



Mellanox ConnectX[®]-4 Firmware Release Notes

Rev 12.17.2052

NOTE:

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



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

© Copyright 2017. Mellanox Technologies Ltd. All Rights Reserved.

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

All other trademarks are property of their respective owners.

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

Table of Contents

Release Update History	5
Chapter 1 Overview	6
1.1 Supported Devices	6
1.2 Supported Cables and Modules	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 QDR/FDR10 Cables	9
1.2.4 Validated and Supported 50Gbs Cables	10
1.2.5 Validated and Supported FDR Cables	11
1.2.6 Validated and Supported EDR/100GB/s Cables	12
1.3 Tested Switches	14
1.3.1 Tested QDR Switches	14
1.3.2 Tested 10/40GbE Switches	14
1.3.3 Tested FDR Switches	16
1.3.4 Tested 100GbE/EDR Switches	16
1.4 Tools, Switch Firmware and Driver Software	16
1.5 Supported FlexBoot	18
1.6 Revision Compatibility	18
Chapter 2 Changes and New Features in Rev 12.17.2052	19
Chapter 3 Known Issues	20
Chapter 4 Bug Fixes History	28
Chapter 5 Firmware Changes and New Feature History	40
Chapter 6 FlexBoot Changes and New Features	48
6.1 FlexBoot Known Issues	49
6.2 FlexBoot Bug Fixes History	54
Chapter 7 UEFI Changes and Major New Features	56
7.1 UEFI Bug Fixes History	56
Chapter 8 Unsupported Features and Commands	57
8.1 Unsupported Features	57
8.2 Unsupported Commands	57
Chapter 9 Supported Non-Volatile Configurations	58

List of Tables

Table 1:	Document Update History	5
Table 2:	Supported Devices	6
Table 3:	Validated and Supported 1GbE Cables	7
Table 4:	Validated and Supported 10/40GbE Cables	7
Table 5:	Validated and Supported QDR/FDR10 Cables	9
Table 6:	Validated and Supported 50Gbs Cables	10
Table 7:	Validated and Supported FDR Cables	11
Table 8:	Validated and Supported EDR/100GB/s Cables	12
Table 9:	Tested QDR Switches	14
Table 10:	Tested 10/40GbE Switches	14
Table 11:	Tested FDR Switches	16
Table 12:	Tested 100GbE/EDR Switches	16
Table 13:	Tools, Switch Firmware and Driver Software	16
Table 14:	Supported FlexBoot	18
Table 15:	Changes and New Features in Rev 12.17.2052	19
Table 16:	Known Issues	20
Table 17:	Bug Fixes History	28
Table 18:	Firmware Changes and New Feature History	40
Table 19:	FlexBoot Changes and New Features	48
Table 20:	FlexBoot Known Issues	49
Table 21:	FlexBoot Bug Fixes History	54
Table 22:	UEFI Changes and Major New Features	56
Table 23:	UEFI Bug Fixes History	56
Table 24:	Per-physical Port Settings	58
Table 25:	Global Settings	58
Table 26:	Per host/function Settings	58

Release Update History

Table 1 - Document Update History

Release	Date	Description
Rev 12.17.2052	April 24, 2017	Initial version of this firmware release.

1 Overview

These are the release notes for the ConnectX[®]-4 adapters firmware Rev 12.17.2052. 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 Devices

Device Part Number	PSID	Device Name	Compiled with FlexBoot	Compiled with UEFI
0068F2_0NNJ2M	DEL2190110032	Mellanox ConnectX-4 Dual Port EDR PCIE Adapter LP	Yes	No
06W1HY_0JJN39	DEL2180110032	Mellanox ConnectX-4 Single Port EDR PCIE Adapter LP	Yes	No
0NHYP5_0XR0K2	DEL2150110033	ConnectX-4 EN network interface card; 100GbE dual-port QSFP28; PCIe3.0 x16; ROHS R6	Yes	No
00272F_0HWTYK	DEL3240110033	Mellanox ConnectX-4 Dual Port 100 GbE QSFP Network Adapter	Yes	Yes
DELL_C6320p_1P_EDR	DEL2010120032	Dell C6320p - ConnectX-4 adapter card; EDR IB (100Gb/s); single-port QSFP28; PCIe3.0 x16	Yes	No

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

1.2 Supported Cables and Modules

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
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

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
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
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

Table 4 - Validated and Supported 10/40GbE Cables

Speed	Cable OPN #	Description
10GB/S	MC3309130-00A	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 0.5M
10GB/S	MC3309130-0A1	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 1.5M
10GB/S	MC3309130-0A2	MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 2.5M
10GB/S	MFM1T02A-LR-F	MELLANOX OPTICAL MODULE ETH 10GBE 10GB/S SFP+ LC-LC 1310NM LR UP TO 10KM
10GB/S	MFM1T02A-SR-F	MELLANOX OPTICAL MODULE ETH 10GBE 10GB/S SFP+ LC-LC 850NM SR UP TO 300M
40GB/S	QSFP-40G-SR-BD	Cisco 40GBASE-SR-BiDi, duplex MMF
40GB/S	QSFP-40G-SR4	Cisco 40GBASE-SR4, 4 lanes, 850 nm MMF
40GB/S	QSFP-H40G-ACU10M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 10-meter, active
40GB/S	QSFP-H40G-AOC10M	Cisco 40GBase-AOC QSFP direct-attach Active Optical Cable, 10-meter
40GB/S	QSFP-H40G-CU1M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 1-meter, passive
40GB/S	QSFP-H40G-CU3M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 3-meter, passive
40GB/S	QSFP-H40G-CU5M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 5-meter, passive
10GB/S	SFP-10G-SR	Cisco 10GBASE-SR SFP+ transceiver module for MMF, 850-nm wavelength, LC duplex connector
10GB/S	SFP-H10GB-CU1M	Cisco 1-m 10G SFP+ Twinax cable assembly, passive
10GB/S	SFP-H10GB-CU3M	Cisco 3-m 10G SFP+ Twinax cable assembly, passive
10GB/S	SFP-H10GB-CU5M	Cisco 5-m 10G SFP+ Twinax cable assembly, passive

1.2.3 Validated and Supported QDR/FDR10 Cables

Table 5 - Validated and Supported QDR/FDR10 Cables

Speed	Cable OPN #	Description
QDR	MC2206125-007	MELLANOX PASSIVE COPPER CABLE IB QDR 40GB/S QSFP 7M
QDR	MC2206126-006	MELLANOX PASSIVE COPPER CABLE IB QDR 40GB/S QSFP 6M

Table 5 - Validated and Supported QDR/FDR10 Cables

Speed	Cable OPN #	Description
FDR10	MC2206128-004	MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 4M
FDR10	MC2206128-005	MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 5M
FDR10	MC2206130-001	MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 1M
FDR10	MC2206130-002	MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 2M
FDR10	MC2206130-003	MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 3M
FDR10	MC2206130-00A	MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 0.5M
FDR10	MC2206310-XXX	MELLANOX ACTIVE FIBER CABLE IB QDR/FDR10 40GB/S QSFP from 3M up to 100M
FDR10	MC2210411-SR4	MELLANOX OPTICAL MODULE 40GB/S QSFP MPO 850NM UP TO 100M
FDR10	MC2210411-SR4E	MELLANOX OPTICAL MODULE 40GB/S QSFP MPO 850NM UP TO 300M
FDR10	MFS4R12CB-XXX	MELLANOX ACTIVE FIBER CABLE VPI UP TO 40GB/S QSFP from 3M up to 100M

1.2.4 Validated and Supported 50Gbs Cables

Table 6 - Validated and Supported 50Gbs Cables

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

1.2.5 Validated and Supported FDR Cables

Table 7 - Validated and Supported FDR Cables

Speed	Cable OPN #	Description
FDR	MC2207126-004	MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 4M
FDR	MC2207128-003	MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 3M
FDR	MC2207128-0A2	MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 2.5M
FDR	MC2207130-001	MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 1M
FDR	MC2207130-002	MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 2M
FDR	MC2207130-00A	MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 0.5M
FDR	MC2207130-0A1	MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 1.5M
FDR	MC2207310-100	MELLANOX ACTIVE FIBER CABLE VPI UP TO 56GB/S QSFP from 3M up to 100M
FDR	MC2207310-XXX	MELLANOX ACTIVE FIBER CABLE VPI UP TO 56GB/S QSFP from 3M up to 100M
FDR	MC2207312-XXX	MELLANOX ACTIVE FIBER CABLE VPI UP TO 56GB/S QSFP from 3M up to 300M
FDR	MC220731V-XXX	Mellanox® Active Fiber Cable, VPI, up to 56Gb/s, QSFP, up to 100m
FDR	MC2207411-SR4L	MELLANOX OPTICAL MODULE IB FDR 56GB/S QSFP MPO 850NM UP TO 30M
FDR	MCP170L-F001	Mellanox® Passive Copper Cable, VPI, up to 56Gb/s, QSFP, LSZH, 1m
FDR	MCP170L-F002	Mellanox® Passive Copper Cable, VPI, up to 56Gb/s, QSFP, LSZH, 2m
FDR	MCP170L-F003	Mellanox® Passive Copper Cable, VPI, up to 56Gb/s, QSFP, LSZH, 3m

1.2.6 Validated and Supported EDR/100GB/s Cables

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

Speed	Cable OPN #	Description
100GB/S	MCP1600-C001	MELLANOX PASSIVE COPPER CABLE ETH 100GBE 100GBS QSFP LSZH 1M
100GB/S	MCP1600-C002	MELLANOX PASSIVE COPPER CABLE ETH 100GBE 100GBS QSFP LSZH 2M
100GB/S	MCP1600-C003	MELLANOX PASSIVE COPPER CABLE ETH 100GBE 100GBS QSFP LSZH 3M
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 ^a	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 1M
EDR	MCP1600-E002 ^a	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 2M
EDR	MCP1600-E003	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 3M
EDR	MCP1600-E00A ^a	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 0.5M
EDR	MCP1600-E01A ^a	Mellanox® Passive Copper cable, VPI, up to 100Gb/s, QSFP, LSZH, 1.5m
EDR	MCP1600-E02A	Mellanox® Passive Copper cable, VPI, up to 100Gb/s, QSFP, LSZH, 2.5m
100GB/S	MFA1A00-C005	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 5m
100GB/S	MFA1A00-C010	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 10m
100GB/S	MFA1A00-C015	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 15m
100GB/S	MFA1A00-C020	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 20m
100GB/S	MFA1A00-C030	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 30m
100GB/S	MFA1A00-C050	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 50m

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

Speed	Cable OPN #	Description
100GB/S	MFA1A00-C100	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 100m
EDR	MFA1A00-E005 ^a	MELLANOX active fiber cable, VPI, up to 100Gb/s, QSFP, 5m
EDR	MFA1A00-E010 ^a	MELLANOX active fiber cable, VPI, up to 100Gb/s, QSFP, 10m
EDR	MFA1A00-E015 ^a	MELLANOX active fiber cable, VPI, up to 100Gb/s, QSFP, 15m
EDR	MFA1A00-E020	MELLANOX active fiber cable, VPI, up to 100Gb/s, QSFP, 20m
EDR	MFA1A00-E030	MELLANOX active fiber cable, VPI, up to 100Gb/s, QSFP, 30m
EDR	MFA1A00-E050	MELLANOX active fiber cable, VPI, up to 100Gb/s, QSFP, 50m
EDR	MFA1A00-E100	MELLANOX active fiber cable, VPI, up to 100Gb/s, QSFP, 100m
100GB/S	MFS1200-C005	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 5m
100GB/S	MFS1200-C010	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 10m
100GB/S	MFS1200-C015	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 15m
100GB/S	MFS1200-C020	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 20m
100GB/S	MFS1200-C030	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 30m
100GB/S	MFS1200-C050	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 50m
100GB/S	MFS1200-C100	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 100m
EDR	MFS1200-E005	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 5m
EDR	MFS1200-E010	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 10m
EDR	MFS1200-E015	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 15m
EDR	MFS1200-E020	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 20m
EDR	MFS1200-E030	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 30m
EDR	MFS1200-E050	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 50m

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

Speed	Cable OPN #	Description
EDR	MFS1200-E100	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 100m
100GB/S	MMA1B00-C100D	Mellanox® transceiver, 100GbE, QSFP28, MPO, 850nm, up to 100m
EDR	MMA1B00-E100	Mellanox® transceiver, IB EDR, up to 100Gb/s, QSFP28, MPO, 850nm, up to 100m
100GB/S	QSFP-100G-AOC-10M	Cisco 100GBase QSFP Active Optical Cable, 10-meter

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

1.3 Tested Switches

1.3.1 Tested QDR Switches

Table 9 - Tested QDR Switches

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

1.3.2 Tested 10/40GbE Switches

Table 10 - Tested 10/40GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
10/40GbE	N/A	3064	48-port 10Gb/40Gb Switch	Cisco
10GbE	N/A	5548	Cisco 10GB ETH switch	Cisco
40GbE	N/A	3132Q	Cisco 40GB ETH switch	Cisco
10/40GbE	N/A	7050Q	16-port 40Gb Switch	Arista
40GbE	N/A	7050QX	32-port 40Gb Switch	Arista
10/40GbE	N/A	7050S	48-port 10Gb/40Gb Switch	Arista

Table 10 - Tested 10/40GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
10GbE	N/A	G8264	BNT 10/40GB ETH switch	BNT
40GbE	N/A	G8316	BNT 40GB RackSwitch G8316	BNT
10GbE	N/A	QFX3500	Juniper 10/40GB ETH switch	Juniper
10GbE	N/A	S4810P-AC	48-port 10Gb/40Gb Switch	Force10
40GbE	N/A	S6000	32-port 40Gb Switch	Dell
10GbE	SwitchX [®]	SX1016X-1BFR	64-Port 10GbE Switch System	Mellanox
40GbE	SwitchX [®]	SX1036B-1BFR	36-Port 40/56GbE Switch System	Mellanox

1.3.3 Tested FDR Switches

Table 11 - Tested FDR Switches

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

1.3.4 Tested 100GbE/EDR Switches

Table 12 - Tested 100GbE/EDR Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
100Gb/s	N/A	7060CX	32-port 100Gb Switch	Arista
100Gb/s	N/A	93180YC-EX	48 x 10/25-Gbps fiber ports and 6 x 40/100-Gbps Quad Small Form-Factor Pluggable 28 (QSFP28) ports	Cisco
100Gb/s	N/A	C3232C	High-Density, 100 Gigabit Ethernet Switch	Cisco
100Gb/s	N/A	CE8860-4C-EI	24x10GE (SFP+) or 25GE (SFP28) and 2x100GE switch	Huawei
EDR	Switch-IB	SB7790-EB2F	36-port EDR 100Gb/s InfiniBand Switch Systems	Mellanox
EDR	Switch-IB 2	SB7800-ES2R	36-port Non-blocking Managed EDR 100Gb/s InfiniBand Smart Switch	Mellanox
100GbE	Spectrum	SN2410-CB2F	48-port 25GbE + 8-port 100GbE Open Ethernet ToR Switch System	Mellanox
100GbE	Spectrum	SN2700-CS2R	32-port Non-blocking 100GbE Open Ethernet Spine Switch System	Mellanox

1.4 Tools, Switch Firmware and Driver Software

Firmware Rev 12.17.2052 is tested with the following tools, Switch firmware, and driver software:

Table 13 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_OFED	3.4-1.0.0.0/3.3-1.0.0.0

Table 13 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_EN (MLNX_OFED based code)	3.4-1.0.0.0/3.3-1.0.0.0
WinOF-2	1.45/1.41
MFT	4.6.0/4.5.0
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"> • Ubuntu 14.04.3 • Ubuntu 14.04.4 • Ubuntu 15.04 • Ubuntu 15.10 • Ubuntu 16.04 • Ubuntu 16.04.1 • Ubuntu 16.10 • SLES12 • SLES12.1 • SLES12.2 • RHEL6.6 • RHEL6.7 • RHEL6.8 • RHEL7.1 • RHEL7.2 • RHEL7.3
Windows Inbox Drivers	Windows Server 2016 Windows Server 2012 R2 Windows Server 2012

1.5 Supported FlexBoot



Please be aware that not all firmware binaries contain FlexBoot (support may vary between cards, see [Section 1.1, “Supported Devices”](#), on page 6).

Firmware Rev 12.17.2052 supports the following FlexBoot:

Table 14 - Supported FlexBoot

Expansion ROM	Supported Version
FlexBoot	3.4.903
UEFI	14.11.34

1.6 Revision Compatibility

Firmware Rev 12.17.2052 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 12.17.2052

Table 15 - Changes and New Features in Rev 12.17.2052

Feature/Change	Description
Bug fixes	See Section 4, “Bug Fixes History” , on page 28

3 Known Issues

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

Table 16 - Known Issues (Sheet 1 of 8)

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

Table 16 - Known Issues (Sheet 2 of 8)

Internal Ref.	Issue
-	Description: When the link is Gen2, entering or exiting L1 state may cause bad CRC or DLLP indication.
	Workaround: N/A
	Keywords: PCI
600534	Description: Configuration of space power management capability <code>PME_EN</code> cannot be set, thus preventing the driver from activating the wake signal.
	Workaround: N/A
	Keywords: PCIe
685062	Description: MultiHost InfiniBand: OpenSM is supported over host0 only and the <code>MAD_IFC</code> usage is limited to host0 only.
	Workaround: Activate OpenSM and the MFT tools via host0
	Keywords: PCI
687113	Description: PF direct pass-through is not supported in InfiniBand (since PF FLR is not supported)
	Workaround: N/A
	Keywords: PF direct pass-through, InfiniBand
591240	Description: Traffic that is loopbacked due to <code>QP.force_loopback</code> being equaled to 1, is steered to the PF.
	Workaround: N/A
	Keywords: Ethernet Network
594964	Description: A minimum of 200 LFM is required in order to cool the MCX4411A-ACAN adapter card.
	Workaround: N/A
	Keywords: Temperature
601485/ 599810	Description: <code>mlxfwreset</code> does not function properly in old MFT versions after upgrading the firmware image.
	Workaround: Upgrade MFT to the latest release or use <code>reboot/power cycle</code> after upgrading firmware.
	Keywords: Firmware Tool

Table 16 - Known Issues (Sheet 3 of 8)

Internal Ref.	Issue
-	<p>Description: Older MFT versions (4.0.0 and 3.8.0) may indicate that the latest GA firmware is old or that it cannot be compared with the existing firmware. A message similar to the below will be displayed upon firmware upgrade stage: # 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>Workaround: Choose one of the options below to upgrade firmware:</p> <ul style="list-style-type: none"> • Upgrade to the latest MFT version (4.1.0) • Type "y" after the note flint provides <p>Run flint with the "-force" flag</p> <p>Keywords: Firmware Upgrade/MFT</p>
-	<p>Description: Windows Server 2016 Inbox driver cannot work with firmware v12.12.0780</p> <p>Workaround: Use WinOF-2 v1.20 out-of-box driver</p> <p>Keywords: Windows Inbox Drivers</p>
655688	<p>Description: When arming SRQ for limit event, the device might issue an event with context_index=0.</p> <p>Workaround: N/A</p> <p>Keywords: RoCE</p>
-	<p>[For customers developing custom low level drivers] Description: VFs internal FLR is not supported in PF teardown HCA command.</p> <p>Workaround: Before unloading the PF driver, PF driver must disable all its active VFs by performing the following:</p> <ol style="list-style-type: none"> 1. Run the disable_hca command on all the function_ids 2. Wait until firmware returns all VFs allocated pages. <p>Keywords: Virtualization, FLR</p>
-	<p>Description: Function (PF/VF) TX port counters are not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>
-	<p>Description: PF driver must work with pages event queue.</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>

Table 16 - Known Issues (Sheet 4 of 8)

Internal Ref.	Issue
597718	Description: Privileged Vport egress traffic is not blocked when Vport is not active
	Workaround: 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.
	Workaround: N/A
	Keywords: Virtualization
597718	Description: Vport number in virtual trap might be reported incorrectly
	Workaround: N/A
	Keywords: Virtualization
691387/ 691415	Description: In a Multihost setup, when running a single TCP stream, you might experience sub optimal throughput.
	Workaround: Use multiple streams to reach optimal results
	Keywords: Multihost setup, Performance, TCP stream
691754	Description: <code>end_padding_mode</code> is required in <code>CREATE_QP</code> and not in <code>INIT_2_RTR</code> command as defined in the PRM
	Workaround: N/A
	Keywords: <code>end_padding_mode</code> , PRM
691490	Description: LR4 cable events are sent although the port is up
	Workaround: N/A
	Keywords: Management
-	Description: QoS must be configured the same for both ports in order for RoCE LAG to function properly.
	Workaround: N/A
	Keywords: RoCE LAG
759571/ 759655	Description: Modifying the <code>encap_id</code> of FTE is not supported.
	Workaround: N/A
	Keywords: <code>encap_id</code> , FTE
756872/ 769604	Description: Flow Counter is supported only for FTE that does not include a <code>flow_tag</code> or for FTE that have TIR as a destination.
	Workaround: N/A
	Keywords: Flow Counter, FTE

Table 16 - Known Issues (Sheet 5 of 8)

Internal Ref.	Issue
756871/ 770208	Description: Using Flow Counters in the FDB Flow Table causes the transmitted IB traffic vport counters not to function properly.
	Workaround: N/A
	Keywords: Flow Counter, FDB Flow Table, vport counters
756870/ 769605	Description: Using Flow Counters in the FDB Flow Table may harm vport counters' clearing functionality.
	Workaround: N/A
	Keywords: Flow Counter, FDB Flow Table, vport counters
748292	Description: When a steering rule in the e-sw FDB includes an encap action and an external port as destination, a transmitted multicast packet that matches the rule is sent to the wire and the loopback and the locally looped back packet will also have an encap header.
	Workaround: N/A
	Keywords: FDB, multicast packet
747967/ 771507	Description: Burning firmware on the same device in parallel from multiple interfaces (e.g. PCIe and MTUSB) is not supported.
	Workaround: N/A
	Keywords: PCIe, MTUSB, burning in parallel
754914	Description: When e-switch FDB is not created, the VF functional loopback traffic is send to vport 0 (PF).
	Workaround: N/A
	Keywords: e-switch FDB, vport, SR-IOV
690890	Description: Updating a non-volatile configuration of port type TLV more than 50 times might cause system to hang.
	Workaround: Run <code>mlxconfig reset</code> after every 50 consecutive updates of port type TLV.
	Keywords: Non-volatile configuration, TLV
783742	Description: In order to raise 50GbE link when using ConnectX-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)
	Workaround: N/A
	Keywords: MLNX-OS, 50G link

Table 16 - Known Issues (Sheet 6 of 8)

Internal Ref.	Issue
776830	<p>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.</p> <p>Workaround: Cold reboot the device instead</p> <p>Keywords: Warm/cold reboot</p>
770824	<p>Description: Pressing the Power Down button resets the server and does not initiate the Standby flow (as init 0 does). As a result, both ports are up due to <code>keep_link_up</code>, which opens the port when the firmware is loaded.</p> <p>Workaround: Use init 0 to start the Standby flow.</p> <p>Keywords: Warm/cold reboot</p>
693832	<p>Description: In an InfiniBand Multihost and SR-IOV setups, traffic should contain GRH (GID index) if the <code>grh_required</code> bit is set in the <code>query_hca_vport_context</code> command.</p> <p>OpenSM should be configured as follow (<code>opensm.conf</code>):</p> <ul style="list-style-type: none"> • <code>virt_enable</code> should be 2 • Enable Qos: • <code>qos TRUE</code> <p>Note: In this case, traffic without GRH will be forwarded to vport0 ("Host0")</p> <p>Workaround: N/A</p> <p>Keywords: Multihost/SR-IOV setups</p>
778257	<p>Description: Performing warm reboot during firmware image burning in VPI/IB devices configured with IB port protocol, might cause the device to disappear from the PCIe.</p> <p>Workaround: Power Cycle the server (cold reboot). Once a cold reboot is performed, the device will reboot with the previous image that was already burned.</p> <p>Keywords: Warm reboot, firmware image burning, VPI/IB devices</p>
758803	<p>Description: The firmware and the hardware do not reset the physical link upon <code>CPortState=down</code>.</p> <p>According to the IB Specification, MANAGEMENT STATE CHANGE COMMANDS: "<i>CPortState... when <code>phy_link=up</code> and <code>CPortState=down</code>, the state machine will transition to the LinkDown state which will reset other link state machines. Since <code>phy_link=up</code>, this will be followed by a transition to the LinkInitialize state. Thus a command to change link port state to down provides a way to re-initialize the link layer...</i>"</p> <p>Workaround: In order to re-train the physical link, <code>sendbug PortInfo.physical_port_state = POLLING</code> is required.</p> <p>Keywords: Physical link, <code>CPortState=down</code></p>

Table 16 - Known Issues (Sheet 7 of 8)

Internal Ref.	Issue
840738	Description: Local loopback traffic might effect vport counters.
	Workaround: N/A
	Keywords: Local loopback traffic, vport counters
862480	Description: Stopping the Rate Limiter while traffic is being transmitted might cause the adapter card to hang.
	Workaround: Stop traffic before: <ul style="list-style-type: none"> • setting rate 0 to the last non-default-rate vport. • issuing <code>destroy_scheduling_element</code> command for the last vport with non-default-rate vport.
	Keywords: Rate Limiter
852744	Description: Mapping an SL to VL 15 is currently not supported. Trying to do so, will cause a health buffer fatal internal error report.
	Workaround: N/A
	Keywords: SL to VL mapping
864200	Description: Running the <code>modify_scheduling_context</code> command does not include checking whether the scheduling element was created or not.
	Workaround: Do not modify non-existing elements
	Keywords: SR-IOV Rate Limiter
854805	Description: Setting/modifying the <code>max_average_bw</code> rate for a function, or setting speeds over the maximum supported speed (as indicated in INI) may result in inaccurate rates, and in an assert.
	Workaround: Set the <code>max_avergae_bw</code> in <code>scheduling_context</code> commands to equal or less than the supported wire speed.
	Keywords: Bandwidth rate, speed
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>TEAR-DOWN_HCA/ FLR VF</code> regardless of the <code>max_average_bw</code> value set.
	Workaround: Verify all rates for all VFs are set to "0" before running <code>TEAR-DOWN_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.
	Workaround: N/A
	Keywords: FDR link, EDR cable

Table 16 - Known Issues (Sheet 8 of 8)

Internal Ref.	Issue
867367	<p>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.</p> <p>Workaround: Use the correct <code>element_type</code> when issuing <code>modify_scheduling_context</code> command.</p> <p>Keywords: SR-IOV Rate Limiter</p>
902828	<p>Description: When using a firmware based LLDP/DCBX software based, LLDP tools (such as <code>lldptool</code> in Linux) should be disabled. When intending to use software based LLDP, firmware LLDP must be disabled by using <code>mlxconfig</code>. Using both the LLDP software and the firmware based LLDP will result in an unexpected results. This applies to both Physical Functions (Bare Metal OS) and Virtual Functions.</p> <p>Workaround: Disable the LLDP software.</p> <p>Keywords: LLDP/DCBX</p>

4 Bug Fixes History

Table 17 lists the bugs fixed in this release.

Table 17 - Bug Fixes History (Sheet 1 of 12)

Internal Ref.	Issue
937018	Description: Fixed a rare issue/race condition that might cause a double command completion or other corrupted completion.
	Keywords: Double command completion, corrupted completion
	Discovered in Release: 12.17.2044
	Fixed in Release: 14.17.2052
996056	Description: Disabled the host management for the following adapter cards: 0068F2_0NNJ2M, 06W1HY_0JJN39
	Keywords: Host management
	Discovered in Release: 12.17.2044
	Fixed in Release: 14.17.2052
996055	Description: Fixed the NC-SI functionality for the following adapter cards: 0068F2_0NNJ2M, 06W1HY_0JJN39
	Keywords: NC-SI
	Discovered in Release: 12.17.2044
	Fixed in Release: 14.17.2052
893261	Description: Fixed the PCIe TX burst at Rec.speed_change state.
	Keywords: PCIe TX Burst
	Discovered in Release: 12.17.2044
	Fixed in Release: 14.17.2052
893261	Description: Fixed the PCIe TX ramp up time at Polling.active state.
	Keywords: PCIe TX Ramp
	Discovered in Release: 12.17.2044
	Fixed in Release: 14.17.2052
963165	Description: Fixed the NCSI LLDP command to increase the maximal allowed data size.
	Keywords: NCSI, LLDP, data size
	Discovered in Release: 12.17.1010
	Fixed in Release: 14.17.2044
959470 902405 947016	Description: Fixed Dynamically Connected Transport issues.
	Keywords: Dynamically Connected Transport
	Discovered in Release: 12.17.1010
	Fixed in Release: 14.17.2044

Table 17 - Bug Fixes History (Sheet 2 of 12)

Internal Ref.	Issue
866181	Description: Fixed an issue which caused system fail when enabled SR-IOV.
	Keywords: SR-IOV
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.17.2020
883830	Description: Fixed the NC-SI "set" Flow Control so it could always report "unsupported command" when operating the smbus.
	Keywords: NC-SI "set" Flow Control
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.17.2020
855533	Description: Fixed ODP flow issues that caused occasional fatal error reporting and RX hanging.
	Keywords: ODP flow
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.17.2020
860574	Description: Fixed performance issues to improve Packet Pacing performance.
	Keywords: Performance, Packet Pacing
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.17.2020
883834	Description: Information is now received per channel port instead of the lowest port which received LLDP.
	Keywords: LLDP
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.17.2020
883834	Description: Changed the maximum TLV size per each TLV to 28B instead of 32B.
	Keywords: TLV size
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.17.2020
883834	Description: Fixed an issue which added a header to the packet which already contained header and data.
	Keywords: Packet headers
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.17.2020

Table 17 - Bug Fixes History (Sheet 3 of 12)

Internal Ref.	Issue
883834	Description: Enabled LLDP by default in the ini after adding the MCTP command.
	Keywords: LLDP, MCTP command
	Discovered in Release: 12.17.1010
	Fixed in Release: 12.17.2020
826702	Description: Fixed an issue which caused RX to hang when a UDP packet with destination port of RoCE v2 arrived and the data matched the DC transport service.
	Keywords: RoCE v2, DC Transport, UDP, RX
	Discovered in Release: 12.16.1010
	Fixed in Release: 12.17.1010
802148	Description: Fixed an issue which caused the link not to come up when the port was toggled in a rapid frequency.
	Keywords: Link up
	Discovered in Release: 12.16.1010
	Fixed in Release: 12.17.1010
736528	Description: On rare occasions during UEFI boot cycles system got stuck while WinPE is loaded. (OS WinPE, system DL160).
	Keywords: WinPE OS, UEFI boot cycles
	Discovered in Release: 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

Table 17 - Bug Fixes History (Sheet 4 of 12)

Internal Ref.	Issue
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
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.
	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 <code>target_link_speed</code> to Gen3 in link control2.
	Key Words: PCIe, Link Speed
	Discovered in Release: 12.14.1006
	Fixed in Release: 12.17.1010

Table 17 - Bug Fixes History (Sheet 5 of 12)

Internal Ref.	Issue
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
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 - Bug Fixes History (Sheet 6 of 12)

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 - Bug Fixes History (Sheet 7 of 12)

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
735315	Description: Reduced Dell card from 32MB VF UAR size to 1MB.
	Keywords: VF UAR size
	Discovered in Release: 14.14.1100
	Fixed in Release: 14.14.2036
684496	Description: Fixed a rare issue which caused the command to hang when moved the QP to RESET and back to RTS.
	Keywords: QP, RTS
	Discovered in Release: 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

Table 17 - Bug Fixes History (Sheet 8 of 12)

Internal Ref.	Issue
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
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 <code>qkey/pkey</code> violation counter via <code>port_info</code> mad.
	Keywords: <code>qkey/pkey</code> violation counter, <code>port_info</code> 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

Table 17 - Bug Fixes History (Sheet 9 of 12)

Internal Ref.	Issue
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
608549	Description: Fixed a link issue which occurred due to wrong port configuration in the INI.
	Keywords: INI, port configuration
	Discovered in Release: 14.12.1240
	Fixed in Release: 14.14.1100
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

Table 17 - Bug Fixes History (Sheet 10 of 12)

Internal Ref.	Issue
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
664528	Description: Modified Dell' iSCSI NC-SI response and reason codes.
	Keywords: iSCSI NC-SI
	Discovered in Release: 14.12.1240
	Fixed in Release: 14.14.1100
657680	Description: Fixed casting of BMC MAC before steering API.
	Keywords: BMC, Steering API
	Discovered in Release: 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

Table 17 - Bug Fixes History (Sheet 11 of 12)

Internal Ref.	Issue
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
-	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

Table 17 - Bug Fixes History (Sheet 12 of 12)

Internal Ref.	Issue
-	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
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 (Sheet 1 of 8)

Feature/Change	Description
Rev. 12.17.2044	
Bug fixes	See Section 4, “Bug Fixes History” , on page 28
Rev. 12.17.2036	
UEFI	See Section 7, “UEFI Changes and Major New Features” , on page 56.
Rev. 12.17.2020	
NC-SI CMD	Enabled NC-SI CMD to report the maximum temperature constantly.
PCIe-VDM	Enabled PCIe-VDM for ConnectX®-4 Lx Dual Port 25 GbE SFP Network Adapter. Dell OPN: 020NJD
Wake-on-LAN (WoL)	Disabled WoL by default.
GENEVE & IP-in-IP Stateless Offload	[Beta] Added support for IP-in-IP and GENEVE network protocols encapsulated into IP frame (L2 tunneling). Encapsulation is suggested as a means to alter the normal IP routing for datagrams, by delivering them to an intermediate destination that would otherwise not be selected based on the (network part of the) IP Destination Address field in the original IP header. Note: For driver support, please see the Release Notes/User Manual of the relevant OS driver.
Bug Fixes	See Section 4, “Bug Fixes History” , on page 28
Rev. 12.17.1010	
Multi-Host LID Base Routing	Added support for Multi-Host LID base routing. This feature requires a new OpenSM (v4.7.1 and above which comes with MLNX_OFED 3.3-2.0.0.0) with the following attributes: <ul style="list-style-type: none"> • qos TRUE • lmc 2 (if there is no quad host in the fabric, you can set the lmc to 1) • virt_enabled 2 Note: Multi-Host LID base routing can be configured by the INI only. The default is 0
Resilient RoCE	Resilient RoCE is the ability to send RoCE traffic over a lossy network (a network without flow control enabled), without the need to enable flow control on the network. The ability is accomplished by enabling ECN on both the Switch and the Host.
Multi-Host L3/L4 Classification	Enables load balancing in the Multi PF Switch layer (MPFS) based on the L3/L4 headers

Table 18 - Firmware Changes and New Feature History (Sheet 2 of 8)

Feature/Change	Description
InfiniBand Multi-Host Isolation	Enabled isolation between separate Hosts using the same HCA. All the Hosts can be rebooted, the driver can be stopped and the FLR signal can be sent independently.
95 Virtual Functions (VF) per Port	<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/multihost 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)	Added support for QP Rate Limit in InfiniBand.
HCA Port Flap Counter	Added support for Port Flap Counter.
Fixed Buffer Size (KSM)	Limits the buffer size for all entries to improve performance. KSM is used when associating Key Length My Virtual Address (KLMs) with fixed memory size.
NULL Mkey	This entry (null_mkey) is use to indicate non-present KLM/KSM entries. When accessing is, it causes the device to generate page fault event.
Out-of-Band Online Firmware Update: Firmware Update over PLDM	PLDM firmware burning is based on the DMTF spec DSP0267 (draft 9). The feature enables upgrading firmware and expansion ROM images using the PLDM protocol over MCTP (over PCIe). By doing so, a supporting BMC can query and upgrade the firmware without using OS based tools.
New Group in Ports Performance Counters (PPCNT)	<p>Added a new physical layer statistics counters group. The new group includes BER counters, FEC error correction, clear time, and additional physical layer counters.</p> <p>For further information, please refer to the Ethernet Adapters Programming Manual (PRM).</p>
Permanent Link Up Mode	<p>Enables the user to set a certain link up state for an unlimited period of time. This mode has 3 states:</p> <ul style="list-style-type: none"> Aux power (standby) Reboot/boot/driver unloaded - the server is active and no driver is up Driver is up - at least one driver is up (the time between init HCA and teardown or FLR)

Table 18 - Firmware Changes and New Feature History (Sheet 3 of 8)

Feature/Change	Description
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 28
Rev. 12.16.1020	
Bug Fixes	See Section 4, “Bug Fixes History” , on page 28
Rev. 12.16.1006	
Explicit Congestion Notification (ECN)	[Beta] Explicit Congestion Notification (ECN) is an extension to the Internet Protocol and to the Transmission Control Protocol. ECN allows end-to-end notification of network congestion without dropping packets.
64 VFs per port	Increased the number of VFs from 32 to 64 per PF. Note: When increasing the number of VFs, the following limitations must be taken into consideration: <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>

Table 18 - Firmware Changes and New Feature History (Sheet 4 of 8)

Feature/Change	Description
RoCE Link Aggregation (RoCE LAG)	<p>[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.</p> <p>For further information, please refer to the PRM.</p>
OVS Offload	<p>Mellanox Accelerated Switching And Packet Processing (ASAP²) 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 ASAP² Direct include packet parsing and matching, forward, drop along with VLAN push/pop or VXLAN encaps/decap and HW based packet/byte flow statistics.</p>
Virtual Extensible LAN (VXLAN) encapsulation/decapsulation	<p>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.</p>
Data Center Bridging Exchange (DCBX)	<p>[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.</p>
FCS no scatter / FCS check	<p>Enables the user to control whether or not to scatter Frame Check Sequence (FCS) or to check FCS functionality.</p>
Packet Pacing	<p>[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.</p>
PRBS Patterns Generation and Tuning	<p>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.</p>
Management Controller Transport Protocol (MCTP) over PCI	<p>Added support for MCTP host management over PCI</p>
OCBB / OCSD support after mlxfwreset	<p>Added support for OCBB/OCSD memory pointers restoration after mlxfwreset</p>
MCTP media migration	<p>Added support for MCTP media migration between SMBUS and PCI</p>
Cables	<p>Removed the RX amplitude configuration on some cable types</p>

Table 18 - Firmware Changes and New Feature History (Sheet 5 of 8)

Feature/Change	Description
Bug Fixes	See Section 4, “Bug Fixes History” , on page 28
Rev. 12.14.2036	
VPD Read	Added v1, v3, v6 tags to VPD read only tag
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 28
Rev. 12.14.1100	
CQE Time Stamping	Keeps track of the creation of a packet. A time-stamping service supports assertions of proof that a datum existed before a particular time.
Priority Flow Control (PFC)	Applies pause functionality to specific classes of traffic on the Ethernet link.
RDMA retransmission counters	Custom port counters provide the user a clear indication about RDMA send/receive statistics and errors.
Link Layer Discovery Protocol (LLDP)	The Link Layer Discovery Protocol (LLDP) is a vendor-neutral Link Layer protocol in the Internet Protocol Suite used by network devices for advertising their identity, capabilities, and neighbors on a IEEE 802 LAN. The protocol is formally defined in IEEE 802.1AB.
1GbE and 56GbE Link Speed	ConnectX-4adapters now support 1Gb/s and 56GbE Ethernet connectivity in addition to 10GigE, 25GigE, 40GigE, 50GigE, and 100GigE
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.

Table 18 - Firmware Changes and New Feature History (Sheet 6 of 8)

Feature/Change	Description
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: <ol style="list-style-type: none"> 1. Changed port configuration which required link re-training (such as speed) 2. PAOS down 3. PAOS up This change, will cause the link to toggle and new configurations to take effect.
Checksum Calculation on Image/Device	Flint utility allows performing an MD5 checksum on the non-persistent sections of the firmware image. For further information, please refer to MFT User Manual .
Rev. 12.12.1240	
Bug Fixes	See Section 4, “Bug Fixes History”, on page 28
Rev. 12.12.1100	
Port Link	Reduced the port link-up time when negotiating according to Clause 73 (DME)
Rev. 12.12.0780	
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

Table 18 - Firmware Changes and New Feature History (Sheet 7 of 8)

Feature/Change	Description
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 9, on page 58.
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
Configuration tools	<ul style="list-style-type: none"> Enabled SR-IOV, NUM_VFS and INT_LOG_MAX_PAYLOAD_SIZE configuration via the <code>mlxconfig</code> 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

Table 18 - Firmware Changes and New Feature History (Sheet 8 of 8)

Feature/Change	Description
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. <p>RoCE is supported only in Reliable Connection (RC) transport</p> <ul style="list-style-type: none"> • 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
Host management	<ul style="list-style-type: none"> • NC-SI over RMIi support
MAD	<ul style="list-style-type: none"> • Config space address in MAD management class 0x09

6 FlexBoot Changes and New Features

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

Table 19 - FlexBoot Changes and New Features

Version	Description
3.4.903	<ul style="list-style-type: none"> • iSCSI re-imaging: Enables the user to install a new image on active iSCSI target • FlexBoot UI: Enabled boot configuration menu in ConnectX-4 when the physical port is IB • FlexBoot UI: Added new configuration for network link type for supported cards (ConnectX-4 VPI cards). • Booting: Enabled booting with non-default Pkey in ConnectX-4 when the physical port is IB • Link Status: Removed link status line printout at boot time • Upstream sync: Synced the source with iPXE (upstream sync) • Boot Menu: Changed the Bus:Device:Function format in boot menu, from PCI-Bus:Dev.Func to 0000:Bus:Dev.Func.
3.4.813	<ul style="list-style-type: none"> • FlexBoot UI: Added debug prints option in the FlexBoot boot menu. For further information, please refer to FlexBoot and UEFI User Manual. • System Diagnosis: Added the ability to diagnose problems in released ROMs by enabling the debug log levels for specific modules. Note: This ability should be used only when debug session is needed. • Interrupts: Added support for ConnectX-4/ConnectX-4 Lx interrupts. • Upstream sync: Synced the source with iPXE (upstream sync)
3.4.722	<ul style="list-style-type: none"> • Bug Fixes, see Section 6.2, “FlexBoot Bug Fixes History”, on page 54
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 • For further information, please refer to: http://ipxe.org/cmd • Synced the source with iPXE (upstream sync)
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 (Sheet 1 of 5)

Internal Ref.	Issue
-	Description: Several BIOS vendors have limited boot-vector space and may not display FlexBoot in their boot menu.
	Workaround: Disable the embedded NIC boot agent in BIOS
	Keywords: BIOS
-	Description: In several BIOS, the server might hang during FlexBoot booting due to wrong configuration of the PMM.
	Workaround: N/A
	Keywords: BIOS
-	Description: Only EBX, ESI, DS, ES registers can be saved in Boot Entry.
	Workaround: N/A
	Keywords: BIOS
-	Description: If a client returned control to the BIOS after a successful connection to an iSCSI target (but did not boot from it), then, unexpected behavior may occur.
	Workaround: Follow the instructions described in the FlexBoot UM for the proper iSCSI boot/install
	Keywords: BIOS
673114	Description: FlexBoot banner might not be shown in some BIOSes.
	Workaround: N/A
	Keywords: BIOS
-	Description: In some cases, PXE boot will not work if the client was given only the filename without next-server (siaddr).
	Workaround: N/A
	Keywords: 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.
	Workaround: N/A
	Keywords: PXE Boot
-	Description: PXE boot after iSCSI boot with static configuration is currently not supported.
	Workaround: N/A
	Keywords: PXE Boot

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

Internal Ref.	Issue
-	Description: Boot over VLAN with IB port is currently not supported.
	Workaround: N/A
	Keywords: PXE Boot
-	Description: Some faulty boot loaders do not close the underlying UNDI device which may result in unexpected behavior and possible system crash after the OS starts to load.
	Workaround: N/A
	Keywords: PXE Boot
-	Description: Chain-loading gPXE stack may result in undesirable behavior.
	Workaround: N/A
	Keywords: PXE Boot
647143	Description: Executing a partial boot loop while only downloading the NBP and selecting localboot is unsupported and may cause undefined behavior.
	Workaround: N/A
	Keywords: PXE Boot
670421	Description: Using filename for PXE boot with rootpath for hooking an iSCSI target (to install) is not supported when the PXE boot loader uses UNDI API, since all traffic must get to the boot loader.
	Workaround: N/A
	Keywords: PXE Boot
-	Description: iSCSI over IB is not tested.
	Workaround: N/A
	Keywords: iSCSI
-	Description: iSCSI over DCB is not supported.
	Workaround: N/A
	Keywords: iSCSI
-	Description: FlexBoot supports only a single active iSCSI connection. Thus, when iSCSI-boot via Port 1 succeeds to connect but fails to boot, it will fail to connect via Port 2.
	Workaround: N/A
	Keywords: iSCSI
-	Description: Boot retries is currently not functional when booting from iSCSI.
	Workaround: N/A
	Keywords: iSCSI

Table 20 - FlexBoot Known Issues (Sheet 3 of 5)

Internal Ref.	Issue
655800	Description: IPv6 is not supported.
	Workaround: N/A
	Keywords: iSCSI
-	Description: Boot menu is displayed as READ ONLY if the HCA card does not support flash configuration.
	Workaround: N/A
	Keywords: User Interface
-	Description: FlexBoot Boot Menu will not be visible in serial output.
	Workaround: N/A
	Keywords: User Interface
-	Description: FlexBoot Boot Menu will not be shown in ConnectX-4 if the physical port is IB.
	Workaround: N/A
	Keywords: User Interface
-	Description: Large Receive Offload (LRO) and iSCSI may not interoperate due to a bug in current Linux kernel distributions.
	Workaround: Disable LRO in the IPoIB module when using iSCSI. See the Mellanox FlexBoot user's manual for details under the Diskless Machines chapter (InfiniBand Ports).
	Keywords: Networking
-	Description: Flexboot supports only 2K MTU.
	Workaround: N/A
	Keywords: Networking
-	Description: 56Gb/s is currently not supported.
	Workaround: N/A
	Keywords: Link Speed
-	Description: Blink LEDs are currently not supported.
	Workaround: N/A
	Keywords: LED
-	Description: Setting the number of Virtual Functions higher than the machine's memory capability may cause memory issues and system instability.
	Workaround: N/A
	Keywords: Virtualization

Table 20 - FlexBoot Known Issues (Sheet 4 of 5)

Internal Ref.	Issue
-	<p>Description: SLAM, FTP, HTTPS and SRP are currently not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Protocols</p>
-	<p>Description: Occasionally, using the Spanning Tree Protocol (STP) in the switches may cause packet drops and boot failure in the system.</p> <p>Workaround: Enable the "edgemode" if disabled on the switch, or use either port-fast or edgemode functionality on the switch ports connected to the NICs.</p> <p>Keywords: Protocols</p>
-	<p>Description: FCoE, BCV are not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Protocols</p>
655800	<p>Description: IPv6 can only run if a RADVD service is running in the network.</p> <p>Workaround: N/A</p> <p>Keywords: Protocols</p>
-	<p>Description: IPv6 over IB is not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Protocols</p>
655800	<p>Description: IPv6 over WDS is not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Protocols</p>
655800	<p>Description: Enabling IPv6 first and then IPv4 is currently not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Protocols</p>
656001	<p>Description: Booting from WDS and Windows DHCP server when only Option 66 is enabled (without Option 67), is not supported.</p> <p>Workaround: N/A</p> <p>Keywords: DHCP</p>
735159	<p>Description: LACP packets handling is not supported during boot.</p> <p>Workaround: N/A</p> <p>Keywords: PXE boot</p>
735159	<p>Description: Interrupts are not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Interrupts</p>

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

Internal Ref.	Issue
-	Description: As ConnectX-4 and Connect-IB HCAs do not support FlexBoot menu, pressing Ctrl-B will not open the UI menu although the Ctrl-B prompt at the Flex-Boot banner is present..
	Workaround: N/A
	Keywords: Ctrl-B, ConnectX-4 and Connect-IB HCAs, FlexBoot menu
841198	Description: FlexBoot fails to boot when the following occurs: <ul style="list-style-type: none"> • Boot priority is set to iSCSI • The iSCSI TCP/IP parameters via DHCP is disabled • iSCSI boot fails or iSCSI boot to target configuration is set to disable
	Workaround: N/A
	Keywords: PXE boot, iSCSI
843377/ 849223	Description: The physical MAC assigned via the boot menu is displayed as zeroes instead of the set MAC when ConnectX-4 VPI adapter card is configured as Infini-Band.
	Workaround: N/A
	Keywords: Physical MAC, Boot menu
776057	Description: Citrix PVS boot is not supported.
	Workaround: N/A
	Keywords: Citrix PVS boot
780862	Description: Set default (CTRL+d key) feature is not supported in Wake-On-Lan setting.
	Workaround: N/A
	Keywords: CTRL+d, WoL
843209	Description: When the first port is ETH in the PXE, the link does not raise in the second port which is set as IB.
	Workaround: N/A
	Keywords: Link up, Ports
754514	Description: Occasionally, working with lpxlinux.0 might result in PXE boot failure due to lpxlinux.0 issue.
	Workaround: N/A
	Keywords: Syslinux

6.2 FlexBoot Bug Fixes History

Table 21 - FlexBoot Bug Fixes History (Sheet 1 of 2)

Version	Issue
847950	Description: Fixed wrong default value of Boot-To-Target in FlexBoot configuration.
	Keywords: Boot-To-Target, FlexBoot configuration
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.903
691148	Description: When connecting a pre-configured port with VLAN to an IB fabric, the port runs as Ethernet port with the VLAN tag.
	Keywords: VLAN, Port Management
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.903
792432	Description: Booting PXE using Grub2.X over HP G9/G8 servers results in system hang.
	Keywords: PXE boot, Grub2.X, HP G9/G8
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.903
737512	Description: If the client gets "PXE boot menu" when contacting the DHCP, it will PXE boot first regardless of the boot priority.
	Keywords: ISCSI, DHCP
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.812
690792	Description: If the PMM fails to allocate memory, the system hangs since Flex-Boot cannot load from the expansion ROM.
	Keywords: PMM, expansion ROM
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.812
697291	Description: In ConnectX-4, the PXE boot time measurement over TFTP Ethernet is 1:30 min for image size of 1GB, TFTP InfiniBand is 1:20 min, and iSCSI boot time measurement is 8 seconds for image size of 25 MB.
	Keywords: PXE Boot
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.812

Table 21 - FlexBoot Bug Fixes History (Sheet 2 of 2)

Version	Issue
689068	Description: In hybrid BIOSes, if the BIOS loads legacy driver without closing the UEFI driver, the legacy driver fails to load.
	Keywords: BIOS, legacy mode
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.812
689068	Description: In hybrid BIOSes, if the BIOS loads legacy driver without closing the UEFI driver, the legacy driver fails to load.
	Keywords: BIOS, legacy mode
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.722
690792	Description: If the PMM fails to allocate memory, the system hangs since Flex-Boot cannot load from the expansion ROM.
	Keywords: PMM, expansion ROM
	Discovered in Release: 3.4.719
	Fixed in Release: 3.4.722
634794	Description: Enabled 'boot_pci_busdevfn' initialization when booting from UNDI loader.
	Keywords: UNDI loader
	Discovered in Release: 3.4.650
	Fixed in Release: 3.4.719
-	Description: Removed the instruction that enabled write-protected section modifications after POST.
	Keywords: PXE Boot
	Discovered in Release: 3.4.650
	Fixed in Release: 3.4.719

7 UEFI Changes and Major New Features

Table 22 - UEFI Changes and Major New Features

Version	Description
14.11.34	<ul style="list-style-type: none"> RM#919435 Removed Network Link Type HII Attribute from the UEFI menu Bug Fixes: See Section 7.1, “UEFI Bug Fixes History”, on page 56

7.1 UEFI Bug Fixes History

Table 23 - UEFI Bug Fixes History

Version	Issue
915043	Description: Fixed a case where the Number of Virtual Functions Supported was limited to 126 after restore to default.
	Keywords: Number of Virtual Functions Supported
	Discovered in Release: 14.11.28
	Fixed in Release: 14.11.34

8 Unsupported Features and Commands

8.1 Unsupported Features

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

- Service types not supported:
 - SyncUMR
 - Mellanox transport
 - PTP
 - RAW IPv6
 - PTP (IEEE 1588)
- INT-A not supported for EQs only MSI-X
- PCI VPD write flow (RO flow supported)
- Streaming Receive Queue (STRQ) and collapsed CQ
- Precise clock synchronization over the network (IEEE 1588)
- 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.

8.2 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

9 Supported Non-Volatile Configurations

Table 24 - 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 25 - 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 26 - Per host/function Settings

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