



Mellanox ConnectX[®]-2 Firmware (fw-ConnectX2) Release Notes

Rev 2.9.1200

Last Updated. October 10, 2013

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™, ExtendX™, FabricIT™, Mellanox Open Ethernet™, Mellanox Virtual Modular Switch™, 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
Chapter 1 Overview	4
1.1 Supported Devices	4
1.2 Tested Cables and Modules	5
1.3 Revision Compatibility	5
Chapter 2 Changes and Major New Features	6
2.1 Changes in Rev 2.9.1200 From Rev 2.9.1000	6
2.2 Changes in Rev 2.9.1000 From Rev 2.8.0600	6
2.3 Changes in Rev 2.8.0600 From Rev 2.8.0000	6
2.4 Changes in Rev 2.8.0000 From Rev 2.7.9112	6
2.5 Changes in Rev 2.7.9112 From Rev 2.7.8100	7
2.6 Changes in Rev 2.7.8100 From Rev 2.7.200	7
2.7 Changes in Rev 2.7.700 From Rev 2.7.200	7
2.8 Changes in Rev 2.7.200 From Rev 2.7.000	7
Chapter 3 History of Bug Fixes	8
Chapter 4 Known Issues	10

1 Overview

These are the release notes for the ConnectX®-2 and ConnectX®-2 EN adapters firmware, fw-ConnectX2 Rev 2.9.1200. This firmware supports the following protocols:

- InfiniBand
- Ethernet
- Fibre Channel over Ethernet (FCoE) – FCoE is at a beta level
- Virtual Protocol Interconnect (VPI) – this capability enables ConnectX®-2 devices to support the InfiniBand, Ethernet and DCE network standards, including auto-sensing of the network protocol to which each device port is connected. This feature is not available with ConnectX®-2 EN.

For the most updated list of adapter cards supported, visit the firmware download pages via <http://www.mellanox.com>.

Note:

1. Firmware Rev 2.9.1200 is compatible with MFT tools v2.7.0a and higher. Using older versions of MFT tools will cause failure during image creation.

1.1 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 (Sheet 1 of 2)

Device Part Number	PSID	Device Name	Supported Protocols
MNPA19-XSR/XTR	MT_0F60110010	ConnectX®-2 EN network interface card, single-port SFP+, PCIe2.0 5.0GT/s, mem-free, short/tall bracket, RoHS R6	Ethernet
MNPH29C-XSR	MT_0DA0110010	ConnectX®-2 EN network interface card, dual-port SFP+, PCIe2.0 x8 5.0GT/s, short bracket, RoHS R6	Ethernet
MNPH29C-XTR	MT_0DB0120010	ConnectX®-2 EN network interface card, dual-port SFP+, PCIe2.0 x8 5.0GT/s, tall bracket, RoHS R6	Ethernet

Table 1 - Supported PSIDs (Sheet 2 of 2)

Device Part Number	PSID	Device Name	Supported Protocols
MNPH29D-XTR	MT_0F30120010	ConnectX®-2 EN network interface card, dual-port SFP+, PCIe2.0 x8 5.0GT/s, tall bracket, RoHS R6	Ethernet
MNPH29D-XSR	MT_0F30120010	ConnectX®-2 EN network interface card, dual-port SFP+, PCIe2.0 x8 5.0GT/s, short bracket, RoHS R6	Ethernet

1.2 Tested 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

1.3 Revision Compatibility

Firmware fw-ConnectX2 Rev 2.9.1200 complies with the following programmer's reference manual:

- *ConnectX Programmer's Reference Manual (PRM), Rev 0.80 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 Changes and Major New Features

2.1 Changes in Rev 2.9.1200 From Rev 2.9.1000

- Enabled `fast_drop` mode
- Improved multicast performance
- Improved signal integrity for short cables
- Added ether-wake support for Wake on LAN (WoL)

2.2 Changes in Rev 2.9.1000 From Rev 2.8.0600

Added support for the following features:

- New steering mechanism according to Ethertype and VLAN
- Header separation
- IPoIB ETHoIB offloads
- PCI Function Level Reset
- MPI performance enhancement: #97601 enigma bit
- Added INI to enable new LED behavior: `dual_color_led_mode_en`
- SRIOV beta support for ConnectX-2. To enable it, set `sriov_en` to TRUE in the INI
- Thermal sensing MAD

2.3 Changes in Rev 2.8.0600 From Rev 2.8.0000

- Bug Fixes - see Section 3, “Known Issues,” on page 8

2.4 Changes in Rev 2.8.0000 From Rev 2.7.9112

Added support for the following features:

- NC-SI over Ethernet
- QSFP-SFP+ module
- Signal integrity improvements in IBTA 1.2.1 mode (Attenuation Based Algorithm)
- Ethernet Vlan stripping
- Eye opener MAD
- Cable info MAD
- CQ.OI to the `MODIFY_CQ` command to enable CORE-Direct
- Thermal warning event
- `SET_NODE` command
- PCIe device Serial Number capability
- LoosyVI MAD

2.5 Changes in Rev 2.7.9112 From Rev 2.7.8100

- Added PCIe Multifunction support

- Added WOL over Ethernet support
- Support for link speed/width changing via SET_PORT

2.6 Changes in Rev 2.7.8100 From Rev 2.7.200

- Added Max and Current temperature measurements
- Support HCA class code configuration
- Support Link and Width change using the SET_PORT command
- Automatic detection of XAUI or XFI using QSFP cables
- Local invalidation enabled on Fibre Channel QPs
- Automatic detection of 1G/10G Ethernet

2.7 Changes in Rev 2.7.700 From Rev 2.7.200

- Added support for up to 7M Direct Attached 10GigE cables
- Added RDMA over Converged Ethernet (RoCE) support

2.8 Changes in Rev 2.7.200 From Rev 2.7.000

- IBTA spec 1.2.1 speed autonegotiation compliance for all standard Mellanox adapter cards
- For ConnectX-2 adapters only: Added the firmware configuration (.ini) file parameter a0_mode_in_cx2. When set (1), this parameter indicates to ConnectX-2 to mimic the ConnectX device's GPIO assignment (Step A0 mode)

3 Known Issues

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

Table 2 - Known Issues

Index	Issue	Description	Workaround	Scheduled Release (fix)
1.	RoCE performance in Lossy fabric	Massive packet drops in RoCE may cause low performance in lossy fabrics	N/A	N/A
2.	UAR Bar is too small for 64k-page machines	The small BAR causes driver loading to fail	Change the "log2_uar_bar_megabytes" .ini parameter under the [HCA] section as follows: log2_uar_bar_megabytes = 5	N/A
1.	GUID of ConnectX-2 Ethernet adapter cards	On ConnectX-2 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>). <code>Mlxburn/flint</code> return <code>0xffff</code> 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 <code>0xffff</code>).	N/A
3.	Change of memory BARs on a disabled system	Changing memory bars size / addresses between <code>SYS_DIS</code> and <code>SYS_EN</code> may cause the device to hang (ID: 24206)	N/A	N/A
4.	BAR resizing on an enabled system	Changing BAR sizes when a system is enabled may cause the device to hang (ID: 24208)	N/A	N/A
5.	Ethernet only: Must query all capabilities upon boot	If not all capabilities are queried upon boot, then the query command may fail. See the <code>QUERY_CAP</code> command in <i>ConnectX Programmer's Reference Manual</i>	Query all capabilities upon boot	N/A

Table 2 - Known Issues (Continued)

Index	Issue	Description	Workaround	Scheduled Release (fix)
6.	Disrupting QDR negotiation may lead to port rising as SDR	Disconnecting an IB cable (or closing the port) during QDR negotiation and then reconnecting (or reopening) may cause the adapter to bring up the port at SDR	Disconnect the cable (or close the port) again and then reconnect (reopen). To avoid this scenario, wait for QDR negotiation to finish prior to disconnecting the cable (or closing the port) and reconnecting (or reopening). The following are two possible methods to verify QDR negotiation is complete: a. The physical (green) LED is on. b. A query of LinkPhyState using a GetPortInfo MAD indicates LinkUp.	N/A
7.	MTNIC driver	MTNIC driver is not supported.	N/A	N/A
8.	InfiniBand Static rate	InfiniBand Static rate is not supported.	N/A	N/A
9.	Transceiver type reporting	Wrong reporting of transceiver type.	N/A	N/A
10.	Autosensing after cable re-connection to different port protocol	If you disconnect a cable from an IB/Eth port and reconnect it to an Eth/IB port (different protocol), the link may not rise	Restart the driver (openibd restart)	N/A
11.	InfiniBand PCIe.class_ID issue	Some boot applications may fail to boot on VPI cards due to the InfiniBand PCIe.class_ID issue	Change class_ID by applying the following setting to the .ini file: Under the [HCA] section, enter: hca_header_class_code = 0x028000 eth_class_en_ib = true	N/A
12.	Cables	Firmware v2.9.1200 does not support 0.5-1m 10GE SFP cables	N/A	N/A

4 History of Bug Fixes

Table 3 - Fixed Bugs List

	Issue	Description	Discovered in	Fixed in
1.	On rare occasions, a PCIe link failure of MLNX 10GbE single and dual port card occurred due to PCIe receiver detection failure	Fixed	2.9.1000	2.9.1200
2.	PCIe FunctionLevelReset occasionally failed when the driver was down	Fixed	2.9.1000	2.9.1200
3.	VMware driver failed to load when attached to Virtual Function one	Fixed	2.9.1000	2.9.1200
4.	Device may failed upon Wake on WOL packets	Fixed	2.9.1000	2.9.1200
5.	LLDP packets were silently dropped by the device	Fixed	2.9.1000	2.9.1200
6.	Possible false multicast miss in RoCE (ID 118145)	Fixed	2.9.1000	2.9.1200
7.	Possible false error completion upon FRWR (ID 115115)	Fixed	2.9.1000	2.9.1200
8.	Issues with the IPoIB checksum offload	Fixed	2.8.0600	2.9.1000
9.	PCIe Select de-emphasis was not functional in a specific environment (ID 85719)	Fixed	2.8.0600	2.9.1000
10.	The wrong event type was displayed in the thermal warning event for drivers that did not support the event	Fixed	2.8.0600	2.9.1000
11.	Wrong UAR mapping in pages larger than 4K: <ul style="list-style-type: none"> MSI-X and Blueflame did not work on UAR pages larger than 4K 	Fixed	2.8.0600	2.9.1000
12.	Mellanox adapters card MHGH29B, may enter a false thermal shutdown state upon driver unload	Fixed	2.8.0600	2.9.1000
13.	40GbE Link may not be established after driver restart	Fixed	2.8.0600	2.9.1000
14.	Congestion Notification Packet may cause HCA to stall under extreme conditions	Fixed	2.8.0600	2.9.1000
15.	Enforced correct IP identification field in transmitted IPoIB/EoIB packets	Fixed	2.8.0600	2.9.1000
16.	Advanced Error Reporting fixes (ID 100570)	Fixed	2.8.0600	2.9.1000
17.	Performance improvement for single port cards (ID 106144)	Fixed	2.8.0600	2.9.1000

Table 3 - Fixed Bugs List

18.	When in unsolicited mode a completion message is not sent on every completion	Fixed	2.8.0600	2.9.1000
19.	QUERY_IF_STAT - Fixed clear number of packets (ID 101245)	Fixed	2.8.0600	2.9.1000
20.	RDMA read message rate improvement (ID 97406)	Fixed	2.8.0600	2.9.1000
21.	If no EEPROM is present in the module/cable or if the identifier field in EEPROM is not set correctly, the link raises only 10G XFI (ID 102651)	Fixed	2.7.700	2.9.1000
22.	Tx rate limiting functionality	Fixed	2.8.0000	2.8.0600
23.	PCIe compatibility issue with a specific PCIe root complex	Fixed	2.8.0000	2.8.0600
24.	Possible failure to link up in 1GigE with several 1GigE devices	Fixed	2.7.000	2.8.0000
25.	PCI express compliancy: The ConnectX-2 PCIe transmitter may stop transmitting for a short period during software reset	Fixed	2.7.000	2.8.0000
26.	HCA may hang upon CLOSE_PORT command during data stress (ID 91657)	Fixed	2.7.000	2.8.0000
27.	Ethernet: Potential packet drops in Pause mode	Fixed	2.7.000	2.8.0000
28.	QUERY_DEV_CAP.log_max_bf_regs_per_page fixed from 8 to 3	Fixed	2.7.000	2.8.0000
29.	PCIe slow handling of configuration cycles may cause NMI	Fixed	2.7.000	2.7.9112
30.	Wrong completion ID in PCI Express INTx message	Fixed	2.7.000	2.7.8100
31.	PCI Express Selective De-emphasis functionality	Fixed	2.7.000	2.7.8100
32.	Potential delay in responding to configuration cycles upon data stress	Fixed	2.7.000	2.7.8100
33.	High Bit Error Rate on active cables	Fixed	2.7.000	2.7.8100
34.	PCI Express compliancy: The ConnectX-2 device may transmit before the PCI Express receiver detect is enabled	Fixed	2.7.000	2.7.8100
35.	Slow VPD performance might cause NMI	Fixed	2.7.000	2.7.8100
36.	Possible wrong link state during link bring up	Fixed	2.7.000	2.7.700
37.	High Bit Error Rate on active cables	Fixed	2.7.000	2.7.700
38.	Potential delay in responding to configuration cycles upon data stress	Fixed	2.7.000	2.7.700

Table 3 - Fixed Bugs List

39.	Ethernet link issues with back-to-back setting on 5m Twinax cables	Fixed	2.7.000	2.7.700
40.	RoCE issue: Wrong GUID in VPI	Fixed	2.7.700	2.8.0000
41.	RoCE issue: Missing GUIDs in Ethernet only cards	Fixed	2.7.700	2.8.0000