



# **Mellanox ConnectX<sup>®</sup>-3 Firmware (fw-ConnectX3) Release Notes**

Rev 2.11.0550

## 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](http://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](http://www.mellanox.com)  
 Tel: +972 (0)74 723 7200  
 Fax: +972 (0)4 959 3245

© Copyright 2012. Mellanox Technologies. All Rights Reserved.

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

Connect-IB™, FabricIT™, MLNX-OST™, ScalableHPC™, Unbreakable-Link™, UFM™ and Unified Fabric Manager™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

# Table of Contents

<b>Chapter 1 Overview</b> .....	<b>5</b>
1.1 Supported Cables and Modules .....	5
1.2 Tools, Switch Firmware and Driver Software .....	5
1.3 Supported Devices .....	5
1.4 Revision Compatibility .....	6
<b>Chapter 2 Supported Features</b> .....	<b>7</b>
2.1 Changes and Major New Features in Rev 2.11.0550 .....	7
2.2 Changes and Major New Features in Rev 2.10.0800 .....	8
2.3 Changes and Major New Features in Rev 2.10.0700 .....	8
2.4 Changes and Major New Features in Rev 2.10.0000 .....	8
<b>Chapter 3 Known Issues</b> .....	<b>9</b>
<b>Chapter 4 Bug Fixes History</b> .....	<b>11</b>

# List of Tables

Table 1: Supported PSIDs .....	6
Table 2: Known Issues .....	9
Table 3: Fixed Bugs List .....	11

# 1 Overview

These are the release notes for the ConnectX®-3 adapters firmware, fw-ConnectX3 Rev 2.11.0550. This firmware supports the following protocols:

- InfiniBand - FDR, QDR, SDR, DDR
- Ethernet - 1GigE, 10GigE and 40GigE
- PCI Express 3.0, supporting backwards compatibility for v2.0 and v1.1



Firmware Rev 2.11.0550 is compatible with MFT tools v2.7.0a 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](http://www.mellanox.com/related-docs/user_manuals/Mellanox_approved_cables.pdf)

## 1.2 Tools, Switch Firmware and Driver Software

Firmware Rev 2.11.0550 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 5.0 Ethernet
- MFT for Linux version: 2.7.0a and higher
- MFT for Windows version: 2.7.1 and higher
- FlexBoot version 3.4.000 for Ethernet adapter cards
- SwitchX® software: MLNX-OS managed version 3.2.0500. A required version for ConnectX® to SwitchX 40GbE link

## 1.3 Supported Devices

This firmware supports the devices and protocols listed in Table 1. 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 1 - Supported PSIDs**

Device Part Number	PSID	Device Name	Supported Protocols
0J05YT	DEL0A10210018	ConnectX-3 Dual-port QDR Mezzanine I/O Card	InfiniBand
0K6V3V	DEL0A20210018	ConnectX-3 Dual-port FDR Mezzanine I/O Card	InfiniBand
0P90JM	DEL0A40000028	ConnectX-3 Dual-port FDR10 Mezzanine I/O Card	InfiniBand

## 1.4 Revision Compatibility

Firmware fw-ConnectX3 Rev 2.11.0550 complies with the following programmer's reference manual:

- *ConnectX Programmer's Reference Manual (PRM), Rev 0.90 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`.

## 2 Supported Features

### 2.1 Changes and Major New Features in Rev 2.11.0550

- Added SRI-OV 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.
- Device managed flow steering  
Described in the PRM.
- Power reduction in PCI Gen3
- Fixed general\_info MAD
- “Bug Fixes History” on page 11.

## 2.2 Changes and Major New Features in Rev 2.10.0800

- No new features
- “Bug Fixes History” on page 11.

## 2.3 Changes and Major New Features in Rev 2.10.0700

- No new features
- “Bug Fixes History” on page 11.

## 2.4 Changes and Major New Features in Rev 2.10.0000

- InfiniBand:
  - FDR
  - FDR10
  - QDR
  - SDR
  - DDR
- Ethernet:
  - 1GigE
  - 10GigE XAUI
  - 10GigE XFI/SFI
  - 40GigE
- PCI Express 3.0, with backwards compatibility with v2.0 and v1.1
- Huge pages



ConnectX®-3 cards with firmware 2.10.0000 will include all ConnectX®-2 cards supported features.



### 3 Known Issues

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

**Table 2 - Known Issues**

Index	Issue	Description	Current Implemented Workaround in FW	Scheduled Release (fix)
1.	sense_port failure	Ethernet cards fail to work with mlnx_ofed unless the do_sense is disabled in the INI	Add 'do_sense = false' in the INI under the [IB] section	Future Release
2.	SBR assertion	SBR should be asserted for a minimum of 50 milliseconds for the ConnectX®-3 adapters	N/A	N/A
3.	NC-SI support	NC-SI is not supported in the current release	N/A	Future Release
4.	40GbE is not supported in Auto-Sensing	Auto-Sensing is not supported with 40GbE connections in VPI cards	N/A	Future Release
5.	Wrong hardware revision ID	Wrong hardware revision ID reported in NodeInfo MAD	N/A	Future Release
6.	Changing from an LLR to non-LLR requires driver restart	Driver restart required when switching from InfiniBand FRDR link with LLR enabled to InfiniBand link w/o LLR (for example: between SwitchX® and GD4036).	N/A	Future Release
7.	VLAN filter	VLAN filter is currently not supported.	N/A	Future Release
8.	Bloom filter	Bloom filter is currently not supported.	N/A	Future Release
9.	PXE	PXE is currently not supported in 40GbE in VPI cards	N/A	Future Release
10.		PXE is currently not supported in QSFP to SFP+ hybrid cable	N/A	Future Release

**Table 2 - Known Issues**

Index	Issue	Description	Current Implemented Workaround in FW	Scheduled Release (fix)
11.	Firmware down-grade	<p>When downgrading from firmware v2.11.0000 and using MFT 3.0.0-3, the following message is displayed due to the mlxconfig tool:</p> <pre>You are trying to override configurable FW by non-con- figurable FW. If you continue, old FW con- figurations 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</pre>	N/A	Future Release

## 4 Bug Fixes History

Table 3 lists the bugs fixed in this release.

**Table 3 - Fixed Bugs List**

Issue	Description	Discovered in Release	Fixed in Release
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	2.11.0550
FDR/QDR vs DDR switch	The link is raised as SDR rather than DDR	2.10.0800	2.11.0550
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	2.11.0550
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	2.11.0550
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	2.11.0550
SDR vs QDR switch	InfiniBand link vs QDR switch rises as DDR or SDR	2.10.0800	2.11.0550
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	2.11.0550
QDR Link stability	QDR Link stability between ConnectX®-3 and InfiniScale® IV	2.10.0800	2.11.0550
Signal integrity issues	Signal integrity issues in all speeds	2.10.0800	2.11.0550
PCI correctable error	Fixed bad PCI reporting	2.10.0800	2.11.0550
RoCE re-transmission	Not re-transmitting from the beginning of the message but from PSN NAK.	2.10.0800	2.11.0550
RoCE	R-RoCE ignored SMAC check	2.10.0800	2.11.0550
Function Level Reset (FLR)	FLR to PPF in SRIOV	2.10.0800	2.11.0550
	FLR in no-driver mode	2.10.0800	2.11.0550
SR-IOV	Comchannel bug fix	2.10.0800	2.11.0550

**Table 3 - Fixed Bugs List**

Issue	Description	Discovered in Release	Fixed in Release
40GbE performance	40GbE full wire speed was not reached with former cards configuration for MCX314 and MCX313 40GbE cards. <b>Fix:</b> Increased core frequency to support 40GbE full wire speed. This change requires 1V input. <b>Affected cards:</b> MCX313A-BCB, MCX314A-BCB	2.10.0700	2.10.0800
	The ConnectX® card was not receiving frames with minimum IPG when configured to 40GbE. <b>Fix:</b> Updated firmware configuration. <b>Affected cards:</b> 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. <b>Fix:</b> The firmware generates the GUIDs from the flashed MACs for ETH only cards. <b>Affected cards:</b> 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. <b>Fix:</b> 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. <b>Fix:</b> A proper notification is sent every time PCIe configuration type1 requests are dropped.	2.10.0000	2.10.0600