



Connect. Accelerate. Outperform.™

Mellanox ConnectX®-3 Firmware Release Notes

Rev 2.32.5100

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (“PRODUCT(S)”) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES “AS-IS” WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER’S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND 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

Mellanox Technologies, Ltd.
 Beit Mellanox
 PO Box 586 Yokneam 20692
 Israel
www.mellanox.com
 Tel: +972 (0)74 723 7200
 Fax: +972 (0)4 959 3245

© Copyright 2014. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, ConnectX®, Connect-IB®, CoolBox®, CORE-Direct®, InfiniBridge®, InfiniHost®, InfiniScale®, MetroX®, MLNX-OS®, TestX®, PhyX®, ScalableHPC®, SwitchX®, UFM®, Virtual Protocol Interconnect® and Voltaire® are registered trademarks of Mellanox Technologies, Ltd.

ExtendX™, FabricIT™, HPC-X™, Mellanox Open Ethernet™, Mellanox PeerDirect™, Mellanox Virtual Modular Switch™, MetroDX™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Chapter 1 Overview	6
1.1 Supported Devices	6
1.2 Supported Cables and Modules	7
1.2.1 Tested Cables and Modules	7
1.3 Tested Switches	13
1.4 Tools, Switch Firmware and Driver Software	15
1.5 Revision Compatibility	15
1.6 Firmware Burning Notes	16
Chapter 2 Supported Features	17
2.1 Firmware Rev 2.32.5100 Changes and New Feature	17
Chapter 3 Known Issues	18
Chapter 4 Bug Fixes History	21
Chapter 5 Firmware Changes and New Feature History	34
Appendix A Flexboot Changes and New Features	38

List of Tables

Table 1:	Release Update History	5
Table 2:	Supported PSIDs	6
Table 3:	Tested Cables and Modules	7
Table 4:	Tested Switches	13
Table 5:	Firmware Rev 2.32.5100 Changes and New Feature	17
Table 6:	Known Issues	18
Table 7:	Fixed Bugs List	21
Table 8:	Firmware Changes and New Feature History	34
Table 9:	FlexBoot Changes and New Feature	38

Release Update History

Table 1 - Release Update History

Release	Date	Description
2.32.5100	December 2014	Initial version

1 Overview

These are the release notes for the ConnectX®-3adapters firmware Rev 2.32.5100. This firmware supports the following protocols:

- InfiniBand – SDR, DDR, QDR, FDR10, FDR
- Ethernet - 1GigE, 10GigE, 40GigE and 56GigE¹
- 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](#). For the most updated list of adapter cards supported, visit the firmware download pages via <http://www.mellanox.com>.



Please contact your Mellanox local FAE for firmware updates to pre-production cards not on this list.

Table 2 - Supported PSIDs

Device Part Number	PSID	Device Name	Supported Protocols
0J05YT	DEL0A10210018	ConnectX®-3 Dual-port QDR Mezzanine I/O Card	InfiniBand
0P90JM	DEL0A40000028	ConnectX®-3 Dual-port FDR10 Mezzanine I/O Card	InfiniBand
0K6V3V	DEL0A20210018	ConnectX®-3 Dual-port FDR10 Mezzanine I/O Card	InfiniBand
08PTD1	DEL0A20000018	ConnectX®-3 Dual-port FDR Mezzanine I/O Card	InfiniBand
0CDMG5	DEL0A40001028	ConnectX®-3 Dual-port FDR10 Mezzanine I/O Card	InfiniBand
0T483W	DEL0A30000019	ConnectX®-3 VPI single port QSFP, FDR IB (56Gb/s) and 40GbE I/O card for Dell DCS	VPI
01T7NW	DEL1090001019	ConnectX®-3 Dual Port VPI FDR Low Profile Adapter	VPI
079DJ3	DEL1100001019	ConnectX®-3 Single Port VPI FDR Low Profile Adapter	VPI

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

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 Tested Cables and Modules

Table 3 - Tested Cables and Modules (Sheet 1 of 6)

Speed	OPN #	Description	Vendor
DDR	MC1204128-005	MT PASSIVE COPPER CABLE 4X CX4 TO QSFP 20GB/S 28AWG 5M	Mellanox
DDR	MC1204130-001	MT PASSIVE COPPER CABLE 4X CX4 TO QSFP 20GB/S 30AWG 1M	Mellanox
DDR	MC1204130-003	MT PASSIVE COPPER CABLE 4X CX4 TO QSFP 20GB/S 30AWG 3M	Mellanox
SDR	MC1104130-001	MT PASSIVE COPPER CABLE 4X CX4 20GB/S 30AWG 1M	Mellanox
SDR	MC1104130-003	MT PASSIVE COPPER CABLE 4X CX4 20GB/S 30AWG 3M	Mellanox
QDR	MC2206310-300	MT ACTIVE FIBER CABLE 4X QSFP 40GB/S 200M	Mellanox
QDR	MC2206125-007	MT PASSIVE COPPER CABLE 4X QSFP 40GB/S 25AWG 7M	Mellanox
QDR	MC2206126-006	MT PASSIVE COPPER CABLE 4X QSFP 40GB/S 26AWG 6M	Mellanox
QDR	MC2206230-010	MT ACTIVE COPPER CABLE IB QDR 40GB/S QSFP 10M	Mellanox
QDR/FDR10	MC2206310-030	MT ACTIVE FIBER CABLE 4X QSFP 40GB/S 30M	Mellanox
QDR / FDR10	MC2206310-100	MT ACTIVE FIBER CABLE 4X QSFP 40GB/S 100M	Mellanox
QDR / FDR10/ 40GE	MFS4R12CB-100	MT ACTIVE FIBER CABLE IB QDR/FDR10 40GB/S QSFP 100M	Mellanox
QDR / FDR10	MC2206128-005	MT PASSIVE COPPER CABLE 4X QSFP 40GB/S 28AWG 5M	Mellanox
QDR / FDR10	MC2206130-003	MT PASSIVE COPPER CABLE 4X QSFP 40GB/S 30AWG 3M	Mellanox
QDR / FDR10	MC2206130-00A	MT PASSIVE COPPER CABLE 4X QSFP 40GB/S 30AWG 0.5M	Mellanox

Table 3 - Tested Cables and Modules (Sheet 2 of 6)

Speed	OPN #	Description	Vendor
QDR / FDR10 on SwitchX only	MC2210511-LR4 - F	Mellanox Optical Module 40Gb/S QSFP LC-LC 1310NM LR4. Up To 10KM	Mellanox
QDR / FDR10/ 40GE	MC2210411-SR4 -T	Mellanox Optical Module 40Gb/S QSFP LC-LC 850NM LR4. Up To 100M	Mellanox
QDR / FDR10/ 40GE	MC2210411-SR4 -F	Mellanox Optical Module 40Gb/S QSFP LC-LC 850NM LR4. Up To 100M	Mellanox
40GE	MC2210126-004	MT PASSIVE COPPER CABLE 4X QSFP 40GE 26AWG 4M	Mellanox
40GE	MC2210126-005	MT PASSIVE COPPER CABLE 4X QSFP 40GE 26AWG 5M	Mellanox
40GE	MC2210128-003	MT PASSIVE COPPER CABLE 4X QSFP 40GE 28AWG 3M	Mellanox
40GE	MC2210130-001	MT PASSIVE COPPER CABLE 4X QSFP 40GE 30AWG 1M	Mellanox
40GE	MC2210310-30-E	Mellanox® active fiber cable, ETH 40GbE, 40Gb/s, QSFP	Mellanox
40GE	MC2210310-100-E	Mellanox® active fiber cable, ETH 40GbE, 40Gb/s, QSFP	Mellanox
40GE- >10GE	MC2609125-005	MT PASSIVE COPPER CABLE QSFP TO 4 SFP+ 10GB/S 26AWG 5M	Mellanox
40GE- >10GE	MC2609130-001	MT PASSIVE COPPER CABLE QSFP TO 4 SFP+ 10GB/S 30AWG 1M	Mellanox
40GE- >10GE	MC2609130-003	MT PASSIVE COPPER CABLE QSFP TO 4 SFP+ 10GB/S 30AWG 3M	Mellanox
40GE- >10GE	MC6709309-050	MT PASSIVE OPTICAL CABLE MULTIMODE SPLITTER MPO TO 8XLC 50 METER	Mellanox
10GE	MC2309124-006	MT PASSIVE COPPER CABLE 1X SFP+ TO QSFP 10GB/S 24AWG 7M	Mellanox
10GE	MC2309124-007	MT PASSIVE COPPER CABLE 1X SFP+ TO QSFP 10GB/S 24AWG 7M	Mellanox
10GE	MC2309130-00A	MT PASSIVE COPPER CABLE 1X SFP+ TO QSFP 10GB/S 30AWG 0.5M	Mellanox
10GE	MC2309130-003	MT PASSIVE COPPER CABLE 1X SFP+ TO QSFP 10GB/S 30AWG 3M	Mellanox

Table 3 - Tested Cables and Modules (Sheet 3 of 6)

Speed	OPN #	Description	Vendor
10GE	MC3309130-00A	MT PASSIVE COPPER CABLE 1X SFP+ 10GB/S 30AWG 0.5M	Mellanox
10GE	MC3309124-005	MT PASSIVE COPPER CABLE 1X SFP+ 10GB/S 24AWG 5M	Mellanox
10GE	MC3309130-003	MT PASSIVE COPPER CABLE 1X SFP+ 10GB/S 30AWG 3M	Mellanox
10GE	MC3309124-007	MT PASSIVE COPPER CABLE 1X SFP+ 10GB/S 24AWG 7M	Mellanox
10GE	MFM1T02A-LR	MELLANOX OPTICAL MODULE ETH 10GBE 10GB/S SFP+ LC-LC 1310NM LR UP TO 10KM	Mellanox
10GE	MFM1T02A-SR	MT OPTICAL MODULE ETH 10GBE 10GB/S SFP+ LC-LC 850NM SR UP TO 300M	Mellanox
1GE	MC3208011-SX	MELLANOX OPTICAL MODULE ETH 1GBE 1GB/S SFP LC-LC SX 850NM UP TO 500M	Mellanox
1GE	MC3208411-T	MELLANOX MODULE ETH 1GBE 1GB/S SFP BASE-T UP TO 100M	Mellanox
10GE	MAM1Q00A-QSA	MELLANOX QSFP TO SFP+ ADAPTER	Mellanox
FDR; 56G VPI	MC2207130-00A	MT PASSIVE COPPER CABLE 4X QSFP 56GB/S 30AWG 0.5M	Voletex - (Taipan)
FDR; 56G VPI	MC2207130-002	MT PASSIVE COPPER CABLE 4X QSFP 56GB/S 30AWG 2M	Voletex - (Taipan)
FDR; 56G VPI	MC2207128-003	MT PASSIVE COPPER CABLE 4X QSFP 56GB/S 28AWG 3M	Voletex - (Taipan)
FDR; 56G VPI	MC2207126-004*	MT PASSIVE COPPER CABLE 4X QSFP 56GB/S 28AWG 4M	Voletex - (Taipan)
FDR; 56 VPI	MC2207312-003	MT ACTIVE FIBER CABLE VPI IB FDR (56GB/S) AND ETH 40GBE QSFP 3M	Voletex - (Taipan)
FDR; 56G VPI	MC2207310-030-E	MT ACTIVE FIBER CABLE 4X QSFP 56GB/S 30M	Mellanox (Emcore)
FDR; 56G VPI	MC2207312-100	MT ACTIVE FIBER CABLE 4X QSFP 56GB/S 100M	Mellanox (Molex)
FDR; 56G VPI	MC2207310-100-E	MT ACTIVE FIBER CABLE 4X QSFP 56GB/S 100M	Mellanox (Emcore)
FDR; 56G VPI	MC2207310-030-F	MT ACTIVE FIBER CABLE 4X QSFP 56GB/S 30M	Mellanox (Finisar)
FDR; 56G VPI	MC2207310-100-F	MT ACTIVE FIBER CABLE 4X QSFP 56GB/S 100M	Mellanox (Finisar)

Table 3 - Tested Cables and Modules (Sheet 4 of 6)

Speed	OPN #	Description	Vendor
10GE	Cisco SFP-H10GB-CU1M	Cisco SFP+ cable	Cisco
10GE	Cisco SFP-H10GB-CU3M	Cisco SFP+ cable	Cisco
10GE	Cisco SFP-H10GB-CU5M	Cisco SFP+ cable	Cisco
40GE	QSFP-H40G-CU1M	Cisco QSFP 40GbE cable	Cisco
40GE	QSFP-H40G-CU3M	Cisco QSFP 40GbE cable	Cisco
40GE	QSFP-H40G-CU5M	Cisco QSFP 40GbE cable	Cisco
40GE->10GE	MC2309124-007	QSFP-4SFP10G	Cisco
40GE	QSFP-40G-SR4	CISCO 40G QSFP Module	Cisco
10GE	SFP-10G-SR	CISCO 10GBASE-SR SFP Module	Cisco
40GE->10GE	MC2309124-007	QSFP-4SFP10G	Cisco
10GE	CAB-SFP-SFP-1M	passive copper cable, SFP+, 10 Gb/s, 1m	Arista
10GE	CAB-SFP-SFP-2M	passive copper cable, SFP+, 10 Gb/s, 2m	Arista
10GE	CAB-SFP-SFP-3M	passive copper cable, SFP+, 10 Gb/s, 3m	Arista
10GE	CAB-SFP-SFP-5M	passive copper cable, SFP+, 10 Gb/s, 5m	Arista
40GE	40GbE QSFP+ to QSFP	QSFP+ copper cable 3M	Arista
40GE	40GbE QSFP+ to QSFP	QSFP+ copper cable 5M	Arista
10GE	QFX-SFP-DAC-1M	SFP+ 10 Gigabit Ethernet Direct Attach Copper (twinax copper cable) 1 m	Juniper
10GE	QFX-SFP-DAC-3M	SFP+ 10 Gigabit Ethernet Direct Attach Copper (twinax copper cable) 3 m	Juniper
10GE	QFX-SFP-DAC-5M	SFP+ 10 Gigabit Ethernet Direct Attach Copper (twinax copper cable) 5 m	Juniper
40GE	90Y9425-N28500A	IBM-Amphenol SFP+ 1m	IBM
40GE	46K6184-L36836B	IBM-Amphenol SFP+ 5m	IBM
40GE	46K6183-L36836B	IBM-Amphenol SFP+ 3m	IBM
40GE	44X1371-N31295E	IBM-Amphenol SFP+ 7m	IBM
40GE	00D5802-N13445C	IBM 40GbE QSFP+ to QSFP+ copper cable 1M	IBM

Table 3 - Tested Cables and Modules (Sheet 5 of 6)

Speed	OPN #	Description	Vendor
40GE	BN-QS-QS-CBL-3M	IBM 40GbE QSFP+ to QSFP+ copper cable 3M	IBM
40GE	BN-QS-QS-CBL-5M	IBM 40GbE QSFP+ to QSFP+ copper cable 5M	IBM
FDR	038-004-066-01	EMC FDR QSFP+ to QSFP+ copper cable 2M	EMCFDR QSFP
FDR	038-004-067-01	EMC FDR QSFP+ to QSFP+ copper cable 3M	EMCFDR QSFP
FDR	038-900-027-01	EMC FDR QSFP+ to QSFP+ copper cable 5M	EMCFDR QSFP
FDR	038-900-030-01	EMC FDR QSFP+ to QSFP+ copper cable 8M	EMCFDR QSFP
40GE	JG325A	HP X140 40G QSFP+ MPO SR4 Transceiver	HP
40GE	JG326A	HP X240 40G QSFP+ QSFP+ 1m DAC Cable	HP
40GE	JG327A	HP X240 40G QSFP+ QSFP+ 3m DAC Cable	HP
40GE	JG328A	HP X240 40G QSFP+ QSFP+ 5m DAC Cable	HP
10GE	J9281B	HP X242 10G SFP+ to SFP+ 1m Direct Attach Copper Cable	HP
10GE	J9283B	HP X242 10G SFP+ SFP+ 3m DAC Cable	HP
10GE	J9285B	HP X242 10G SFP+ SFP+ 7m DAC Cable	HP
10GE	JD096B	HP X240 10G SFP+ SFP+ 1.2m DAC Cable	HP
10GE	JD095B	HP X240 10G SFP+ SFP+ 0.65m DAC Cable	HP
10GE	JD097B	HP X240 10G SFP+ SFP+ 3m DAD Cable	HP
10GE	487649-B21	HP BLc SFP+ .5m 10GbE Copper Cable	HP
10GE	487652-B21	HP BLc SFP+ 1m 10GbE Copper Cable	HP
10GE	487655-B21	HP BLc SFP+ 3m 10GbE Copper Cable	HP
10GE	537963-B21	HP BLc SFP+ 5m 10GbE Copper Cable	HP
10GE	487658-B21	HP BLc SFP+ 7m 10GbE Copper Cable	HP
10GE	AP784A	HP 3m C-series Passive Copper SFP+ Cable	HP
10GE	AP785A	HP 5m C-series Passive Copper SFP+ Cable	HP
10GE	AP818A	HP 1m B-series Active Copper SFP+ Cable	HP
10GE	AP819A	HP 3m B-series Active Copper SFP+ Cable	HP
1GB	453151-B21	HP BLc VC 1Gb SX SFP Opt Kit	HP
1GB	453154-B21	HP BLc VC 1Gb RJ-45 SFP Opt Kit	HP

Table 3 - Tested Cables and Modules (Sheet 6 of 6)

Speed	OPN #	Description	Vendor
10GE	455883-B21	HP BLc 10Gb SR SFP+ Opt	HP
10GE	455886-B21	HP BLc 10Gb LR SFP+ Opt	HP
10GE	J9150A	HP X132 10G SFP+ LC SR Transceiver	HP
10GE	J9151A	HP X132 10G SFP+ LC LR Transceiver	HP

1.3 Tested Switches

Table 4 - Tested Switches (Sheet 1 of 3)

Speed	OPN # /Name	Description
IB SDR	F-X430060	24-port SDR-Switch
IB DDR	F-X430044	24-port DDR-Switch
IB DDR	9024	24-port DDR-Switch
IB QDR	12300	36-Port 40Gb QDR Infiniband Switch, Management Module, Dual Power
IB QDR	MIS5025Q-1SFC	InfiniScale® IV QDR InfiniBand Switch, 36 QSFP ports, 1 Power Supply, Unmanaged, PSU side to connector side airflow, Standard depth, Rail Kit and RoHS5
IB QDR	MIS5024Q-1BFR	InfiniScale® IV QDR InfiniBand Switch, 36 QSFP ports, 1 power supply, Unmanaged, PSU side to Connector side airflow, no FRUs, with rack rails, Short Depth Form Factor
IB QDR	4036	InfiniScale® IV QDR Mellanox Grid Director 4036 36-Port QDR InfiniBand Switch - Part ID: VLT-30011
IB FDR	MSX6036F-1BFR	SwitchX™ based FDR InfiniBand Switch, 36 QSFP ports, 1 Power Supply, Short depth, Managed, PSU side to Connector side airflow, Rail Kit and RoHS6
IB FDR10	MSX6025T-1SFR	SwitchX™ based FDR10 Infiniband Switch, 36 QSFP ports, 1 Power Supply, Standard depth, Unmanaged, PSU side to Connector side airflow, Rail Kit and RoHS6
40Gb/s	MSX1036B-1BFR	SwitchX™ based 36-port QSFP 40GigE 1U Ethernet
10Gb/s	MSX1016X-1BFR	SwitchX™ based 64-port SFP+ 10GigE, 1U Ethernet switch
10/40Gb/s	MSX1024B-1BFS	SwitchX®-2 based 48-port SFP+ 10GbE, 12 port QSFP 40GbE, 1U Ethernet switch
10/40Gb/s	S4810P-AC	48-port 10Gb/40Gb Switch
10/40Gb/s	3064	48-port 10Gb/40Gb Switch
40GbE	3016	Cisco 40GB ETH switch
10Gb/s	5548	Cisco 10GB ETH switch
40GbE	3132	Cisco 40GB ETH switch
10/40Gb/s	7050Q	16-port 40Gb Switch

Table 4 - Tested Switches (Sheet 2 of 3)

Speed	OPN # /Name	Description
10/40Gb/s	7050S	48-port 10Gb/40Gb Switch
40GbE	7050QX	32-port 40Gb Switch
10Gb/s	Fujitsu 10GbE (XAUI)	Ethernet Switch 24 ports, 20xCX4 and 4xQSFP
10Gb/s	6600-24XG	24-port 10GbE switch
10/40Gb/s	IBM G8264	IBM 10/40GB ETH switch
40GbE	IBM G8316	IBM 40GB RackSwitch G8316
40GbE	S6000	32-port 40Gb Switch
10/40Gb/s	S5000	10GbE switch
10Gb/s	8024F	10GbE switch
1Gb/s	7024F	1/10GbE switch
10Gb/s	8132F	10GbE switch
10Gb/s	8164F	10GbE switch
10Gb/s	EX2500	Juniper 10GB ETH switch
10Gb/s	EX4550	Juniper 10GB ETH switch
10/40Gb/s	EX3500	Juniper 10/40GB ETH switch
1/10Gb/s	X650	Extreme 10GB ETH switch
10Gb/s	8000	Brocade 10GB ETH switch
10/40Gb/s	ASF5900	HP 10GB ETH switch
1Gb/s	2810-24G.	HP 1GB ETH switch
IB FDR	M4001F	SwitchX® 56Gb/s 16+16 port InfiniBand switch blade for the Dell M1000e Blade System
IB FDR10	M4001T	SwitchX® 40Gb/s 16+16 port InfiniBand switch blade for the Dell M1000e Blade System
IB QDR	M3601Q	40Gb/s 16+16 port InfiniBand switch blade for the Dell M1000e Blade System
IB DDR	M2401G	Infiniscale III 24-Port 20Gb/s InfiniBand Switch for Dell M1000E Blade System
40GbE	689638-B21	Mellanox SX1018HP Enet Switch 40G Ethernet
10GbE	516733-B21	HP ProCurve 6120XG 10GbE Ethernet Blade Switch

Table 4 - Tested Switches (Sheet 3 of 3)

Speed	OPN # /Name	Description
10GbE	6125XLG Blade Switch	HP 6125XLG 10/40G Ethernet Blade Switch
10GbE	538113-B21	HP 10GbE Pass-Through Module (PTM)
1GbE	3020X	Cisco Catalyst 3020X 1GbE switch blade
1GbE	3020	Cisco Catalyst 3020 1GbE switch blade
1GbE	438030-B21	HP 1GbE switch module - GbE2c Layer 2/3 Ethernet Blade Switch
1GbE	6120G	HP ProCurve 6120G/XG 1GbE switch blade
IB FDR	648311-B21	HP BLc 4X FDR IB Switch
IB QDR	489184-B21	HP BLc 4X QDR IB Switch
IB DDR	410398-B21	HP BLc 4X DDR IB Switch
40GbE	90Y3477	BM Flex System EN6131 40Gb Ethernet Switch
IB FDR	90Y3452	IBM Flex System IB6131 Infiniband Switch
10GbE	B22	IBM B22 10 Gigabit Scalable Switch Module
10GbE	EN4093	IBM PureFlex System Fabric EN4093 10 Gigabit Scalable Switch Module

1.4 Tools, Switch Firmware and Driver Software

Firmware Rev 2.32.5100 is tested with the following tools, SwitchX® firmware, and driver software:

- Driver versions:
 - MLNX-OFED 2.3-1.0.0 and higher
 - MLNX_EN 2.2-1.0.1 and higher
 - WinOF 4.80 and higher
- MFT: 3.7.0 and higher
- PXE version 3.4.306 and higher
- UEFI version 11.5.24

1.5 Revision Compatibility

Firmware Rev 2.32.5100 complies with the following programmer's reference manual:

- *ConnectX Programmer's Reference Manual (PRM), Rev 1.7 or later*, which has Command Interface Revision 0x3. The command interface revision can be retrieved by means of the QUERY_FW command and is indicated by the field `cmd_interface_rev`.

1.6 Firmware Burning Notes

- Firmware Family Version (FFV)

As of firmware v2.30.8000, all firmware images have the FFV field populated. The FFV value is identical to the firmware version but in a different format.

FFV format example:

```
FW version:      2.30.8000
FFV:             02.30.80.00
```

- Updating EXP_ROM

Updating only the EXP_ROM (FlexBoot/UEFI) for firmware images which contain FFV requires an additional MFT flag: "-allow_rom_change"

The following is an example for removing the EXP_ROM from the binary image using Flint (a Mellanox device located at PCI bus function 05:00.0):

```
$ flint -d 05:00.0 -allow_rom_change drom
```


2 Supported Features

2.1 Firmware Rev 2.32.5100 Changes and New Feature

Table 5 - Firmware Rev 2.32.5100 Changes and New Feature

Category	Description
Ethtool Improvements	Added support to query PTYS, PTOS registers through ACCESS_REG PRM command.
Non-volatile Random Access Memory (NVRAM)	Added support for CLP access to NVRAM.
Packet Steering	Added support for more than 22 QPs per MCG in DMFS.
	Added support for high rate steering mode (a.k.a Simplified Steering)
PRM	Added support for reading current hardware mode through the QUERY_PORT PRM command.
	Added CSUM mode reporting in QUERY_DEV_CAP command.
	Added additional configuration options for UPDATE_QP command.
	Added support for 128 Byte stride for CQ/EQ.
	Enabled module EEPROM access using command I/F.
Device Reset and Error recovery	Reset Flow improvements and graceful handling of error caused by Virtual Functions.
Performance	RX performance optimization for single port cards.
	Promiscuous mode performance improvements.
Secure Host	Added support for Secure Host mode for ConnectX-3/ConnectX-3 Pro adapter cards.
Non-Volatile device configuration	Added Port protocol configuration option.
GPIO	Added support for GPIO swap.
Signal Integrity (SI)	40GbE SI improvements.
MAD	Added support for Temp Sensing Vendor specific MAD.
NC-SI	Added Temp Sensing NC-SI cmd.
	Added support for AEN.
SMBUS	Added new command to report firmware revision.

3 Known Issues

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

Table 6 - Known Issues

Index	Issue	Description	Workaround	Scheduled Release (fix)
1.	Downgrade to previous GA requires server reboot.	Downgrading from v2.30.8000 or later to an earlier version than 2.30.8000 requires server reboot.	Reboot the server.	Future Release
2.	GUID ConnectX®-3 Ethernet adapter cards	On ConnectX-3 Ethernet adapter cards, there is a mismatch between the GUID value returned by firmware management tools and that returned by fabric/driver utilities that read the GUID via device firmware (e.g., using <code>ibstat</code>). Mlxburn/flint return 0xffff as GUID while the utilities return a value derived from the MAC address. For all driver/firmware/software purposes, the latter value should be used.	N/A. Please use the GUID value returned by the fabric/driver utilities (not 0xffff).	N/A
3.	SBR assertion	SBR should be asserted for a minimum of 50 milliseconds for the ConnectX-3 adapters	N/A	N/A
4.	PCIe	On Pilot1 SL230, PCIe link occasionally does not come up at Gen3 speed	Production SL230 should be used for PCIe Gen3 operation	N/A
5.	Kernel panic in SR-IOV with RH6.3 Inbox driver and VPI cards	RH6.3 Inbox driver causes kernel panic when SR-IOV is enabled on VPI cards due to driver compatibility issue.	Set the <code>"do_sense=false"</code> parameter in the [IB_TAB] in the INI of the VPI card	N/A
6.	Side band Management compatibility with SR-IOV	In advanced steering mode, side band management connectivity may be lost when having more than 8 QP per mcg.	N/A	Future Release
7.	SR-IOV disabled in the BIOS	When SR-IOV is disabled in the system BIOS, a PCI issue is noticed in Ubuntu v12.04.3 with Linux kernel v3.8 which affects NICs of several manufacturers including Mellanox's, preventing them from operating.	Enable SR-IOV in the BIOS	Future Release

Table 6 - Known Issues

Index	Issue	Description	Workaround	Scheduled Release (fix)
8.	MFT locking of flash semaphore	MFT tools might leave the flash semaphore locked if the tool operation is forced stopped. The locked semaphore prevents the firmware from accessing the flash and causes firmware hang.	Clear the semaphore using MFT command: 'flint -clear_semaphore'	Future Release
9.	MC2210411-SR4 module with Cable Info MAD	Cable Info MAD reports a wrong cable info when using the MC2210411-SR4 module	N/A	Future Release
10.	PCIe failure on temperature shock 10C/min	Gen2 failure at temperature sweep up to 10C/min (for MT27518A1-FDIR-BV only).	N/A	Future Release
11.	PCIe Gen2 link	PCIe Gen2 link unstable at temperature sweep of 10C/min for MT27518A1-FDIR-BV	N/A	Future Release
12.	Changing from an LLR to non-LLR requires driver restart	Driver restart required when switching from InfiniBand FDR link with LLR enabled to InfiniBand link w/o LLR (for example: between SwitchX® and GD4036).	N/A	Future Release
13.	Bloom filter	Bloom filter is currently not supported.	N/A	Future Release
14.	Firmware downgrade	When downgrading from firmware v2.11.0000 and using MFT 3.0.0-3, the following message is displayed due to the mlxconfig tool: You are trying to override configurable FW by non-configurable FW. If you continue, old FW configurations will be cleared, do you want to continue ? (y/n) [n] : y You are trying to restore default configuration, do you want to continue ? (y/n) [n] : y	N/A	Future Release
15.	DMFS steering mode with IB in Linux	DMFS should not be enabled when working with InfiniBand on MLNX_OFED-2.0.3	Upgrade to MLNX_OFED-2.1-x.x.x or later	Future Release
16.	VPD read-only fields	VPD read-only fields are writable.	Do not write to read-only fields if you wish to preserve them	Future Release

Table 6 - Known Issues

Index	Issue	Description	Workaround	Scheduled Release (fix)
17.	Increasing SymbolErrorCounter	When working in VPI mode with port1 FDR and port2 40G, error counters misbehave and increase rapidly	N/A	Future Release
18.	128 Byte CQ/EQ stride compatibility with side-band Management	Setting the device to 128Byte CQ/EQ stride will cause misbehavior of side-band management resulting in communication loss.	N/A	Future Release
19.	128 Byte CQ/EQ stride	CQ and EQ cannot be configured to different stride sizes.	N/A	Future Release

4 Bug Fixes History

Table 7 lists the bugs fixed in this release.

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
1.	IBDump performance	Fixed performance degradation when running IBDump	2.30.8000	2.32.5100
2.	PCIe link Disable/Enable	Occasionally, a link training timeout occurred in EQ phase0 during disable/enable test.	2.31.5050	2.32.5100
3.	40GbE QoS	Improved strict bandwidth mode functionality	2.30.8000	2.32.5100
4.	Port Counters reporting	Fixed an issue with the PortRcvPkts counter always displaying zero value.	2.31.5050	2.32.5100
5.	GMP MADs in SecureHost	Fixed an issue with processing GMP MADs with SET method in SecureHost mode.	2.31.5050	2.32.5100
6.	NC-SI over IPv6	Fixed an issue causing a wrong usage of MCG size when configuring Global Multicast filter	2.31.5050	2.32.5100
7.	NC SI link failure	Disabling the first port occasionally causes second port TX failure.	2.31.5050	2.32.5100
8.	10GbE link failure	Fixed a mismatch in links status reported. The adapter reports links as down while the switch perceives them as up	2.31.5050	2.32.5100
9.	Link failure	Fixed an occasional 40GbE link failure with SCM5 Switch blade	2.31.5050	2.32.5100
10.	ExtPortInfo MAD	Fixed a wrong FDR10 speed reporting in MAD	2.31.5050	2.32.5100
11.	IB link failure	Fixed an issue preventing the ports to rise up when set to FDR10 vs QDR	2.31.5050	2.32.5100
12.	40GbE link failure	Fixed an occasional link failure vs Arista switch	2.31.5050	2.32.5100
13.	RDMA Write retransmission	Retransmission started from the first PSN of message instead of the last acknowledged PSN	2.30.3200	2.32.5100
14.	Firmware burning	Firmware hangs when receiving GeneralInfoMad during inline firmware burning	2.30.3200	2.32.5100
15.	PCIe PML1	L1 flow adjustments and threshold tuning	2.31.5050	2.32.5100

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
16.	PCIe reset	Fixed a rare hanging issue during PERST_assertion	2.31.5050	2.32.5100
17.	PCIe Gen3 EQ	Wrong coefficients were reported during phase3	2.31.5050	2.32.5100
18.	Boot	Fixed an issue causing wrong behavior due to reset timing	2.31.5050	2.32.5100
19.	SMBUS	Fixed long timeout issues	2.31.5050	2.32.5100
20.	NVRAM	Fixed NVRAM write issues in driver-less mode	2.31.5050	2.32.5100
21.	40GbE Link support	Fixed 40GbE link support in aux mode	2.31.5050	2.32.5100
22.	NC-SI	Dropped commands with non-existing channel ID	2.31.5050	2.32.5100
23.	PRM PortInfo command	Fixed issues in extended speed reporting	2.31.5050	2.32.5100
24.	Trap 257/8(IB)	Fixed bad QP reporting in trap 257/8	2.30.8000	2.32.5100
25.	Bad Q_KEY errors	Fixed an issue causing false bad q_key error messages	2.30.8000	2.32.5100
26.	PFC	Fixed Pause Frame opcode mismatch	2.30.8000	2.32.5100
27.	Sideband Communication	Fixed communication loss upon PCIe error detection	2.31.5050	2.32.5100
28.	NC-SI	Fixed wrong channel value in the SELECT/DESELECT PACKAGE commands	2.30.8000	2.31.5050
29.		Fixed an issue caused response packet to include 4 extra bytes	2.30.8000	2.31.5050
30.		Fixed wrong reason code value returned when using Set Link command with unsupported speed.	2.30.8000	2.31.5050
31.		Added protection from bad MAC address given by BMC	2.30.8000	2.31.5050
32.		False Link Indication	Fixed an issue causing the device to report false link up when no cable is connected.	2.30.8000

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
33.	PCIe	Removed false TX pulse after PERST_ deassertion	2.30.8000	2.31.5050
34.		Fixed FLR capability bit inconsistency when SR-IOV is enabled.	2.30.3200	2.31.5050
35.		Fixed an issue with the device not reporting PCIe related errors.	2.30.8000	2.31.5050
36.	SDR instead of DDR ConnectX-3 to SX6036	When a link is configured to DDR in a setup of ConnectX-3 to SX6036, SDR link is established instead.	2.30.8000	2.31.5050
37.	VXLAN	VXLAN used the wrong default UDP port. the UDP port number was changed to 4789.	2.30.8000	2.31.5050
38.		Fixed wrong setting of the UDP destination port for VXLAN.	2.30.8000	2.31.5050
39.	Flow Steering	Fixed an internal error caused when moving to the DMS mode with IPMI/NC-SI enabled.	2.30.8000	2.31.5050
40.	FDR speed degradation with 0.5m cables	In a back-to-back setup of FDR cards connected with a 0.5m FDR cable, a link may be established as FDR10 instead of FDR.	2.30.3200	2.31.5050
41.	PCI interrupt	Fixed issues related to working with PCI legacy interrupts.	2.30.8000	2.31.5050
42.	TCP/UDP Checksum	Fixed wrong checksum calculation for short packets which were padded by the software.	2.30.8000	2.31.5050
43.	MFT tool deadlock	Reading PCIe configuration space after using the MFT flint tool caused the device to crash.	2.10.0000	2.31.5050
44.	Side band packet loss	Fixed occasional packet loss over IPMI	2.30.8000	2.31.5050
45.	Eye opening MAD	Fixed wrong values reported in the Eye opening MAD.	2.30.8000	2.31.5050
46.	PCIe Link width	Fixed occasional link width degrades during link negotiation and link transitions from L1 state.	2.30.8000	2.31.5050
47.	PCIe signal detect	Fixed adjust signal detect thresholds	2.30.8000	2.31.5050
48.	Error counters	PortExtendedSpeedsCounters MAD counters were mistakenly increased while LLR was active	2.30.8000	2.31.5050
49.	PCIe Gen3 Equalization	Lane reversal was not considered when configured TX parameters	2.30.8000	2.31.5050

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
50.	Reset On LAN (ROL)	Fixed ROL factory MAC usage when a FlexBoot address was given.	2.30.8000	2.31.5050
51.	Flow Control	Fixed Pause frames factory MAC usage when a FlexBoot address was given.	2.30.3200	2.31.5050
52.	WOL/ROL	The device did not different between WOL/ROL packets.	2.30.8000	2.31.5050
53.	PortInfo MAD	Fixed a set of extended fields in PortInfo MAD which did not function.	2.10.0000	2.31.5050
54.	LLR cell size	Adjusted LLR cell size according to the MLPN negotiation of ib_128b_llr	2.30.8000	2.31.5050
55.	Link max speed	The max speed restriction was active in full power mode instead of standby mode only.	2.30.8000	2.31.5050
56.	InfiniBand Automatic Path Migration	The InfiniBand Path migration did not work with GRH. http://webdev01:8080/commit/ConnectX.git/a9c37ee4c31038f2c1179d4d9e79c9337e0ab5c7	2.10.0000	2.31.5050
57.	Packet steering	Reading MGM after writing it returned wrong members count.	2.30.8000	2.31.5050
58.	RSS QP context	Fixed corruption of the RSS hash key given by the driver.	2.30.8000	2.31.5050
59.	10Gb/s QoS	Fixed QoS rate limit bandwidth offset.	2.30.3200	2.31.5050
60.	ExtendedPortInfo MAD	Fixed FDR10 speed_en reporting.	2.30.8000	2.31.5050
61.	Management link	Fixed long management link com loss.	2.30.8000	2.31.5050
62.	PRM Query_Port Command	The command results reported both link types active at the same time.	2.30.3200	2.31.5050
63.	Link not raising	Fixed collision between forcing phy type and port sensing.	2.30.8000	2.31.5050
64.	Core clock reporting	Fixed a wrong core clock freq reporting in QUERY_HCA command.	2.30.3200	2.31.5050
65.	56GbE link issues	Fixeds occasional link failure when 56GbE is enabled	2.30.8000	2.31.5050
66.	RX calibration	Fixed max eye margins to be per protocol.	2.30.8000	2.31.5050
67.	VPI symbol errors	perfquery reported wrong error symbol on ConnectX©-3 VPI mode: IB, ETH.	2.30.8000	2.31.5050

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
68.	Symbol error on ConnectX-3 dual-port QDR with MC2207312-030 AOCs	On ConnectX-3 dual-port QDR and FDR/FDR10 switch setups, symbol errors may occur with MC2207312-030 AOCs.	2.30.8000	2.31.5050
69.	Symbol error on Falcon QDR against FDR switches with MC2207126-004 copper cables	Symbol errors occur on ConnectX-3 dual-port QDR connected to FDR switches with MC2207126-004 copper cables.	2.30.8000	2.31.5050
70.	PCIe correctable errors in speed change	When PCIe Gen3 is enabled, temporary correctable errors might occur when changing speed between PCIe Gen1 and PCIe Gen2.	2.10.0000	2.30.8000
71.	Incompatibility between Device managed Flow steering and NC-SI	Device managed Flow steering and NC-SI cannot be enabled simultaneously.	2.11.0500	2.30.8000
72.	40GbE is not supported in Auto-Sensing	Auto-Sensing is not supported with 40GbE connections in VPI cards	2.10.0000	2.30.8000
73.	InfiniBand port_rcv_pkts counter	InfiniBand port_rcv_pkts counter over counts when LLR is enabled on the port. The port_rcv_pkts counter continues to count packets even when no traffic is received.	2.10.0000	2.30.8000
74.	PXE	PXE is currently not supported in 40GbE in VPI cards	2.10.0000	2.30.8000
75.		PXE is currently not supported in QSFP to SFP+ hybrid cable	2.10.0000	2.30.8000
76.	sense_port failure	Ethernet cards failed to work with MLNX_OFED unless the do_sense was disabled in the INI	RH6.4 driver	MLNX_OFED 2.0-3.0.0
77.	Link errors	BER of 10^{-11} with 7M copper SFP+ 10GbE cable against Arista switch	2.30.3200	2.30.8000
78.	Linkup Failure vs SwitchX®-2 based switch	Port failed to link up in 10GbE if it was previously linked up in 40GbE vs. SwitchX®-2 based switch	2.30.3200	2.30.8000
79.	NC-SI 40GbE reporting	Added 40GbE reporting in get_link_status NC-SI command	2.30.3000	2.30.8000

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
80.	Packets drop in receive when DMFS enabled	Steering entries overlapping caused packets to drop in the receive due to wrong hash size calculation of QP hash folding.	2.30.3000	2.30.8000
81.	PCIe speed degrade	Occasionally, PCIe speed degraded during speed change test	2.30.3000	2.30.8000
82.	Link failure vs Cisco	Device failed to raise the link against Cisco b-22 Blade switch	2.30.3000	2.30.8000
83.	False RX drops indication	Fixed an issue causing RX drop counters to falsely increase when using MLNX_OFED 2.0.-3.0.0	2.30.3000	2.30.8000
84.	NC-SI wrong command info	Wrong port information provided in get link status command.	2.30.3000	2.30.8000
85.	Port interfaces stay down	After firmware upgrade, the device failed to raise port interface.	2.30.3000	2.30.8000
86.	SR-IOV guest failure	Guest VM failed to execute firmware commands operations and crashed.	2.30.3000	2.30.8000
87.	ConnectX®-3 Pro: MCG write timeout	PRM WRITE_MCG command caused the device to hang.	2.30.3000	2.30.8000
88.	ConnectX®-3 Pro: Driver start failure	Driver could not start when NIC was configured for NC-SI SNP.	2.30.3000	2.30.8000
89.	SR-IOV command timeouts	Guest MSIX vectors were not assigned properly.	2.30.3000	2.30.8000
90.	PCIe speed degrade/link down	Occasionally, the PCIe link experienced speed degrading or link falling in driver restart/reboot	2.30.3000	2.30.8000
91.	QP Sniffer issue	Fixed a hash fold issue for sniffer QPs	2.30.3000	2.30.8000
92.	Long link up time	A long link up time is experienced in the HCA connected to a 10GbE cable against the MSX1012 switch	2.30.3000	2.30.8000
93.	PCIe TS parity bit	In recovery.EQLZ state TSs used incorrect parity bit calculation	2.30.3000	2.30.8000
94.	PRM Init_port failure	Init port command may fail on a system with NC-SI	2.30.3000	2.30.8000
95.	PortInfo MAD link width support	Wrong link_width_support is reported occasionally in PORT_INFO MAD	2.30.3000	2.30.8000
96.	Wrong Virtual Function completer ID	In SR-IOV a wrong completer ID is added in some of the VF completion packets	2.30.3000	2.30.8000

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
97.	PCIe PML1 failures	Fixed occasional failures upon entering and exiting L1 state in PCIe Gen1 & 2 speeds	2.30.3000	2.30.8000
98.	IB: APM failures in ConnectX-3 Pro	Occasional FSM transition timeouts are seen on APM requests	2.30.3000	2.30.8000
99.	ipmitool OOB commands	On rare occasions, ipmitool OOB commands failed upon send payload.	2.30.3000	2.30.8000
100.	Sideband communication	On rare occasions, after stress of power cycles, side- band communication might disconnect.	2.30.3000	2.30.8000
101.	Expansion ROM partition	Expansion ROM partition not found	2.30.3000	2.30.8000
102.	AC power cycle issue	In certain servers, AC power cycle may cause BMC connectivity loss.	2.30.3000	2.30.8000
103.	PXE issue	Occasionally, during DC cycle stress, failure occurred in PXE due to race condition	2.30.3000	2.30.8000
104.	SMBUS communication	SMBUS communication lost during AC/DC cycle	2.30.3000	2.30.8000
105.	MTU configuration	Temporal wrong MTU configuration during initialization may cause Serial over LAN disconnection	2.30.3000	2.30.8000
106.	IPMI connectivity	IPMI OOB communication lost during stress	2.30.3000	2.30.8000
107.	iperf stress test	Packet drops during iperf stress w/ different MSS	2.30.3000	2.30.8000
108.	Loss of connection to BMC upon firmware upgrade	Upgrading from firmware v2.10.3898 may cause loss of connection to BMC.	2.30.3000	2.30.8000
109.	Running rmod may cause unexpected behavior	Removing the mlx4_en driver using the "rmod" command, may cause unexpected behavior	2.30.3000	2.30.8000
110.	IPMI connectivity	On rare occasions, after stress of BMC cold reset, link failure might occur	2.30.3000	2.30.8000
111.	Flexboot 3.4.100-UEFI-4.0.410 BIOS menu	Enter the BIOS menu while using Flex-Boot 3.4.100- UEFI-4.0.410 may cause the server to stop responding (hang)	2.30.3000	2.30.8000

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
112.	ROL failure	ROL failure after disassembling the driver on the SUT	2.30.3000	2.30.8000
113.	IPMI link failure	IPMI link failure after disabling the WoL or disassembling the driver	2.30.3000	2.30.8000
114.	WoL and RoL issues	WoL and RoL issues caused when the IPMI is disabled	2.30.3000	2.30.8000
115.	A link flapping issue	Alignment marker arrival can no longer drop the link.	2.30.3000	2.30.8000
116.	IPMI - SOL traffic performance improvement	SOL with multiple data streams occasionally hang	2.30.3000	2.30.8000
117.	RoCE	RoCE does not function properly after running "ethtool ethX"	2.30.3000	2.30.8000
118.	PCI link errors	PCI link errors false indication. Cleared errors during PCIe link retraining	2.30.3000	2.30.8000
119.	PCIe speed change	Fixed a false indication for incoming PCIe speed change request	2.30.3000	2.30.8000
120.	No NC-SI after PXE teardown	Management transport was not supported, after PXE driver teardown (unload). Fixed the flow to issue software reset, after the driver was stopped	2.30.3000	2.30.8000
121.	Ports PLL calibration issue	PLL calibration were affected by operating point configuration	2.30.3000	2.30.8000
122.	Cable Info MAD issue	Wrong cable info was received when using the MC2210411-SR4 module	2.30.3000	2.30.8000
123.	Revision ID	Updated the Revision ID of Node Info and General Info MADs to reflect HW Rev ID instead of 0	2.30.3000	2.30.8000
124.	Port Error counters reset	Port error counters were not cleared upon XAUI/SGMII link up	2.30.3000	2.30.8000
125.	RDMA retransmission on ConnectX®-3 Pro	RDMA retransmission failed in specific scenario of receiving NAKs on ConnectX®-3 Pro due to bad static configuration	2.30.3000	2.30.8000
126.	NMI on PCIe Gen2 server	Fixed a PCIe Gen2 firmware flow to prevent NMI during hibernation on PCIe Gen2 server	2.30.3000	2.30.8000

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
127.	Access to closed resources	Fixed a possible access to unmapped resource memory	2.30.3000	2.30.8000
128.	Slow link establishment in NC-SI	Slow link establishment in NC-SI caused due to slow cable reading in boot	2.30.3000	2.30.8000
129.	PCIe speed change	Occasionally, a failure occurred in speed when changing to Gen2	2.30.3000	2.30.8000
130.	PXE teardown issue	PXE might halt during teardown	2.30.3000	2.30.8000
131.	InfiniBand loopback	InfiniBand loopback was blocked during link negotiation on the same port	2.11.0500	2.30.8000
132.	Voltage scaling	Fixed process voltage scaling issue	2.30.3000	2.30.8000
133.	DMA address 0x0	Fixed a possible read access to DMA address 0x0	2.11.0500	2.30.8000
134.	cqe issue	Fixed miss cqe issue due to interrupt moderation	2.11.0500	2.30.8000
135.	Cable reading issue	Fixed a rare cable reading issue upon cable insertion	2.30.3000	2.30.8000
136.	Linkup issue	Fixed a linkup issue against MSX60XX FDR switch	2.30.3000	2.30.8000
137.	LLR Vendor Specific MAD	LLR Dropped cell counter reported CSN error	2.11.0500	2.30.8000
138.	PXE boot failure	On rare occasions, PXE boot fails due to a firmware issue interfering with the PXE load from the card's flash. Upon failure, the following message is received: "Payload inaccessible - cannot continue"	2.30.3000	2.30.8000
139.	MLNX_OFED v2.0-2.0.3 reports ETH RX errors	MLNX_OFED v2.0-2.0.3 reported RX errors when the driver operated in the ETH only mode.	MLNX_OFED v2.0.2.0.3	2.30.3200
140.	SR-IOV guest communication channel error	Under certain conditions, SR-IOV guest experienced request timeouts and got stuck.	2.11.0500	2.30.3200
141.	No NC-SI after PXE teardown	Management transport was not supported, after PXE driver teardown (unload). Fixed the flow to issue software reset, after the driver was stopped.	2.11.0500	2.30.3000
142.	Revision ID	Updated the Revision ID of Node Info and General Info MADs to reflect hardware Rev ID instead of 0	2.11.0500	2.30.3000

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
143.	Ports PLL calibration issue	PLL calibration were affected by operating point configuration	2.11.0500	2.30.3000
144.	Port Error counters reset	Port error counters were not cleared upon XAUI/SGMII link up	2.11.0500	2.30.3000
145.	Receiver SerDes tuning	Enhanced the receiver SerDes tuning for 10GE, to support specific 10GbE QSFP to SFP+ splitter cables	2.11.0500	2.30.3000
146.	Access to closed resources	Fixed a possible access to unmapped resource memory	2.11.0500	2.30.3000
147.	Slow link establishment in NC-SI	Slow link establishment in NC-SI caused due to slow cable reading in boot	2.11.0500	2.30.3000
148.	PCIe speed change	Occasionally, a failure occurred in speed when changing to Gen2	2.11.0500	2.30.3000
149.	PXE teardown issue	PXE might halt during teardown	2.11.0500	2.30.3000
150.	InfiniBand loopback	InfiniBand loopback was blocked during link negotiation on the same port	2.11.0500	2.30.3000
151.	Voltage scaling	Fixed process voltage scaling issue	2.11.0500	2.30.3000
152.	DMA address 0x0	Fixed a possible read access to DMA address 0x0	2.11.0500	2.30.3000
153.	cqe issue	Fixed miss cqe issue due to interrupt moderation	2.11.0500	2.30.3000
154.	Cable reading issue	Fixed a rare cable reading issue upon cable insertion	2.11.0500	2.30.3000
155.	PXE boot failure	On rare occasions, PXE boot fails due to a firmware issue interfering with the PXE load from the card's flash. Upon failure, the following message is received: "Payload inaccessible - cannot continue"	2.11.0500	2.30.3000
156.	Momentarily packet drop on one port while the other port goes down	When the same VLAN is configured for both ports and one port goes down, the second port may suffer a momentarily packet drop	2.11.0500	2.30.3000
157.	Advanced Error Reporting	Fixes to Advanced Error Reporting according to the PTC (PCIe compliancy) Test failures	2.11.0500	2.30.3000
158.	10GbE link remains down after changing to a 40GbE link	When changing link connection between 40Gbe to 10GbE, the port might stay down until the next driver restart.	2.11.0500	2.30.3000

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
159.	Race in PCIe L1 flow	The device may enter an L1 power state before completing an incoming configuration request if it arrived before the power state change ack for Root Complex	2.11.0500	2.30.3000
160.	Wrong link speed after several cable re-insertions	During repeated cable reinsertion, the link may raise in a lower speed than expected/	2.11.0500	2.30.3000
161.	Changing port protocol from ETH to IB requires driver restart	Bad configuration of <code>ib_protocol</code> when setting the port to InfiniBand after exchanging it from Ethernet/RoCE on the same port	2.11.0500	2.30.3000
162.	Rare NMI error on HP servers when using PXE	On rare occasions, an NMI error is received when stopping PXE boot in the middle of an action on HP servers	2.11.0500	2.30.3000
163.	Race in PCI configuration handling	The system was unresponsive when a race between PCI configuration cycles handling in hardware and firmware occurred.	2.11.0500	2.30.3000
164.	RoCE breaks IPv6 traffic	IPv6 packets dropped while RoCE was enabled	2.11.0500	2.30.3000
165.	Link reset, HCA to SwitchX®	On rare link reset occasions, an HCA to SwitchX® link may remain in ARM state after SwitchX reboot, and not reach the Active state.	2.11.0500	2.30.3000
166.	40GbE switches link down upon repeated insertion and removing of cables	When repeatedly plugging out and plugging in cables to/from a 40GbE switch, the link remains down.	2.10.0800	2.11.0500
167.	FDR/QDR vs DDR switch	The link is raised as SDR rather than DDR	2.10.0800	2.11.0500
168.	0.5M, 1M Copper FDR cable support	0.5M, 1M Copper FDR cables (Paralink11, Paralink14) are not supported in back to back connection.	2.10.0800	2.11.0500
169.	Changing FDR cables from Fiber to Copper	Changing FDR cables from Fiber to Copper and vice versa might cause link speed degradation from FDR to FDR10 or FDR10 to QDR.	2.10.0800	2.11.0500
170.	Changing 40GbE connection	Changing 40GbE connection from back to back to a switch port and vice versa might cause the link to stay down in the ConnectX®-3 side.	2.10.0800	2.11.0500
171.	SDR vs QDR switch	InfiniBand link vs QDR switch rises as DDR or SDR	2.10.0800	2.11.0500

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
172.	QDR link up time	On rare occasions, the QDR link up time between an FDR card and a QDR device was up to 4 minutes	2.10.0800	2.11.0500
173.	QDR Link stability	QDR Link stability between ConnectX®-3 and InfiniScale® IV	2.10.0800	2.11.0500
174.	Signal integrity issues	Signal integrity issues in all speeds	2.10.0800	2.11.0500
175.	PCI correctable error	Fixed bad PCI reporting	2.10.0800	2.11.0500
176.	RoCE re-transmission	Not re-transmitting from the beginning of the message but from PSN NAK.	2.10.0800	2.11.0500
177.	RoCE	R-RoCE ignored SMAC check	2.10.0800	2.11.0500
178.	Function Level Reset (FLR)	FLR to PPF in SRIOV	2.10.0800	2.11.0500
179.		FLR in no-driver mode	2.10.0800	2.11.0500
180.	SR-IOV	Comchannel bug fix	2.10.0800	2.11.0500
181.	40GbE performance	40GbE full wire speed was not reached with former cards configuration for MCX314 and MCX313 40GbE cards. Fix: Increased core frequency to support 40GbE full wire speed. This change requires 1V input. Affected cards: MCX313A-BCB, MCX314A-BCB	2.10.0700	2.10.0800
182.		The ConnectX® card was not receiving frames with minimum IPG when configured to 40GbE. Fix: Updated firmware configuration. Affected cards: All cards capable of 40GbE speed	2.10.0700	2.10.0800
183.	GUIDs for RoCE	Ethernet only cards did not have GUIDs flashed on the device. This caused RoCE to fail once two or more cards were active on the same machine. Fix: The firmware generates the GUIDs from the flashed MACs for ETH only cards. Affected cards: All ETH only cards.	2.10.0700	2.10.0800
184.	40GbE signal integrity	Signal integrity improved in 40GbE speed.	2.10.0000	2.10.0700
185.				

Table 7 - Fixed Bugs List

Index	Issue	Description	Discovered in Release	Fixed in Release
186.	PCIe Gen3 receiver sensitivity due to BIOS changes	Due to BIOS changes required by latest Intel guidelines for PCIe Gen3, a sensitivity in PCIe eye_centering algorithm caused link degradation on some platforms.	2.10.0600	2.10.0620
187.	FDR, FDR10, 10GbE XFI/SFI signal integrity	Signal integrity improved for the following speeds: FDR, FDR10, 10GbE XFI/SFI.	2.10.0000	2.10.0600
188.				
189.	VDD_MODE GPIO changed after reset	In boards with VDD_MODE GPIO, the GPIO was dropped after reset hence, causing voltage to be reduced to 0.9v. Fix: The GPIO in boards with VDD_MODE GPIO is no longer dropped across resets.	2.10.0000	2.10.0600
190.	PCIe configuration type1 requests were dropped after a PCIe hot reset	After a PCIe hot reset, PCIe configuration type1 requests were dropped instead of being replied as unsupported. Fix: A proper notification is sent every time PCIe configuration type1 requests are dropped.	2.10.0000	2.10.0600

5 Firmware Changes and New Feature History

Table 8 - Firmware Changes and New Feature History

Firmware Version	Description
2.31.5050	<ul style="list-style-type: none"> • MAD: <ul style="list-style-type: none"> • Added support for GeneralInfo SMP MAD • Updated capability mask in GeneralInfo SMP/GMP MAD • Added support for PortCountersVL MAD • Added support for PortSamplesControl/PortSamplesResults/PortSamplesExtended MADs • INI: <ul style="list-style-type: none"> • Added Exponential Backoff Timer support. It is enabled via the <code>rtm_ini</code> parameter. The default value is 0 • Flow Steering (DMFS): <ul style="list-style-type: none"> • Added VLAN steering to Device Managed Flow Steering (DMFS) • Non-Volatile configuration tool: <ul style="list-style-type: none"> • Added support for Non-Volatile configuration of TLVs to set device attributes: <ul style="list-style-type: none"> • Query and set of configurations is available through PRM ACCEES_REG command • PRM ACCESS_REG command is now also supported through the <code>tools_hcr</code> command interface • Added support for MTF <code>mlxconfig</code> tool • Management protocols: <ul style="list-style-type: none"> • Added IPv6 support for NC-SI and IPMI Pass-Through • Added support for the same unicast MAC simultaneously for both IPMI and NC-SI • PCIe: <ul style="list-style-type: none"> • Added enhancements for receiver equalization in Gen3: <ul style="list-style-type: none"> • Enhancements are enabled by the INI. The default value is disabled. Please contact Mellanox support if required to enable it. • PCIe power optimizations for 8X/4X links • Side Band protocols: <ul style="list-style-type: none"> • SMBUS optimizations • Physical and Virtual Functions reset flows: <ul style="list-style-type: none"> • Added new Physical and Virtual Functions reset flows support • PXE support: <ul style="list-style-type: none"> • Added support for 64Bit BIOS mode • PRM: <ul style="list-style-type: none"> • Added IEEE802.3 CL73 autoneg support to the <code>QUERY_PORT</code> command. • Added factory MAC address reporting to the <code>Query_Port</code> command. • Added support for reverting virtual MAC configuration per port and restoring to factory MAC through <code>MOD_STAT_CFG</code> command. • Added support for inline TLV read through <code>MOD_STAT_CFG</code> command. • Added current MTU reporting to the <code>QUERY_PORT</code> command. • Added support query for additional MAC addresses per port (up to 7) through the <code>QUERY_PORT</code> command.

Table 8 - Firmware Changes and New Feature History

Firmware Version	Description
Rev 2.30.8000	<ul style="list-style-type: none"> • DMFS and GRE steering: Rule insertion adjustments • Removed DIF support from reported capabilities in QUERY_DEV_CAP PRM command • Flow control by DSCP priority for IPv4 • DMFS improvements: Insertion scheme enforcement and block loopback for InfiniBand • Added I2C resiliency support • Support for NC-SI over MCTP over SMBus • Added a flash access interface for persistent (non-volatile) configuration support • Added port BW arbitration configuration through the CONFIG_DEV command • Added IP-in-IP TCP checksum offload support • PCI Express compliancy Tx and Rx adjustments • Removed software limitations that were required for the use of Mellanox-certified FDR InfiniBand cables with Mellanox FDR InfiniBand adapters and switches. Please refer to "Memo: FDR 56Gb/s InfiniBand Cables" that was released on Dec/2013. Mellanox will offer an EXTENDED diagnostics support plan which will be available for mixed environments only and that will help identify issues they may encounter with the FDR installations. • Added support for 40GbE in WoL and pre-OS driver modes To enable this, add/change the following flags in the INI file in the IB and HCA tabs respectively: <ul style="list-style-type: none"> • restrict_max_eth_standby_speed = NO_RESTRICTION • slow_clock_enable = 0 • Bug fixes - see “Bug Fixes History” on page 21
2.30.3200	<ul style="list-style-type: none"> • Added support for FDR AOC MC2207312-XXX • Bug Fixes, see “Bug Fixes History” on page 21

Table 8 - Firmware Changes and New Feature History

Firmware Version	Description
2.30.3000	<ul style="list-style-type: none"> • Added support for the UPDATE_QP PRM command • Added support for resetting the modified MAC addresses in the standby mode by the MOD_STAT_CFG command • Added support for receiving TCP and UDP truncated packets of certain type • Added support for 56GbE in all devices supporting 40GbE <ul style="list-style-type: none"> • Establishing 56GbE link with SwitchX® requires 56GbE enablement on the switch • Establishing 56GbE link on back to back setup requires additions to the INI. • For further information, please contact Mellanox Support • mlxconfig tool in not supported in this release • RDMA/RoCE read retransmission improvement • PFC thresholds improvements • PCI speed_change flow improvement • Added support for DIAG_RPRT per port • Added PCIe Polling Compliance mode • RoCE default configuration fixes: <ul style="list-style-type: none"> • ethertype now updated per port at SET_PORT • The default value of rroce.ip_next_protocol is 0xfe • Increased the number of extended interface counters (max_if_cnt_extended) to 0x80, as reported in QUERY_DEV_CAP PRM command • Improved link parallel detection calibration of 40GbE • Modified Dell Baldur INI • Added support for PFC counters in DUMP_ETH_STAT PRM command • Fixed wrong reporting of RSS context in QUERY_FW of RSS QP • Added Sniffer QP support on Port#1 • Device managed Flow Steering performance enhancements • Disabled "pkt.dmac==qp.mac" for RoCE/RoCE over IP • Added missing loopback blocking for device managed Flow Steering • Fixed SET_PORT.mac_table configuration issue which caused minor packet loss on Port A when working in bonding mode and closing Port B. • Fixed issues with NC-SI commands reason codes • Fixed the insertion of L4 head rule in device managed Flow Steering • Added to the INI for 10/40GbE parallel detect Serdes parameters • Added support for "reset upon parity error" • Added support for 40Gb/s MC2210411-SR4 optical module • Fixed interoperability issue with the Intel 12300 switch using firmware version 6.1.0.1.11 • Improved QDR link stability when connected to InfiniScale® IV and SwitchX® switches

Table 8 - Firmware Changes and New Feature History

Firmware Version	Description
2.11.0500	<ul style="list-style-type: none"> • Added SR-IOV support • Added VPI auto-sensing support • Mellanox Link Property Negotiation (MLPN) • Enables ports to negotiate link properties between Mellanox devices. • The MLPN is activated by INI. • CR4 + KR4 • Auto-negotiation 40GBASE CR4 and KR4 as described in IEEE 802.3. • CR4 + KR4 is enabled by the INI. • 1GbE Clause37 • Auto-negotiation 1000BASE-X as described in IEEE 802.3 clause 37. • 1GbE Clause37 is enabled by default. • cable_info MAD extension for more I2C addressed • Enables different address for cable access through I2C. • cable_info MAD extension is enabled by the INI. • secure_host • smp firewall described in the PRM • Activated by the INI (active by default). • cq_2_eq mapping command • Modifies EQ by MODIFY_CQ command, described in the PRM • multi-function reserved lkey • Described in the PRM. • increase CQE timestamp to 48bit • Described in the PRM. • 56Gb Ethernet (proprietary) - Beta level • Activated by the INI (disabled by default). • mlxconfig - Beta level (requires MFT 3.0.0-3 or above) • Modifies the device cfg • FMR for SRIOV - Beta level • Described in the PRM. • Power reduction in PCI Gen3 • Fixed general_info MAD • “Bug Fixes History” on page 21
2.10.0800	<ul style="list-style-type: none"> • Bug fixes - see “Bug Fixes History” on page 21
2.10.0700	<ul style="list-style-type: none"> • Bug fixes - see “Bug Fixes History” on page 21
2.10.0000	<ul style="list-style-type: none"> • InfiniBand: <ul style="list-style-type: none"> • FDR • FDR10 • QDR • SDR • DDR • Ethernet: <ul style="list-style-type: none"> • 1GigE • 10GigE XAUI • 10GigE XFI/SFI • 40GigE • PCI Express 3.0, with backwards compatibility with v2.0 and v1.1 • Huge pages • ConnectX®-3 firmware includes all ConnectX®-2 cards supported features

Appendix A: Flexboot Changes and New Features

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

Table 9 - FlexBoot Changes and New Feature

Version	Description
Rev 3.4.306	<ul style="list-style-type: none"> • Added validation script for the released ROMs • Added the option to always keep SAN hook to enable WIN install on iSCSI target • Added compilation flag around the flash readout. • Added URI Boot retry. Default retries = 0. • Added Unmap MPT command in teardown. • Added support for HII iSCSI configuration. • Added 64-bit PCI BAR support (Large bar). • Added the option added for running PXE with promiscuous VLAN. • Re-added COMBOOT image support by default. • Enabled pages-function handling in Connect-IB initialization stage to work according to the PRM. • Applied additional patches from ipxe.org • Updated the window even if ACK does not acknowledge new data. • Modified the error print to debug print. • Modified the printed string when initializing devices. • Modified the error print. Added additional information to make the output more user-friendly. • Changed the size of the domain name array to 0xfd. • Disabled the waiting period for link up on trunk-net-device when VLAN is enabled on port. • Removed unsupported EQ event in Connect-IB® • Fixed an issue for TLV with length 0. • Fixed an issue related to sync VLAN IRQ operation with trunk IRQ operation. • Fixed an issue which enabled a netdevice (VLAN) to open/close twice. • Fixed an issue which prevented the iSCSI initiator's name from being received from HII. • Fixed an issue related to dual port adapters; occasionally, booting from the second port resulted in TFTP download failure when the first port was already linked up with DHCP, and has received a TFTP address. • Fixed an issue which caused PXE boot failure when using a filename if iSCSI rootpath is set. • Fixed an issue which prevented the device to PXE boot from the 2nd port if first port was already downloaded. • Fixed compilation issue. • Fixed a broken VLAN issue. • Fixed a retry issue when the value is infinite.

Table 9 - FlexBoot Changes and New Feature

Version	Description
Rev 3.4.225	<ul style="list-style-type: none"> • Added additional information to the error print output • Added compilation flag around the flash readout • Added URI Boot retry. Default retries = 0 • Added Unmap MPT command in teardown • Added 64-bit PCI BAR support • Added an option for running PXE with promiscuous VLAN • Added support for HII iSCSI configuration • Enlarged the mailbox size to 4kb • Enlarged the number of WQE to 64 (from 4) • Enabled multiple DHCP offers to be received before proceeding to request state • Changed the size of the domain name array to 0xfd • Changed error print to debug print • Changed printed string when initializing devices • Kept the SAN connection permanently open to enable Windows install on iSCSI target even when the iSCSI target is empty • Re-added COMBOOT image support by default • Prevented a netdevice (VLAN) from opening/closing twice • Removed unsupported EQ event in Connect-IB® • Disabled the waiting time for link up on trunk net device when VLAN is enabled on a port • Fixed sync VLAN IRQ operation with trunk IRQ operation • Fixed iSCSI initiator's name retrieval from HII issue • Fixed an issue caused in dual port adapters, when the first port was already linked up with DHCP, and had received a TFTP address. Booting from the second port resulted in TFTP download failure. • Fixed retry issue when the value is infinite • Fixed a TLV with length 0 issue • Fixed a PXE boot failure issue occurred when using a filename when iSCSI rootpath is set • Fixed "Impossible to PXE boot from 2nd port if first port already downloaded." issue • Fixed compilation issue • Fixed broken VLAN support issues
Rev 3.4.151	<ul style="list-style-type: none"> • Enlarged the mailbox size to 4kb • Enlarged the number of WQE to 64 (from 4) • Enabled multiple DHCP offers to be received before proceeding to request state
Rev 3.4.146	<ul style="list-style-type: none"> • Fixed memory corruption issues • Modified TLV flash access • Added additional WQ
Rev 3.4.142	<ul style="list-style-type: none"> • Enabled firmware to handle the link state with the Subnet Manager • Updated the DHCP class code to NONE • Added flash access capability for reading software-to-software configurations • Enabled DHCP validation of MAC address and XID for a unique tuple • Improved randomness algorithm for DHCP XID

Table 9 - FlexBoot Changes and New Feature

Version	Description
Rev 3.4.112	<ul style="list-style-type: none"> • Broadcast responses for firewall support • Enabled request broadcast responses from DHCP server to support firewall.
Rev 3.4.100	<ul style="list-style-type: none"> • OCSD activation initiation change • Moved the OCSD activation initiation from the FlexBoot to the CLP code. This enables the OCSD • activation to no longer be dependent on the FlexBoot being enabled in the servers's BIOS configuration. • Messages' improvement • Made the FlexBoot on-screen notification messages more informative and user friendly. • FlexBoot and CLP merge improvement • Improved the process of merging the FlexBoot and CLP codes together. • PXE and UFI merge capability • Added the ability to merge the PXE image with a UFI image. • Supported servers • Added FlexBoot support capabilities to several new non-HP servers. • Use of newer iPXE version • Moved to use a newer iPXE version as the basis for the Flexboot release. • Fixed "no more network devices" issues during Flexboot.