



Mellanox ConnectX®-3 Firmware Release Notes

Rev 2.11.1250

NOTE:

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



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

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 2013. Mellanox Technologies. All Rights Reserved.

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

Connect-IB™, FabricIT™, Mellanox Open Ethernet™, MetroX™, MetroDX™, ScalableHPC™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Table of Contents	3
List of Tables	4
Chapter 1 Overview	5
1.1 Supported Cables and Modules	5
1.2 Tested Cables and Modules	6
1.3 Tested Switches	7
1.4 Tools, Switch Firmware and Driver Software	9
1.5 Revision Compatibility	9
1.6 Supported Devices	9
Chapter 2 Changes and Major New Features	10
2.1 Changes in Rev 2.11.1250 From Rev 2.11.0500	10
2.2 Changes in Rev 2.11.0500 From Rev 2.11.0000	10
2.3 Changes in Rev 2.11.0000 From Rev 2.10.5380	11
2.4 Changes in Rev 2.10.5380 From Rev 2.10.2284	11
2.5 Changes in Rev 2.10.2284 From Rev 2.10.2280	11
2.6 New Features in Rev 2.10.2280	12
Chapter 3 Known Issues	13
Chapter 4 Bug Fixes History	14

List of Tables

Table 1:	Tested Cables and Modules	6
Table 2:	Switches Tested	7
Table 3:	Supported PSIDs	9
Table 4:	Known Issues	13
Table 5:	Fixed Bugs List	14

1 Overview

These are the release notes for firmware Rev 2.11.1250. Firmware revision Rev 2.11.1250 supports the following features:

- InfiniBand - FDR, QDR, SDR
- Ethernet - 40GbE, 10GbE and 1GbE
- PCIe 3.0 and PCIe 2.0
- PXE
- UEFI



Firmware Rev 2.11.1250 is compatible with MFT tools v2.7.1 and above. Using older versions of MFT tools will cause failure during image creation.

1.1 Supported Cables and Modules

Please refer to the Mellanox Products Approved Cable Lists document (Doc Nr. 3796) for the list of supported cables.

http://www.mellanox.com/related-docs/user_manuals/Mellanox_approved_cables.pdf

1.2 Tested Cables and Modules

Table 1 - Tested Cables and Modules

Speed	OPN #	Description	Vendor
FDR	MC2207130-00A-V	FDR InfiniBand QSFP passive copper cable, 0.5m	Voalex
FDR	MC2207130-001-V	FDR InfiniBand QSFP passive copper cable, 1m	Voalex
FDR	MC2207130-002-V	FDR InfiniBand QSFP passive copper cable, 2m	Voalex
FDR	MC2207128-003-V	FDR InfiniBand QSFP passive copper cable, 3m	Voalex
FDR	MC2207126-004-V	FDR InfiniBand QSFP assembled optical cable, 4m	Voalex
FDR	MC2207310-010-E	FDR InfiniBand QSFP assembled optical cable, 10m	Emcore - Optic FDR
FDR	MC2207310-030-E	FDR InfiniBand QSFP assembled optical cable, 30m	Emcore optic FDR
FDR	MC2207310-100-E	FDR InfiniBand QSFP assembled optical cable, 100m	Emcore optic FDR
40GbE	MC2210130-001	40GE QSFP passive copper cable, 1m	Voalex
40GbE	MC2210130-002	40GE QSFP passive copper cable, 2m	Voalex
40GbE	MC2210128-003	40GE QSFP passive copper cable, 3m	Voalex
40GbE	MC2210126-004	40GE QSFP passive copper cable, 4m	Voalex
40GbE	MC2210310-020	40GE QSFP assembled optical cable, 20m	Emcore optic
QDR/ FDR10	MC2206130-00A	Mellanox copper cable, up to IB QDR/FDR10 (40Gb/s), 4X QSFP, 0.5 meter	Voalex
QDR/ FDR10	MC2206130-001	Mellanox copper cable, up to IB QDR/FDR10 (40Gb/s), 4X QSFP, 1 meter	Voalex
QDR/ FDR10	MC2206130-003	Mellanox copper cable, up to IB QDR/FDR10 (40Gb/s), 4X QSFP, 3 meter	Voalex
QDR/ FDR10	MC2206126-005	Mellanox copper cable, up to IB QDR/FDR10 (40Gb/s), 4X QSFP, 5 meter	Voalex
QDR/ FDR10	MC2206310-010-E	QDR/FDR10 InfiniBand QSFP assembled optical cable, 10m	Emcore optic QDR/FDR10
QDR/ FDR10	MC2206310-030-E	QDR/FDR10 InfiniBand QSFP assembled optical cable, 30m	Emcore optic QDR/FDR10
QDR/ FDR10	MC2206310-100-E	QDR/FDR10 InfiniBand QSFP assembled optical cable, 100m	Emcore optic QDR/FDR10

Table 1 - Tested Cables and Modules

Speed	OPN #	Description	Vendor
SDR/DDR	MC1104130-002-A	Mellanox CX4 Passive copper SDR/DDR cables 2m	Volex
SDR/DDR	MC1104130-003-A	Mellanox CX4 Passive copper SDR/DDR cables 3m	Volex
SDR/DDR	MC1104128-005-A	Mellanox CX4 Passive copper SDR/DDR cables 5m	Volex
SDR/DDR	MC1204130-001-V	Mellanox QSFP to CX4 (Q2C) Passive copper SDR/DDR cables 1m	Volex
SDR/DDR	MC1204130-002-V	Mellanox QSFP to CX4 (Q2C) Passive copper SDR/DDR cables 2m	Volex
SDR/DDR	MC1204130-003-V	Mellanox QSFP to CX4 (Q2C) Passive copper SDR/DDR cables 3m	Volex
SDR/DDR	MC1204128-005-V	Mellanox QSFP to CX4 (Q2C) Passive copper SDR/DDR cables 5m	Volex

1.3 Tested Switches

Table 2 - Switches Tested

Speed	OPN # /Name	Description
10/40Gb/s	Force 10 S4810P-AC	48-port 10Gb/40Gb Switch
10/40Gb/s	Cisco 3064	48-port 10Gb/40Gb Switch
10/40Gb/s	Arista 7050Q	16-port 40Gb Switch
10/40Gb/s	Arista 7050S	48-port 10Gb/40Gb Switch
10Gb/s	Fujitsu 10GbE (XAUI)	Ethernet Switch 24 ports, 20xCX4 and 4xQSFP
10Gb/s	HP ProCurve 6600-24XG	24-port 10GbE switch
40Gb/s	Mellanox MSX1036B-1BFR	SwitchX® based 36-port QSFP 40GigE 1U Ethernet
IB DDR	F-X430044	24-port DDR-Switch
IB SDR	F-X430060	24-port SDR-Switch
IB QDR	QLogic 12300	36-Port 40Gb QDR Infiniband Switch, Management Module, Dual Power

Table 2 - Switches Tested

Speed	OPN # /Name	Description
IB QDR	Mellanox MIS5030Q	InfiniScale® IV QDR IS5030/IS5035 36-port 40Gb/s InfiniBand Switch Systems
IB QDR	Mellanox QDR-Switch 4036	InfiniScale® IV QDR Mellanox Grid Director 4036 36-Port QDR InfiniBand Switch - Part ID: VLT-30011
IB FDR	Mellanox 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	Mellanox 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

1.4 Tools, Switch Firmware and Driver Software

Firmware Rev 2.11.1250 is tested with the following tools, SwitchX® firmware, and driver software:

- Driver versions:
 - MLNX-OFED 1.5.3-3.0.0 and above
 - WinOF 3.0.0 and above
 - VMware ESXi 5.1 Ethernet
- MFT for Linux version: 2.7.0a and higher
- MFT for Windows version: 2.7.1 and higher
- FlexBoot version 3.4.112 for Ethernet adapter cards
- SwitchX® software: MLNX-OS managed version 3.2.0500 and higher. A required version for ConnectX® to SwitchX 40GbE link

1.5 Revision Compatibility

Firmware fw-ConnectX3 Rev 2.11.1250 complies with the following programmer's reference manual:

- *ConnectX Programmer's Reference Manual (PRM), Rev 1.0 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 Supported Devices

This firmware supports the devices and protocols listed in Table 3. 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 3 - Supported PSIDs

Device Part Number	PSID	Device Name	Supported Protocols
0X8GJ4	DEL1120000023	MCX343A-XCAA ConnectX-3 Eth; single port SFP+ ; 10Gb/s card for Dell	Ethernet
0T483W	DEL0A30000019	MCX383A-FCxx ConnectX-3 VPI single port QSFP; FDR IB (56Gb/s) and 40GigE I/O card	VPI
01965W	DEL1110000023	MCX384A-BCAA ConnectX-3 Eth; dual port QSFP; 40Gb/s Mezz card for Dell	Ethernet
0K75R1	DEL1110000019	MCX384A-FCAA ConnectX-3 VPI; dual port QSFP; FDR (56Gb/s) Mezz card for Dell	VPI

2 Changes and Major New Features

2.1 Changes in Rev 2.11.1250 From Rev 2.11.0500

- Increased the number of extended interface counters (`max_if_cnt_extended`), as reported in `QUERY_DEV_CAP PRM` command to 0x80
- Added to the INI for 10/40GbE parallel detect Serdes parameters
- Added IPMI over SMBUS support
- Added NC-SI support
- 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 -Beta level
- 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 NCSI commands reason codes
- Fixed the insertion of L4 head rule in device managed flow steering

2.2 Changes in Rev 2.11.0500 From Rev 2.11.0000

- 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.
- `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).
- multi-function reserved lkey
Described in the PRM.
- Increased CQE timestamp to 48bit
Described in the PRM.
- 56Gb Ethernet (proprietary) - Beta level
Activated by the INI (disabled by default).
- FMR for SRIOV - Beta level
Described in the PRM.

- Device managed flow steering
Described in the PRM.
- Fixed general_info MAD

2.3 Changes in Rev 2.11.0000 From Rev 2.10.5380

- Added a workaround for QPtimers failure in write-back timer ctx
- Added support for changing the QP dynamic SL in rts2rts command as described in PRM 1.0
- Added a new link fail policy when LLR enabled
- Enabled verspace completion sending with the correct completer ID
- Added msix vectors mask refreshing for Virtual Function after cfg write cycle
- See “Bug Fixes History” on page 14.

2.4 Changes in Rev 2.10.5380 From Rev 2.10.2284

Firmware version 2.10.5380 includes the following features:

- Device managed flow steering
- 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.
- cq_2_eq mapping command
Modifies EQ by MODIFY_CQ command, described in the PRM
- reserved_lkey
- cq_timestamp
- 56GbE - Beta
- mlxconfig - Beta level (requires MFT 3.0.0-3 or above)
Modifies the device cfg.
- FMR for SRIOV - Beta
- RoCE LAG
- Dual 2x 10GbE on the same quad
- Power reduction in PCI Gen3

2.5 Changes in Rev 2.10.2284 From Rev 2.10.2280

- Fixed general_info MAD

2.6 New Features in Rev 2.10.2280

- Initial release

3 Known Issues

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

Table 4 - Known Issues

Index	Issue	Description	Current Implemented Workaround in FW	Scheduled Release (fix)
1.	Changing from an LLR to non-LLR requires driver restart	Driver restart required when switching from InfiniBand link with LLR enabled to InfiniBand link w/o LLR (for example: between SwitchX® and GD4036).	N/A	Future Release
2.	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
3.	VLAN filter	VLAN filter is currently not supported.	N/A	Future Release
4.	PXE	PXE is currently not supported in 40GbE in VPI cards	N/A	Future Release
5.	PXE	PXE is currently not supported in 10GbE QSFP to SFP+ hybrid cable	N/A	Future Release
6.	40GbE link up waiting time	Occasionally, 40GbE link takes longer to establish link when CR4 is used.	N/A	Future Release
7.	PCIe failure on temperature shock 10C/min	Gen2 failure at temperature sweep at 10C/min	N/A	Future Release
8.	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

4 Bug Fixes History

Table 5 lists the bugs fixed in this release.

Table 5 - Fixed Bugs List

Issue	Description	Discovered in Release	Fixed in Release
NodeInfo device revision	NodeInfo device revision field incorrect on ConnectX®-3 card	2.10.0800	2.11.1200
NMI issues	Rare Red Screen NMI when OCSD is enabled	2.11.0500	2.11.1200
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
4036 Switch	Occasionally, the link up time of the 4036 switch might take up to 4 minutes.	2.10.0800	2.11.0500
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.	2.10.0800	2.11.0500
FDR/QDR vs DDR switch	The link is raised as SDR rather than DDR	2.10.0800	2.11.0500
NCSI	NCSI compliancy	2.10.0800	2.11.0500
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
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
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
SDR vs QDR switch	InfiniBand link vs QDR switch rises as DDR or SDR	2.10.0800	2.11.0500
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
QDR Link stability	QDR Link stability between ConnectX®-3 and InfiniScale® IV	2.10.0800	2.11.0500
Signal integrity issues	Signal integrity issues in all speeds	2.10.0800	2.11.0500
PCI correctable error	Fixed bad PCI reporting	2.10.0800	2.11.0500
RoCE re-transmission	Not re-transmitting from the beginning of the message but from PSN NAK.	2.10.0800	2.11.0500
RoCE	R-RoCE ignored SMAC check	2.10.0800	2.11.0500

Table 5 - Fixed Bugs List

Issue	Description	Discovered in Release	Fixed in Release
Function Level Reset (FLR)	FLR to PPF in SRIOV	2.10.0800	2.11.0500
	FLR in no-driver mode	2.10.0800	2.11.0500
SR-IOV	Comchannel bug fix	2.10.0800	2.11.0500
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
	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
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
40GbE signal integrity	Signal integrity improved in 40GbE speed.	2.10.0000	2.10.0700
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
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
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
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