLOAD_SIZE configuration via the <code>mlxconfig</code> tool

Table 18 - Firmware Changes and New Feature History

Firmware Version	Description
12.0100.6440	<ul style="list-style-type: none"> • Initial Release of ConnectX®-4 adapter cards • InfiniBand port speed up to EDR • Ethernet port speed up to 100GigE • Function per port • Dynamically Connected transport • 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 <p>Note: All the Ethernet features listed below are at Beta level.</p> <ul style="list-style-type: none"> • 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 • 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 • NC-SI over RMIi support • 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 Feature

Version	Description
Rev 3.4.719	<ul style="list-style-type: none"> • Added IPv6 support • Added x64 architecture support in ConnectX-4 and Connect-IB adapter cards • Removed support for the following SHELL CLI commands: <ul style="list-style-type: none"> • Non-volatile option storage commands • SAN boot commands • Menu commands • Login command • Sync command • DNS resolving command • Time commands • Image crypto digest commands • Loopback testing commands • VLAN commands • PXE commands • Reboot command <p>For further information, please refer to: http://ipxe.org/cmd</p> <ul style="list-style-type: none"> • Synced the source with iPXE (upstream sync)
Rev 3.4.650	<ul style="list-style-type: none"> • Added support for .mrom images larger than 128kB • Added support for ConnectX-4 EN and ConnectX-4 Lx EN • Synced the source with iPXE (upstream sync) • Moved to flat real mode when calling INT 1a,b101 to avoid BIOSes issues • Added support for detecting Spanning Tree Protocol non-forwarding ports (RSTP/MSTP)

6.1 FlexBoot Known Issues

Table 20 - FlexBoot Known Issues

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

Table 20 - FlexBoot Known Issues

Internal Ref.	Description
-	Description: Boot over VLAN with IB port is currently not supported.
	WA: N/A
	Key Words: 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.
	WA: N/A
	Key Words: PXE Boot
-	Description: Chain-loading gPXE stack may result in undesirable behavior.
	WA: N/A
	Key Words: PXE Boot
647143	Description: Executing a partial boot loop while only downloading the NBP and selecting localboot is unsupported and may cause undefined behavior.
	WA: N/A
	Key Words: 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.
	WA: N/A
	Key Words: PXE Boot
-	Description: iSCSI over IB is not tested.
	WA: N/A
	Key Words: iSCSI
-	Description: iSCSI over DCB is not supported.
	WA: N/A
	Key Words: 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.
	WA: N/A
	Key Words: iSCSI

Table 20 - FlexBoot Known Issues

Internal Ref.	Description
-	Description: Boot retries is currently not functional when booting from iSCSI.
	WA: N/A
	Key Words: iSCSI
655800	Description: IPv6 is not supported.
	WA: N/A
	Key Words: iSCSI
-	Description: Boot menu is displayed as READ ONLY if the HCA card does not support flash configuration.
	WA: N/A
	Key Words: User Interface
-	Description: FlexBoot Boot Menu will not be visible in serial output.
	WA: N/A
	Key Words: User Interface
-	Description: FlexBoot Boot Menu will not be shown in ConnectX-4 if the physical port is IB.
	WA: N/A
	Key Words: User Interface
-	Description: FlexBoot Boot Menu will not be shown in ConnectX-4 and ConnectX-4 Lx if the physical port is IB
	WA: N/A
	Key Words: User Interface
-	Description: Large Receive Offload (LRO) and iSCSI may not interoperate due to a bug in current Linux kernel distributions.
	WA: 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).
	Key Words: Networking
-	Description: Flexboot supports only 2K MTU.
	WA: N/A
	Key Words: Networking

Table 20 - FlexBoot Known Issues

Internal Ref.	Description
-	Description: 56Gb/s is currently not supported.
	WA: N/A
	Key Words: Link Speed
-	Description: Blink LEDs are currently not supported.
	WA: N/A
	Key Words: LED
-	Description: Setting the number of Virtual Functions higher than the machine's memory capability may cause memory issues and system instability.
	WA: N/A
	Key Words: Virtualization
-	Description: SLAM, FTP, HTTPS and SRP are currently not supported.
	WA: N/A
	Key Words: Protocols
-	Description: Occasionally, using the Spanning Tree Protocol (STP) in the switches may cause packet drops and boot failure in the system.
	WA: Enable the "edgemode" if disabled on the switch, or use either portfast or edgemode functionality on the switch ports connected to the NICs.
	Key Words: Protocols
-	Description: FCoE, BCV are not supported.
	WA: N/A
	Key Words: Protocols
655800	Description: IPv6 can only run if a RADVD service is running in the network.
	WA: N/A
	Key Words: Protocols
-	Description: IPv6 over IB is not supported.
	WA: N/A
	Key Words: Protocols
655800	Description: IPv6 over WDS is not supported.
	WA: N/A
	Key Words: Protocols

Table 20 - FlexBoot Known Issues

Internal Ref.	Description
655800	Description: Enabling IPv6 first and then IPv4 is currently not supported.
	WA: N/A
	Key Words: Protocols
656001	Description: Booting from WDS and Windows DHCP server when only Option 66 is enabled (without Option 67), is not supported.
	WA: N/A
	Key Words: DHCP
691148	Description: When connecting a pre-configured port with VLAN to an IB fabric, the port runs as Ethernet port with the VLAN tag.
	WA: Disable VLAN in the boot menu before connecting the Ethernet port to the Infini-Band fabric.
	Key Words: VLAN, Port Management
690792	Description: If the PMM fails to allocate memory, the system hangs since FlexBoot cannot load from the expansion ROM.
	WA: N/A
	Key Words: PMM, expansion ROM
689068	Description: In hybrid BIOSes, if the BIOS loads legacy driver without closing the UEFI driver, the legacy driver fails to load.
	WA: Run the BIOS in legacy only mode.
	Key Words: BIOS, legacy mode

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)
- Data integrity validation of control structures
- SM is not supported on VFs
- DC is not supported in: SR-IOV, and RoCE
- FLR is not supported in Multihost InfiniBand

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

8 Supported Non-Volatile Configurations

Table 21 - Per-physical Port Settings

Name	Parameter Index
VPI settings	0x12
RoCE CC	0x107
RoCE CC ECN	0x108

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

Table 23 - Per host/function Settings

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