



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

Rev 2.9.1000

www.mellanox.com

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
Sunnyvale, CA 94085
U.S.A.
www.mellanox.com
Tel: (408) 970-3400
Fax: (408) 970-3403

Mellanox Technologies, Ltd.
PO Box 586 Hermon Building
Yokneam 20692
Israel
Tel: +972-4-909-7200
Fax: +972-4-959-3245

© Copyright 2011. Mellanox Technologies. All rights reserved.

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

FabricIT and SwitchX 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
Chapter 2 Revision Compatibility	5
Chapter 3 Changes and Major New Features	6
3.1 Changes in Rev 2.9.1000 From Rev 2.8.0600 -----	6
3.2 Changes in Rev 2.8.0600 From Rev 2.8.0000 -----	6
3.3 Changes in Rev 2.8.0000 From Rev 2.7.9112 -----	6
3.4 Changes in Rev 2.7.9112 From Rev 2.7.8100 -----	6
3.5 Changes in Rev 2.7.8100 From Rev 2.7.200 -----	7
3.6 Changes in Rev 2.7.700 From Rev 2.7.200 -----	7
3.7 Changes in Rev 2.7.200 From Rev 2.7.000 -----	7
Chapter 4 Bug Fixes	8
Chapter 5 Known Issues	10

1 Overview

These are the release notes for the ConnectX®-2 and ConnectX®-2 EN adapters firmware, fw-ConnectX2 Rev 2.9.1000. 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>.

Notes

1. Firmware Rev 2.9.1000 is compatible with MFT tools v2.6.0 and above. Using older versions of MFT tools will cause failure during image creation.
2. After burning new firmware to an adapter card, reboot the machine so that the new firmware can take effect. If you do not reboot, you will get an error in the RUN_FW command.

The document consists of the following sections:

- “Revision Compatibility” on page 5
- “Changes and Major New Features” on page 6
- “Bug Fixes” on page 8
- “Known Issues” on page 10
- “History of Fixed Issues” on page 11

2 Revision Compatibility

Firmware fw-ConnectX2 Rev 2.9.1000 complies with the following programmer's reference manual:

- *ConnectX Programmer's Reference Manual (PRM), Rev 0.39 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*.

3 Changes and Major New Features

3.1 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

3.2 Changes in Rev 2.8.0600 From Rev 2.8.0000

- Bug Fixes - see Section 4, “Bug Fixes,” on page 8

3.3 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

3.4 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`

3.5 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

3.6 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

3.7 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)

4 Bug Fixes

Table 1 - Firmware v2.9.1000 Bug Fixes

	Issue	Description	Discovered in	Fixed in
1.	Issues with the IPoIB checksum offload	Fixed	2.8.0600	2.9.1000
2.	PCIe Select de-emphasis was not functional in a specific environment (ID 85719)	Fixed	2.8.0600	2.9.1000
3.	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
4.	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
5.	Mellanox adapters card MHGH29B, may enter a false thermal shutdown state upon driver unload	Fixed	2.8.0600	2.9.1000
6.	40GbE Link may not be established after driver restart	Fixed	2.8.0600	2.9.1000
7.	Congestion Notification Packet may cause HCA to stall under extreme conditions	Fixed	2.8.0600	2.9.1000
8.	Enforced correct IP identification field in transmitted IPoIB/EoIB packets	Fixed	2.8.0600	2.9.1000
9.	Advanced Error Reporting fixes (ID 100570)	Fixed	2.8.0600	2.9.1000
10.	Performance improvement for single port cards (ID 106144)	Fixed	2.8.0600	2.9.1000
11.	When in unsolicited mode a completion message is not sent on every completion	Fixed	2.8.0600	2.9.1000
12.	QUERY_IF_STAT - Fixed clear number of packets (ID 101245)	Fixed	2.8.0600	2.9.1000
13.	RDMA read message rate improvement (ID 97406)	Fixed	2.8.0600	2.9.1000
14.	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
15.	Tx rate limiting functionality	Fixed	2.8.0000	2.8.0600
16.	PCIe compatibility issue with a specific PCIe root complex	Fixed	2.8.0000	2.8.0600
17.	Possible failure to link up in 1GigE with several 1GigE devices	Fixed	2.7.000	2.8.0000
18.	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

Table 1 - Firmware v2.9.1000 Bug Fixes

19.	HCA may hang upon CLOSE_PORT command during data stress (ID 91657)	Fixed	2.7.000	2.8.0000
20.	Ethernet: Potential packet drops in Pause mode	Fixed	2.7.000	2.8.0000
21.	QUERY_DEV_CAP.log_max_bf_regs_per_page fixed from 8 to 3	Fixed	2.7.000	2.8.0000
22.	PCIe slow handling of configuration cycles may cause NMI	Fixed	2.7.000	2.7.9112
23.	Wrong completion ID in PCI Express INTx message	Fixed	2.7.000	2.7.8100
24.	PCI Express Selective De-emphasis functionality	Fixed	2.7.000	2.7.8100
25.	Potential delay in responding to configuration cycles upon data stress	Fixed	2.7.000	2.7.8100
26.	High Bit Error Rate on active cables	Fixed	2.7.000	2.7.8100
27.	PCI Express compliancy: The ConnectX-2 device may transmit before the PCI Express receiver detect is enabled	Fixed	2.7.000	2.7.8100
28.	Slow VPD performance might cause NMI	Fixed	2.7.000	2.7.8100
29.	Possible wrong link state during link bring up	Fixed	2.7.000	2.7.700
30.	High Bit Error Rate on active cables	Fixed	2.7.000	2.7.700
31.	Potential delay in responding to configuration cycles upon data stress	Fixed	2.7.000	2.7.700
32.	Ethernet link issues with back-to-back setting on 5m Twi-nax cables	Fixed	2.7.000	2.7.700
33.	RoCE issue: Wrong GUID in VPI	Fixed	2.7.700	2.8.0000
34.	RoCE issue: Missing GUIDs in Ethernet only cards	Fixed	2.7.700	2.8.0000

5 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	Possible 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	N/A	Change the "log2_uar_bar_megabytes" .ini parameter under the [HCA] section as follows: log2_uar_bar_megabytes = 5	N/A
3.	Change of memory BARs on a disabled system	Changing memory bars size / addresses between SYS_DIS and SYS_EN may cause the device to hang (ID: 24206)	N/A	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	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 QUERY_CAP command in <i>ConnectX Programmer's Reference Manual</i>	N/A	Query all capabilities upon boot	N/A
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	N/A	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 is not supported				

Table 2 - Known Issues (Continued)

Index	Issue	Description	Current Implemented Workaround in FW	Possible Workaround	Scheduled Release (fix)
8.	InfiniBand Static rate is not supported				
9.	Wrong reporting of transceiver type				
10.	Autosensing after cable reconnection 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	N/A	Restart the driver (openibd restart)	N/A
11.	Some boot applications may fail to boot on VPI cards due to the infiniband PCIe.class_ID issue		N/A	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