



Mellanox ConnectX[®]-4 Firmware Release Notes

Rev 12.17.1502

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 2016. 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®, 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

Chapter 1	Overview	6
1.1	Supported Devices	6
1.2	Supported Cables and Modules	6
1.2.1	Validated and Supported 1GbE Cables	6
1.2.2	Validated and Supported 10/40GbE Cables	6
1.2.3	Validated and Supported QDR/FDR10 Cables	9
1.2.4	Validated and Supported 50Gbs Cables	9
1.2.5	Validated and Supported FDR Cables	10
1.2.6	Validated and Supported EDR/100GBs Cables	10
1.3	Supported Switches	12
1.3.1	Validated and Supported QDR Switches	12
1.3.2	Validated and Supported 10/40GbE Switches	12
1.3.3	Validated and Supported FDR Switches	12
1.3.4	Validated and Supported 100GbE/EDR Switches	13
1.4	Tools, Switch Firmware and Driver Software	14
1.5	Supported Fcode	15
1.6	Revision Compatibility	15
Chapter 2	Changes and New Features in Rev 12.17.1502	16
Chapter 3	Known Issues	17
Chapter 4	Bug Fixes History	24
Chapter 5	Firmware Changes and New Feature History	34
Chapter 6	Unsupported Features and Commands	42
6.1	Unsupported Features	42
6.2	Unsupported Commands	42
Chapter 7	Supported Non-Volatile Configurations	43

List of Tables

Table 1:	Release Update History	5
Table 2:	Supported PSIDs	6
Table 3:	Validated and Supported 1GbE Cables	6
Table 4:	Validated and Supported 10/40GbE Cables	6
Table 5:	Validated and Supported QDR/FDR10 Cable	9
Table 6:	Validated and Supported 50Gbs Cables	9
Table 7:	Validated and Supported FDR Cables	10
Table 8:	Validated and Supported EDR/100GBs Cables	10
Table 9:	Validated and Supported QDR Switches	12
Table 10:	Validated and Supported 10/40GbE Switches	12
Table 11:	Validated and Supported FDR Switches	12
Table 12:	Validated and Supported 100Gb/s/EDR Switches	13
Table 13:	Tools, Switch Firmware and Driver Software	14
Table 14:	Compiled Fcode version	15
Table 15:	Firmware Rev 12.17.1502 Changes and New Feature	16
Table 16:	Known Issues	17
Table 17:	Fixed Bugs List	24
Table 18:	Firmware Changes and New Feature History	34
Table 19:	Per-physical Port Settings	43
Table 20:	Global Settings	43
Table 21:	Per host/function Settings	43

Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 12.17.1502	November 15, 2016	Updated Section 1.5, “Supported Fcode” , on page 15 : update the Fcode version to 150
	October 2016	Initial version of this firmware release.

1 Overview

These are the release notes for the ConnectX®-4 adapters firmware Rev 12.17.1502. 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

Device Part Number	PSID	Device Name	FlexBoot	UEFI
00WT041	IBM2150110033	2-port 100GbE RoCE QSFP28 PCIe 3.0 x16 Adapter	No	No

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	Mellanox Optical module, SX, 850nm
1GB/S	MC3208411-T	Mellanox Optical module, Base-T

1.2.2 Validated and Supported 10/40GbE Cables

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
10GB/S	CAB-SFP-SFP-1M	Arista 10GBASE-CR SFP+ Cable 1 Meter
10GB/S	CAB-SFP-SFP-3M	Arista 10GBASE-CR SFP+ Cable 3 Meter
10GB/S	CAB-SFP-SFP-5M	Arista 10GBASE-CR SFP+ Cable 5 Meter
NA	MAM1Q00A-QSA	MELLANOX 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

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 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
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	MC2309124-007	Mellanox® passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 7m
10GB/S	MC2309130-001	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 1M
10GB/S	MC2309130-002	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 2M
10GB/S	MC2309130-003	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 3M
10GB/S	MC2309130-00A	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 0.5M
10GB/S	MC2609125-004	MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 4M
10GB/S	MC2609125-005	MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 5M
10GB/S	MC2609130-001	MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 1M
10GB/S	MC2609130-002	MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 2M
10GB/S	MC2609130-003	MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 3M

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
10GB/S	MC2609130-0A1	MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 1.5M
10GB/S	MC3309124-004	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 4M
10GB/S	MC3309124-005	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 5M
10GB/S	MC3309124-006	Mellanox® Passive Copper Cable, ETH 10GbE, 10Gb/s, SFP+, 6m
10GB/S	MC3309124-007	Mellanox® Passive Copper Cable, ETH 10GbE, 10Gb/s, SFP+, 7m
10GB/S	MC3309130-001	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 1M
10GB/S	MC3309130-002	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 2M
10GB/S	MC3309130-003	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 3M
10GB/S	MC3309130-00A	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 0.5M
10GB/S	MC3309130-0A1	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 1.5M
10GB/S	MC3309130-0A2	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 2.5M
10GB/S	MFM1T02A-LR-F	MELLANOX OPTICAL MODULE ETH 10GBE 10GB/S SFP+ LC-LC 1310NM LR UP TO 10KM
10GB/S	MFM1T02A-SR-F	MELLANOX OPTICAL MODULE ETH 10GBE 10GB/S SFP+ LC-LC 850NM SR UP TO 300M
40GB/S	QSFP-40G-SR-BD	Cisco 40GBASE-SR-BiDi, duplex MMF
40GB/S	QSFP-40G-SR4	Cisco 40GBASE-SR4, 4 lanes, 850 nm MMF
40GB/S	QSFP-H40G-ACU10M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 10-meter, active
40GB/S	QSFP-H40G-AOC10M	Cisco 40GBase-AOC QSFP direct-attach Active Optical Cable, 10-meter
40GB/S	QSFP-H40G-CU1M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 1-meter, passive
40GB/S	QSFP-H40G-CU3M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 3-meter, passive
40GB/S	QSFP-H40G-CU5M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 5-meter, passive
10GB/S	SFP-10G-SR	Cisco 10GBASE-SR SFP+ transceiver module for MMF, 850-nm wavelength, LC duplex connector
10GB/S	SFP-H10GB-CU1M	Cisco 1-m 10G SFP+ Twinax cable assembly, passive
10GB/S	SFP-H10GB-CU3M	Cisco 3-m 10G SFP+ Twinax cable assembly, passive

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
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 Cable

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/100GBs Cables

Table 8 - Validated and Supported EDR/100GBs 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

Table 8 - Validated and Supported EDR/100GBs Cables

Speed	Cable OPN #	Description
100GE	MCP1600-C003-AM	Mellanox® Passive Copper Cable ETH 100GbE 100GB/S QSFP 3M 30AWG
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	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 1M
EDR	MCP1600-E002	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	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 0.5M
EDR	MCP1600-E01A	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-CXXX	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, up to 100m
EDR	MFA1A00-EXXX	Mellanox active fiber cable, VPI, up to 100Gb/s, QSFP, up to 100m
100GB/S	MFS1200-CXXX	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, up to 100m
EDR	MFS1200-EXXX	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, up to 100m
100GB/S	MMA1B00-C100D	Mellanox® transceiver, 100GbE, QSFP28, MPO, 850nm, up to 100m
EDR	MMA1B00-E100	Mellanox® transceiver, IB EDR, up to 100Gb/s, QSFP28, MPO, 850nm, up to 100m
100GB/S	QSFP-100G-AOC10M	Cisco 100GBase QSFP Active Optical Cable, 10-meter

1.3 Supported Switches

1.3.1 Validated and Supported QDR Switches

Table 9 - Validated and Supported 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 Validated and Supported 10/40GbE Switches

Table 10 - Validated and Supported 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 Validated and Supported FDR Switches

Table 11 - Validated and Supported 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 Validated and Supported 100GbE/EDR Switches

Table 12 - Validated and Supported 100Gb/s/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 12.17.1502 is tested with the following tools, Switch firmware, and driver software:

Table 13 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_OFED	3.4-1.0.4.0/3.3-1.0.0.0
MLNX_EN (MLNX_OFED based code)	3.4-1.0.4.0/3.3-1.0.0.0
WinOF-2	1.45/1.41
MFT	4.5.0/4.4.0
WMware	4.15.6.17/4.5.6.17
MLNX-OS	<ul style="list-style-type: none"> • SwitchX: 3.6.1930 • Switch-IB: 3.6.1930 • Switch-IB 2: 3.6.1930 • Spectrum: 3.6.0530
SwitchX®/SwitchX®-2 Firmware	9.4.1050
Spectrum™ Firmware	13.1200.0084
SwitchX-IB™ Firmware	11.1200.0084
SwitchX-IB 2 Firmware	15.1200.0084
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 • RHEL7.2 • Ubuntu 12.04 • Ubuntu 14.04 • SLES11.2 • SLES11.3 • SLES12.0
Windows Inbox Drivers	Windows Server 2016

1.5 Supported Fcode

Firmware Rev 12.17.1502 is compiled with the following Fcode:

Table 14 - Compiled Fcode version

	Supported Version
Fcode	150

1.6 Revision Compatibility

Firmware Rev 12.17.1502 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 12.17.1502

Table 15 - Firmware Rev 12.17.1502 Changes and New Feature

Feature/Change	Description
Fcode	Updated the Fcode version from 140 to 150.
Packet Pacing	Added burst size bound exposure for packet pacing.
Bug Fixes	See Section 4, “Bug Fixes History” , on page 24

3 Known Issues

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

Table 16 - Known Issues (Sheet 1 of 7)

Internal Ref.	Issue
-	<p>Description: Bit Error Rate is not optimal on QDR links</p> <p>WA: N/A</p> <p>Keywords: 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>Keywords: 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>Keywords: 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>Keywords: 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>Keywords: 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>Keywords: Cables</p>
-	<p>Description: PCIe capability “Device S/N” returns false value.</p> <p>WA: N/A</p> <p>Keywords: PCI</p>

Table 16 - Known Issues (Sheet 2 of 7)

Internal Ref.	Issue
-	<p>Description: When the link is Gen2, entering or exiting L1 state may cause bad CRC or DLLP indication.</p> <p>WA: N/A</p> <p>Keywords: PCI</p>
600534	<p>Description: Configuration of space power management capability <code>PME_EN</code> cannot be set, thus preventing the driver from activating the wake signal.</p> <p>WA: N/A</p> <p>Keywords: PCIe</p>
685062	<p>Description: MultiHost InfiniBand: OpenSM is supported over host0 only and the <code>MAD_IFC</code> usage is limited to host0 only.</p> <p>WA: Activate OpenSM and the MFT tools via host0</p> <p>Keywords: PCI</p>
687113	<p>Description: PF direct pass-through is not supported in InfiniBand (since PF FLR is not supported)</p> <p>WA: N/A</p> <p>Keywords: PF direct pass-through, InfiniBand</p>
-	<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 Current FW version on flash: 12.1100.6630 New FW version: 12.0012.0572 Note: The new FW version is not newer than the current FW version on flash. Do you want to continue ? (y/n) [n] : y</pre> <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>Keywords: Firmware Upgrade/MFT</p>
591240	<p>Description: Traffic that is loopbacked due to <code>QP.force_loopback</code> being equaled to 1, is steered to the PF.</p> <p>WA: N/A</p> <p>Keywords: Ethernet Network</p>

Table 16 - Known Issues (Sheet 3 of 7)

Internal Ref.	Issue
594964	Description: A minimum of 200 LFM is required in order to cool the MCX4411A-ACAN adapter card.
	WA: N/A
	Keywords: Temperature
601485/ 599810	Description: mlxfwreset does not function properly in old MFT versions after upgrading the firmware image.
	WA: Upgrade MFT to the latest release or use <code>reboot/power cycle</code> after upgrading firmware.
	Keywords: Firmware Tool
-	Description: Windows Server 2016 Inbox driver cannot work with firmware v12.12.0780
	WA: Use WinOF-2 v1.20 out-of-box driver
	Keywords: Windows Inbox Drivers
655688	Description: When arming SRQ for limit event, the device might issue an event with <code>context_index=0</code> .
	WA: N/A
	Keywords: RoCE
-	[For customers developing custom low level drivers] Description: VFs internal FLR is not supported in PF teardown HCA command.
	WA: Before unloading the PF driver, PF driver must disable all its active VFs by performing the following: 1. Run the <code>disable_hca</code> command on all the <code>function_ids</code> 2. Wait until firmware returns all VFs allocated pages.
	Keywords: Virtualization, FLR
-	Description: Function (PF/VF) TX port counters are not supported.
	WA: N/A
	Keywords: Virtualization
-	Description: PF driver must work with pages event queue.
	WA: N/A
	Keywords: Virtualization
597718	Description: Privileged Vport egress traffic is not blocked when Vport is not active
	WA: N/A
	Keywords: Virtualization
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
	Keywords: Virtualization

Table 16 - Known Issues (Sheet 4 of 7)

Internal Ref.	Issue
597718	Description: Vport number in virtual trap might be reported incorrectly
	WA: N/A
	Keywords: Virtualization
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
	Keywords: Multihost setup, Performance, TCP stream
691754	Description: end_padding_mode is required in CREATE_QP and not in INIT_2_RTR command as defined in the PRM
	WA: N/A
	Keywords: end_padding_mode, PRM
691490	Description: LR4 cable events are sent although the port is up
	WA: N/A
	Keywords: Management
-	Description: QoS must be configured the same for both ports in order for RoCE LAG to function properly.
	WA: N/A
	Keywords: RoCE LAG
759571/ 759655	Description: Modifying the encap_id of FTE is not supported.
	WA: N/A
	Keywords: encap_id, FTE
756872/ 769604	Description: Flow Counter is supported only for FTE that does not include a flow_tag or for FTE that have TIR as a destination.
	WA: 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.
	WA: 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.
	WA: N/A
	Keywords: Flow Counter, FDB Flow Table, vport counters

Table 16 - Known Issues (Sheet 5 of 7)

Internal Ref.	Issue
748292	Description: When a steering rule in the e-sw FDB includes an encap action and an external port as destination, a transmitted multicast packet that matches the rule is sent to the wire and the loopback and the locally looped back packet will also have an encap header.
	WA: 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.
	WA: 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).
	WA: 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.
	WA: 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-4 firmware v12.16.1006 or newer, the following conditions must be met: <ul style="list-style-type: none"> • The minimum ConnectX-4 firmware version should be 12.16.1006 • The minimum ConnectX-4 Lx firmware version should be 14.16.1006 • The minimum MLNX-OS version should be 3.6.1000 (firmware v13.1100.0026)
	WA: 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.
	WA: 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.
	WA: Use <code>init 0</code> to start the Standby flow.
	Keywords: Warm/cold reboot

Table 16 - Known Issues (Sheet 6 of 7)

Internal Ref.	Issue
693832	<p>Description: In an InfiniBand Multihost and SR-IOV setups, traffic should contain GRH (GID index) if the grh_required bit is set in the query_hca_vport_context command. OpenSM should be configured as follow (opensm.conf):</p> <ul style="list-style-type: none"> • virt_enable should be 2 • Enable Qos: • qos TRUE <p>Note: In this case, traffic without GRH will be forwarded to vport0 ("Host0")</p> <p>WA: 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>WA: 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: Warm reboot, firmware image burning, VPI/IB devices</p>
758803	<p>Description: The firmware and the hardware do not reset the physical link upon CPort-State=down.</p> <p>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>WA: In order to re-train the physical link, sendbug PortInfo.physical_port_state = POLLING is required.</p> <p>Keywords: Physical link, CPortState=down</p>
840738	<p>Description: Local loopback traffic might effect vport counters.</p> <p>WA: N/A</p> <p>Keywords: Local loopback traffic, vport counters</p>
862480	<p>Description: Stopping the Rate Limiter while traffic is being transmitted might cause the adapter card to hang.</p> <p>WA: Stop traffic before:</p> <ul style="list-style-type: none"> • setting rate 0 to the last non-default-rate vport. • issuing destroy_scheduling_element command for the last vport with non-default-rate vport. <p>Keywords: Rate Limiter</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>WA: N/A</p> <p>Keywords: SL to VL mapping</p>

Table 16 - Known Issues (Sheet 7 of 7)

Internal Ref.	Issue
864200	Description: Running the <code>modify_scheduling_context</code> command does not include checking whether the scheduling element was created or not.
	WA: Do not modify non-existing elements
	Keywords: SR-IOV Rate Limiter
854805	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.
	WA: Set the <code>max_avergae_bw</code> in <code>scheduling_context</code> commands to equal or less than the supported wire speed.
	Keywords: Bandwidth rate, speed
865373/ 865820	Description: Although the ref counter design is to increment on every <code>max_average_bw != 0</code> (limited), the eSwitch <code>max_average_bw</code> ref counter decrements in <code>TEARDOWN_HCA/ FLR VF</code> regardless of the <code>max_average_bw</code> value set.
	WA: Verify all rates for all VFs are set to "0" before running <code>TEARDOWN_HCA/ FLR VF</code> . This is applicable only if a rate is set for any VF.
	Keywords: Bandwidth rate, VFs, <code>TEARDOWN_HCA/ FLR VF</code>
827444	Description: FDR link can raise with symbol errors on optic EDR cable longer than 30M.
	WA: N/A
	Keywords: FDR link, EDR cable
867367	Description: When running the <code>modify_scheduling_context</code> command, <code>scheduling_context.element_type</code> is taken into consideration with performing verifications, although the field is reserved.
	WA: Use the correct <code>element_type</code> when issuing <code>modify_scheduling_context</code> command.
	Keywords: SR-IOV Rate Limiter
882234	Description: On rare occasions, an assert is reported to the driver when the DC transport is enabled in the following cases: retransmission occurs and RX receives the same packet twice.
	WA: N/A
	Keywords: DC transport

4 Bug Fixes History

Table 17 lists the bugs fixed in this release.

Table 17 - Fixed Bugs List

Internal Ref.	Issue
868960	Description: Modified the TX configuration in the INI to resolve an EMI issue.
	Key Words: INI, EMI
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.17.1502
736528	Description: On rare occasions during UEFI boot cycles system got stuck while WinPE is loaded. (OS WinPE, system DL160).
	Keywords: WinPE OS, UEFI boot cycles
	Discovered in Release: 12.16.1010
	Fixed in Release: 12.17.1010
689471	Description: Single FTE that catches both untagged and prio-tagged packets (by giving an FTE with match_value.vlan_tag = 0 and match_value.vid = 0) is currently not supported.
	Keywords: Ethernet Network
	Discovered in Release: 12.14.1006
	Fixed in Release: 12.17.1010
-	Description: Flashing the firmware requires server reboot. Firmware cannot be flashed twice without server reboot after first flashing
	Keywords: Upgrading/Downgrading
	Discovered in Release: 12.14.1006
	Fixed in Release: 12.17.1010
854206	Description: If the vport state is DOWN and a packet is sent in local loopback, the sx_sniffer tool will not function.
	Keywords: sx_sniffer, vport
	Discovered in Release: 12.16.1010
	Fixed in Release: 12.17.1010
828608	Description: Fixed an issue causing bubbles to appear as symbol errors when link raised FDR 1x.
	Keywords: FDR 1x
	Discovered in Release: 12.16.1010
	Fixed in Release: 12.17.1010

Table 17 - Fixed Bugs List

Internal Ref.	Issue
677359	Description: When Clause 74 Fire-Code FEC is active, and there are FC corrected errors, both the FC_correctable counter and the FC_uncorrectable counter are increment.
	Keywords: Clause 74 Fire-Code FEC, Counters
	Discovered in Release: 12.16.1010
	Fixed in Release: 12.17.1010
687113	Description: Some Port Control Register do not return to the default value after the last port owner host restarts the driver.
	Keywords: PRM
	Discovered in Release: 12.14.1006
	Fixed in Release: 12.17.1010
827579/ 826702	Description: Fixed an issue which caused RX to hang when the UDP packet had a reserved UDP destination port.
	Key Words: UDP packet, RX
	Discovered in Release: 12.14.1006
	Fixed in Release: 12.17.1010
835735	Description: Fixed DMAC reporting mapping per host.
	Key Words: DMAC reporting
	Discovered in Release: 12.14.1006
	Fixed in Release: 12.17.1010
846520	Description: Fixed an EEH error from PCI which caused firmware to hang.
	Key Words: EEH error
	Discovered in Release: 12.14.1006
	Fixed in Release: 12.17.1010
857644	Description: Fixed the default value of the PCIe target_link_speed to Gen3 in link control2.
	Key Words: PCIe, Link Speed
	Discovered in Release: 12.14.1006
	Fixed in Release: 12.17.1010
783733/ 774373	Description: Fixed an issue which prevented LEDs from blinking when the traffic was less than 0.1% of the link speed.
	Key Words: LEDs, Blink, Link Speed
	Discovered in Release: 12.14.1006
	Fixed in Release: 12.17.1010

Table 17 - Fixed Bugs List

Internal Ref.	Issue
785128	Description: Fixed an issue which caused the mlxconfig configuration of VF_LOG_BAR_SIZE to be ignored and to be set to 5 (32MB).
	Key Words: mlxconfig configuration, VF_LOG_BAR_SIZE
	Discovered in Release: 12.14.1020
	Fixed in Release: 12.17.1010
689503/ 688670	Description: Fixed an issues which occasionally caused the driver to hang during unload on some VLs when configuring the SM with a VL weight 0 and running traffic on it.
	Keywords: VL, SM
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.17.1010
781039	Description: A server getting into a Standby mode while Packet-Pacing is enabled might cause firmware to hang and driver call-trace.
	Keywords: Packet-Pacing
	Discovered in Release: 12.14.1020
	Fixed in Release: 12.17.1010
753349	Description: Fixed an issue which caused unexpected QoS functionality in case of multiple sources to single destination traffic transmission.
	Key Words: QoS
	Discovered in Release: 12.14.1020
	Fixed in Release: 12.17.1010
801374	Description: Fixed an issue which occasionally caused the RX traffic to hang in DC when received a PCI error on WQE fetch.
	Key Words: RX traffic, DC
	Discovered in Release: 12.14.1020
	Fixed in Release: 12.17.1010
773110	Description: Fixed OOB connection issue during Intel's ITP inject errors test.
	Keywords: OOB, ITP inject errors test
	Discovered in Release: 14.14.2036
	Fixed in Release: 12.17.1010
773724	Description: Modified PCIe settings.
	Key Words: PCIe
	Discovered in Release: 12.14.1006
	Fixed in Release: 12.16.1020

Table 17 - Fixed Bugs List

Internal Ref.	Issue
670185	Description: Added protection from IOPX thermal diode destabilization to prevent UEFI IPv6 PXE boot failure on ConnectX-4 Lx 25GE OCP card.
	Key Words: UEFI, OCP card
	Discovered in Release: 12.14.2036
	Fixed in Release: 12.16.1006
780651	Description: Fixed an issue causing single port devices to query and write Physical Port TLVs to Port 2.
	Keywords: Physical Port TLVs, single port device
	Discovered in Release: 12.14.2036
	Fixed in Release: 12.16.1006
752392	Description: Enabled mlxfwreset to work using the PCIe Secondary Bus Reset.
	Keywords: mlxfwreset
	Discovered in Release: 12.14.2036
	Fixed in Release: 12.16.1006
775393	Description: Fixed an issue causing link flapping as a result, incorrect link settings.
	Keywords: Link flapping, link settings
	Discovered in Release: 12.14.2036
	Fixed in Release: 12.16.1006
756570	Description: Fixed an issue causing wrong alignment markers to be used when running 50G with Clause91 FEC enabled.
	Keywords: Clause91 FEC, 50G
	Discovered in Release: 12.14.2036
	Fixed in Release: 12.16.1006
767281	Description: Reduced the default BAR size for VF (SR-IOV) from 5 (32 MB) to 1 (2MB).
	Keywords: BAR size for VF (SR-IOV)
	Discovered in Release: 12.14.2036
	Fixed in Release: 12.16.1006
735159/ 747595/ 752533	Description: Added legacy interrupts support in FlexBoot.
	Keywords: Interrupts, PXE
	Discovered in Release: 12.14.2036
	Fixed in Release: 12.16.1006

Table 17 - Fixed Bugs List

Internal Ref.	Issue
752343	Description: Modified the TX configuration to support EMI crossing margins in 16Ghz
	Keywords: TX configuration, EMI
	Discovered in Release: 12.14.2036
	Fixed in Release: 12.16.1006
691194	Description: In some cases, a Bit Error Rate is not optimal on 10G/40G links.
	Keywords: 10G/40G links, Bit Error Rate
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036
689788	Description: Instability of Link Training Flow occurs during 100G Auto-Negotiation.
	Keywords: Link Training Flow, Auto-Negotiation
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036
684496	Description: Fixed a rare issue which caused the command to hang when moved the QP to RESET and back to RTS.
	Keywords: QP, RTS
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036
665089	Description: Improved RDMA READ bandwidth under packet lost scenario.
	Keywords: RDMA READ bandwidth
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036
736195	Description: Added support for pnat = 1 in HCA <code>access_reg</code> command as required by the <code>ibdiagnet</code> tool.
	Keywords: <code>access_reg</code> command, <code>ibdiagnet</code>
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036
696486	Description: Increased the steering hash tables static size from 128 to a maximum of 32K entries.
	Keywords: Steering hash tables static size
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036

Table 17 - Fixed Bugs List

Internal Ref.	Issue
691649	Description: Prevented miscalculation of module temperature when using 100Gb/s cables (OPN: MFA1A00-Cxxx for 100GbE and MFA1A00-Exxx for IB EDR).
	Keywords: 100Gb/s cables
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036
687096	Description: Fixed an issue which caused the device to hang when resetting qkey/pkey violation counter via port_info mad.
	Keywords: qkey/pkey violation counter, port_info mad
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036
693446	Description: Reduced one hop for Unicast RX steering, steering pipes balancing.
	Keywords: Ethernet Steering performance
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036
690614	Description: Non-volatile configuration of Port Type TLV more than 50 times might cause system hang.
	Keywords: Non-volatile configuration, Port Type TLV
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036
691043	Description: Enabled RoCE IPv4 Multicast prevents MCG command from failing when an IPv4 is mapped to an IPv6 address.
	Keywords: RoCE IPv4 Multicast
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036
649696/ 690681	Description: If the PF driver or the tool (e.g. ethtool) use PAOS DOWN command (e.g. by ifconfig down or ip link set down), loopback traffic is blocked for all functions on this port (PF<->VFs / VF<->VF) In 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).
	Keywords: Multihost loopback
	Discovered in Release: 12.14.1100
	Fixed in Release: 12.14.2036

Table 17 - Fixed Bugs List

Internal Ref.	Issue
659307	Description: Fixed a 25G and 50G link issue when Clause 91 RS FEC was active.
	Keywords: 25G and 50G link, Clause 91 RS FEC
	Discovered in Release: 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.
	Keywords: 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.
	Keywords: 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.
	Keywords: 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.
	Keywords: 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.
	Keywords: Virtualization
	Discovered in Release: 12.12.1240
	Fixed in Release: 12.14.1100
664558	Description: Fixed an issue preventing driver loading or TX traffic sending upon reboot, after ungraceful driver unload.
	Keywords: Driver Load
	Discovered in Release: 12.12.1240
	Fixed in Release: 12.14.1100

Table 17 - Fixed Bugs List

Internal Ref.	Issue
657680	Description: Fixed casting of BMC MAC before steering API.
	Keywords: 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.
	Keywords: 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.
	Keywords: 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.
	Keywords: Flow Control
	Discovered in Release: 12.12.0780
	Fixed in Release: 12.12.1240
631225	Description: Fixed an issue causing the firmware to hang when running ibdiagnet. The received DiagData MAD included the following values: <ul style="list-style-type: none"> • Clear_all = 1 • PageNum = 0 • Port_select = 0 To prevent the firmware from hanging, a port check was added to Set() as well.
	Keywords: ibdiagnet
	Discovered in Release: 12.12.1100
	Fixed in Release: 12.12.1240
638024	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
	Keywords: Port Link
	Discovered in Release: 12.12.1100
	Fixed in Release: 12.12.1240

Table 17 - Fixed Bugs List

Internal Ref.	Issue
-	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
	Keywords: Cables
	Discovered in Release: 12.1100.6630
	Fixed in Release: 12.12.0780
565905	Description: Fixed a LED issue in adapter cards with bi-color LEDs. The LEDs were activated simultaneously due to a firmware issue.
	Keywords: LEDs
	Discovered in Release: 12.1100.6630
	Fixed in Release: 12.12.0780
-	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.
	Keywords: Port Link
	Discovered in Release: 12.1100.6440
	Fixed in Release: 12.12.0780
592712	Description: Fixed a Phy manager PCS event handling when the port's next state was disable.
	Keywords: 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.
	Keywords: 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
	Keywords: 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.
	Keywords: Diagnostic Tools
	Discovered in Release: 12.1100.6440
	Fixed in Release: 12.12.0780

Table 17 - Fixed Bugs List

Internal Ref.	Issue
552227	Description: Reduced the VF ICM footprint for VFs.
	Keywords: 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

Feature/Change	Description
Rev. 12.17.1010	
Multi-Host LID Base Routing	<p>Added support for Multi-Host LID base routing. This feature requires a new OpenSM (v4.7.1 and above which comes with MLNX_OFED 3.3-2.0.0.0) with the following attributes:</p> <ul style="list-style-type: none"> • qos TRUE • lmc 2 (if there is no quad host in the fabric, you can set the lmc to 1) • virt_enabled 2 <p>Note: Multi-Host LID base routing can be configured by the INI only. The default is 0</p>
Resilient RoCE	<p>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.</p>
Multi-Host L3/L4 Classification	<p>Enables load balancing in the Multi PF Switch layer (MPFS) based on the L3/L4 headers</p>
InfiniBand Multi-Host Isolation	<p>Enabled isolation between separate Hosts using the same HCA. All the Hosts can be rebooted, the driver can be stopped and the FLR signal can be sent independently.</p>
95 Virtual Functions (VF) per Port	<p>Increased the number of VFs from 64 to 95 per Physical Function (PF).</p> <p>Note: When increasing the number of VFs, the following limitations must be taken into consideration:</p> <pre>server_total_bar_size >= (num_pfs) * (2log_pf_uar_bar_size + 2log_vf_uar_bar_size * total_vfs) server_total_msix >= (num_pfs) * (num_pf_msix + num_vfs_msix * total_vfs)</pre> <p>Note: For the maximum number of VFs supported by your driver, please refer to your drivers' Release Notes or User Manual.</p>
QoS per VFs	<p>[InfiniBand Only] Added support for multiple VLs in SR-IOV/multi-host environments.</p> <p>Note: The number of VLs can be configured by the NVCONFIG. The default VL number is 4 VLs.</p>
InfiniBand Rate Limit per QP (static rate)	<p>Added support for QP Rate Limit in InfiniBand.</p>
HCA Port Flap Counter	<p>Added support for Port Flap Counter.</p>
Fixed Buffer Size (KSM)	<p>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.</p>

Table 18 - Firmware Changes and New Feature History

Feature/Change	Description
NULL Mkey	This entry (null_mkey) is use to indicate non-present KLM/KSM entries. When accessing is, it causes the device to generate page fault event.
Out-of-Band Online Firmware Update: Firmware Update over PLDM	PLDM firmware burning is based on the DMTF spec DSP0267 (draft 9). The feature enables upgrading firmware and expansion ROM images using the PLDM protocol over MCTP (over PCIe). By doing so, a supporting BMC can query and upgrade the firmware without using OS based tools.
New Group in Ports Performance Counters (PPCNT)	Added a new physical layer statistics counters group. The new group includes BER counters, FEC error correction, clear time, and additional physical layer counters. For further information, please refer to the Ethernet Adapters Programming Manual (PRM) .
Permanent Link Up Mode	Enables the user to set a certain link up state for an unlimited period of time. This mode has 3 states: <ul style="list-style-type: none"> • Aux power (standby) • Reboot/boot/driver unloaded - the server is active and no driver is up • Driver is up - at least one driver is up (the time between init HCA and teardown or FLR)
No Driver NIC (NODNIC) Performance Improvement	Added support for Doorbell from User Access Region (UAR).
SR-IOV: Rate Limit Per Function	[Beta] Added support for maximum rate limit per function in SR-IOV.
Firmware Resiliency: Suppress Pauses	Allows the user to configure the adapter card to stop sending pauses after x when the receive port is unavailable (in a hang state).
Performance Back-pressure Counters	[Beta] Added support for new performance counters.
Data Center Bridging Exchange (DCBX)	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
Bug Fixes	See Section 4, “Bug Fixes History”, on page 24

Table 18 - Firmware Changes and New Feature History

Feature/Change	Description
Rev. 12.16.1020	
Bug Fixes	See Section 4, “Bug Fixes History” , on page 24
Rev. 12.16.1006	
Explicit Congestion Notification (ECN)	[Beta] Explicit Congestion Notification (ECN) is an extension to the Internet Protocol and to the Transmission Control Protocol. ECN allows end-to-end notification of network congestion without dropping packets.
64 VFs per port	Increased the number of VFs from 32 to 64 per PF. Note: When increasing the number of VFs, the following limitations must be taken into consideration: <code>server_total_bar_size >= (num_pfs)*(2log_pf_uar_bar_size + 2log_vf_uar_bar_size*total_vfs)</code> <code>server_total_msix >= (num_pfs)*(num_pf_msix + num_vfs_msix *total_vfs)</code>
RoCE Link Aggregation (RoCE LAG)	[Beta] RoCE Link Aggregation provides failover and link aggregation capabilities. In this mode, only one IB port, that represents the two physical ports, is exposed to the application layer. For further information, please refer to the PRM.
OVS Offload	Mellanox Accelerated Switching And Packet Processing (ASAP ²) Direct technology allows to offload OVS by handling OVS data-plane in Mellanox ConnectX-4 / ConnectX-4 Lx NIC hardware (Mellanox Embedded Switch or eSwitch) while maintaining OVS control-plane unmodified. The current actions supported by ASAP ² Direct include packet parsing and matching, forward, drop along with VLAN push/pop or VXLAN encaps/decap and HW based packet/byte flow statistics.
Virtual Extensible LAN (VXLAN) encapsulation/decapsulation	Virtual Extensible LAN (VXLAN) is a network virtualization technology that improves scalability problems associated with large cloud computing deployments. It tunnels Ethernet frames within Ethernet + IP + UDP frames. Mellanox implements VXLAN encapsulation and decapsulation in the hardware.
Data Center Bridging Exchange (DCBX)	[Beta] DCBX is used by DCB devices to exchange configuration information with directly connected peers. DCBX uses Link Layer Discovery Protocol (LLDP) to exchange parameters between two link peers. For further information, please refer to the PRM.
FCS no scatter / FCS check	Enables the user to control whether or not to scatter Frame Check Sequence (FCS) or to check FCS functionality.

Table 18 - Firmware Changes and New Feature History

Feature/Change	Description
Packet Pacing	[Beta] Send Queues (SQ/ Send queue of QP) may be individually rate limited, thus, allowing granular rate control over specific SW-defined flows. A rate-limited flow is allowed to transmit a few packets before its transmission rate is evaluated, and the next packet is scheduled for transmission accordingly.
PRBS Patterns Generation and Tuning	A new PHY test mode in which the device can generate different PRBS patterns for SerDes tuning purpose. For further information, please refer to PRM registers: PPAOS, PPTT, PPRT.
Management Controller Transport Protocol (MCTP) over PCI	Added support for MCTP host management over PCI
OCBB / OCSD support after mlxfwreset	Added support for OCBB/OCSD memory pointers restoration after mlxfwreset
MCTP media migration	Added support for MCTP media migration between SMBUS and PCI
Cables	Removed the RX amplitude configuration on some cable types
Bug Fixes	See Section 4, “Bug Fixes History”, on page 24
Rev. 12.14.2036	
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.
Bug Fixes	See Section 4, “Bug Fixes History”, on page 24
Rev. 12.14.1100	
CQE Time Stamping	Keeps track of the creation of a packet. A time-stamping service supports assertions of proof that a datum existed before a particular time.
Priority Flow Control (PFC)	Applies pause functionality to specific classes of traffic on the Ethernet link.
RDMA retransmission counters	Custom port counters provide the user a clear indication about RDMA send/receive statistics and errors.
Link Layer Discovery Protocol (LLDP)	The Link Layer Discovery Protocol (LLDP) is a vendor-neutral Link Layer protocol in the Internet Protocol Suite used by network devices for advertising their identity, capabilities, and neighbors on a IEEE 802 LAN. The protocol is formally defined in IEEE 802.1AB.
1GbE and 56GbE Link Speed	ConnectX-4adapters now support 1Gb/s and 56GbE Ethernet connectivity in addition to 10GigE, 25GigE, 40GigE, 50GigE, and 100GigE
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 18 - Firmware Changes and New Feature History

Feature/Change	Description
table-miss Flow	A flow table may include a table-miss flow entry, which renders all Match Fields wildcards. If a packet does not match a flow entry in a flow table, this is a table miss. The behavior on a table miss depends on the table configuration. A table-miss flow entry in the flow table may specify how to process unmatched packets.
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 .
Link params modification via access registers	The change includes the following: 1. Changed port configuration which required link re-training (such as speed) 2. PAOS down 3. PAOS up This change, will cause the link to toggle and new configurations to take effect.
Checksum Calculation on Image/Device	Flint utility allows performing an MD5 checksum on the non-persistent sections of the firmware image. For further information, please refer to MFT User Manual .
Rev. 12.12.1240	
Bug Fixes	See Section 4, “Bug Fixes History” , on page 24
Rev. 12.12.1100	
Port Link	Reduced the port link-up time when negotiating according to Clause 73 (DME)
Rev. 12.12.0780	

Table 18 - Firmware Changes and New Feature History

Feature/Change	Description
Ethernet Network	<ul style="list-style-type: none"> • Large Receive Offload (LRO) • Large Send Offload (LSO) • Receive Side Scaling (RSS) • Global Pause • RoCEv1.0/RoCEv2.0 • Flow Steering • Sniffer Ethernet • Rate Limiter (at Beta level) • Multi packet WQE • Minimal Bandwidth Guarantee (ETS) • Explicit Congestion Notification (ECN) • Priority Flow Control (PFC)
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"> • Self Loopback support • Transport Domain support • CQ2EQ remapping • Added support for the following commands: <ul style="list-style-type: none"> • MODIFY/QUERY_ESW_VPORT_CONTEXT • QUERY/MODIFY_CONG_STATUS • QUERY/MODIFY_CONG_PARAMS • QUERY_CONG_STATISTICS • ADD/DELETE_VXLAN_UDP_DPORT
Virtualization	<ul style="list-style-type: none"> • VXLAN/NVGRE Stateless offload In this release, this feature is supported through Windows ONLY • SR-IOV EN (at Beta level)
Performance	<ul style="list-style-type: none"> • CQE zipping
InfiniBand Network	<ul style="list-style-type: none"> • Dynamically Connected (DC) transport
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 7, 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.
Rev. 12.1100.6630	
Virtualization	<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>
PRM	<ul style="list-style-type: none"> • Updated virtualization command set according to PRM 0.26

Table 18 - Firmware Changes and New Feature History

Feature/Change	Description
Configuration tools	<ul style="list-style-type: none"> Enabled SR-IOV, NUM_VFS and INT_LOG_MAX_PAYLOAD_SIZE configuration via the mlxconfig tool
Rev. 12.0100.6440	
All	<ul style="list-style-type: none"> Initial Release of ConnectX®-4 adapter cards
Port Speed	<ul style="list-style-type: none"> InfiniBand port speed up to EDR Ethernet port speed up to 100GigE
Virtualization	<ul style="list-style-type: none"> Function per port
InfiniBand Network	<ul style="list-style-type: none"> 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
Ethernet Network	<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
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

Table 18 - Firmware Changes and New Feature History

Feature/Change	Description
Host management	<ul style="list-style-type: none"><li data-bbox="635 369 979 401">• NC-SI over RMIi support
MAD	<ul style="list-style-type: none"><li data-bbox="635 417 1318 449">• Config space address in MAD management class 0x09

6 Unsupported Features and Commands

6.1 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
- RoCE LAG for VFs and Multihost are not supported in RoCE LAG
- QoS per VFs feature is supported up to 14 VFs per PF in dual port device with 8 VLs.

6.2 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

7 Supported Non-Volatile Configurations

Table 19 - 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 20 - 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 21 - Per host/function Settings

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