



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

Rev 2.9.1314

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

© Copyright 2013. 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™, MetroX™, 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 Content

| | |
|---|-----------|
| Table of Content | 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 2.9.1314 From Rev 2.9.1200 | 6 |
| 2.2 Changes in Rev 2.9.1200 From Rev 2.9.1000 | 6 |
| 2.3 Changes in Rev 2.9.1000 From Rev 2.8.0600 | 6 |
| 2.4 Changes in Rev 2.8.0600 From Rev 2.8.0000 | 6 |
| 2.5 Changes in Rev 2.8.0000 From Rev 2.7.9112 | 6 |
| 2.6 Changes in Rev 2.7.9112 From Rev 2.7.8100 | 7 |
| 2.7 Changes in Rev 2.7.8100 From Rev 2.7.200 | 7 |
| 2.8 Changes in Rev 2.7.700 From Rev 2.7.200 | 7 |
| 2.9 Changes in Rev 2.7.200 From Rev 2.7.000 | 7 |
| Chapter 3 Known Issues | 8 |
| Chapter 4 History of Bug Fixes | 10 |

1 Overview

These are the release notes for the ConnectX®-2 and ConnectX®-2 EN adapters firmware, fw-ConnectX2 Rev 2.9.1314. 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 2.9.1314 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 |
|--------------------|---------------|--|
| 74Y3674 | IBM0F30000010 | ConnectX®-2 EN Dual-port SFP+ PCIe Gen2.0 x8 2.5GT/s RoHS R6 Adapter |
| 90Y3460 | IBM0FB0140018 | QDR mezz for next gen BladeCenter; Two ConnectX-2 devices |
| | IBM0FB0141018 | |
| 74Y2987 | IBM0FA0000009 | ConnectX®-2 VPI Dual-port 40Gb/s QSFP PCI-E 2.0 Adapter |
| 90Y3480 | IBM0FB0110023 | 10Gb Ethernet mezz for next gen BladeCenter; Two ConnectX-2 devices |
| | IBM0FB0111023 | |
| 60Y0928 | IBM08B0000009 | 2-Port 40Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter |
| | IBM08B0130009 | |

Table 1 - Supported PSIDs (Sheet 2 of 2)

| Device Part Number | PSID | Device Name |
|---------------------|---------------|---|
| 00E1600/ 00E1601 | IBM0F30001010 | Dual port 10GbE for pSeries low/high profile w/IBM optics |

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 2.9.1314 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 2.9.1314 From Rev 2.9.1200

- Bug Fixes - see [Section 4, “History of Bug Fixes”, on page 10](#)

2.2 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)

2.3 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

2.4 Changes in Rev 2.8.0600 From Rev 2.8.0000

- Bug Fixes - see [Section 4, “History of Bug Fixes”, on page 10](#)

2.5 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

2.6 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

2.7 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

2.8 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

2.9 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 | 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 | 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 |

Table 2 - Known Issues (Continued)

| Index | Issue | Description | Current Implemented Workaround in FW | Possible 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 | 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 | MTNIC driver is not supported. | N/A | N/A | N/A |
| 8. | InfiniBand Static rate | InfiniBand Static rate is not supported. | N/A | N/A | N/A |
| 9. | Transceiver type reporting | Wrong reporting of transceiver type. | N/A | 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 | N/A | 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 | 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 |
| 12. | Cables | Firmware v2.9.1200 does not support 0.5-1m 10GE SFP cables | N/A | N/A | N/A |

4 History of Bug Fixes

Table 3 - Fixed Bugs List

| | Issue | Discovered in | Fixed in |
|-----|--|---------------|----------|
| 1. | Under heavy rate of completions, when completion queue numbers are >0x8000, the firmware activated a flow that has as a side effect reads to address 0 | 2.9.1000 | 2.9.1314 |
| 2. | The <code>write_mcg</code> of InfiniBand MCG might fail after many insertions/removals of MCGs | 2.9.1000 | 2.9.1314 |
| 3. | ETH link fix - The <code>link_up</code> to driver was not reported while receiving <code>remote_fault</code> symbols | 2.9.1000 | 2.9.1314 |
| 4. | <code>General_Info</code> MAD fix - Vkey, PSID and firmware versions were reported incorrectly | 2.9.1000 | 2.9.1314 |
| 5. | VLAN and MAC table init fixes - Prevented table blocking by non-valid entries after many insert/remove operations | 2.9.1000 | 2.9.1314 |
| 6. | SR-IOV: <code>msix</code> vectors mask did not refreshed for VF after <code>cfg</code> write cycle | 2.9.1000 | 2.9.1314 |
| 7. | SR-IOV: Completion were not sent with the correct completor ID (according to the function index) | 2.9.1000 | 2.9.1314 |
| 8. | MalayaP and MalayaP-Net - Reduced the number of QDR Serdes sets to allow faster QDR link up time | 2.9.1000 | 2.9.1314 |
| 9. | Wrong reporting of RSS context in <code>QUERY_FW</code> of RSS qp | 2.9.1000 | 2.9.1314 |
| 10. | On rare occasions, a PCIe link failure of MLNX 10GbE single and dual port card occurred due to PCIe receiver detection failure | 2.9.1000 | 2.9.1200 |
| 11. | PCIe <code>FunctionLevelReset</code> occasionally failed when the driver was down | 2.9.1000 | 2.9.1200 |
| 12. | VMware driver failed to load when attached to Virtual Function one | 2.9.1000 | 2.9.1200 |
| 13. | Device may failed upon Wake on WOL packets | 2.9.1000 | 2.9.1200 |
| 14. | LLDP packets were silently dropped by the device | 2.9.1000 | 2.9.1200 |
| 15. | Possible false multicast miss in RoCE (ID 118145) | 2.9.1000 | 2.9.1200 |
| 16. | Possible false error completion upon FRWR (ID 115115) | 2.9.1000 | 2.9.1200 |
| 17. | Issues with the IPoIB checksum offload | 2.8.0600 | 2.9.1000 |
| 18. | PCIe Select de-emphasis was not functional in a specific environment (ID 85719) | 2.8.0600 | 2.9.1000 |

Table 3 - Fixed Bugs List

| | Issue | Discovered in | Fixed in |
|-----|---|----------------------|-----------------|
| 19. | The wrong event type was displayed in the thermal warning event for drivers that did not support the event | 2.8.0600 | 2.9.1000 |
| 20. | Wrong UAR mapping in pages larger than 4K: • MSI-X and Blueflame did not work on UAR pages larger than 4K | 2.8.0600 | 2.9.1000 |
| 21. | Mellanox adapters card MHGH29B, may enter a false thermal shut-down state upon driver unload | 2.8.0600 | 2.9.1000 |
| 22. | 40GbE Link may not be established after driver restart | 2.8.0600 | 2.9.1000 |
| 23. | Congestion Notification Packet may cause HCA to stall under extreme conditions | 2.8.0600 | 2.9.1000 |
| 24. | Enforced correct IP identification field in transmitted IPoIB/EoIB packets | 2.8.0600 | 2.9.1000 |
| 25. | Advanced Error Reporting fixes (ID 100570) | 2.8.0600 | 2.9.1000 |
| 26. | Performance improvement for single port cards (ID 106144) | 2.8.0600 | 2.9.1000 |
| 27. | When in unsolicited mode a completion message is not sent on every completion | 2.8.0600 | 2.9.1000 |
| 28. | QUERY_IF_STAT - Fixed clear number of packets (ID 101245) | 2.8.0600 | 2.9.1000 |
| 29. | RDMA read message rate improvement (ID 97406) | 2.8.0600 | 2.9.1000 |
| 30. | 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) | 2.7.700 | 2.9.1000 |
| 31. | Tx rate limiting functionality | 2.8.0000 | 2.8.0600 |
| 32. | PCIe compatibility issue with a specific PCIe root complex | 2.8.0000 | 2.8.0600 |
| 33. | Possible failure to link up in 1GigE with several 1GigE devices | 2.7.000 | 2.8.0000 |
| 34. | PCI express compliancy: The ConnectX-2 PCIe transmitter may stop transmitting for a short period during software reset | 2.7.000 | 2.8.0000 |
| 35. | HCA may hang upon CLOSE_PORT command during data stress (ID 91657) | 2.7.000 | 2.8.0000 |
| 36. | Ethernet: Potential packet drops in Pause mode | 2.7.000 | 2.8.0000 |
| 37. | QUERY_DEV_CAP.log_max_bf_regs_per_page fixed from 8 to 3 | 2.7.000 | 2.8.0000 |

Table 3 - Fixed Bugs List

| | Issue | Discovered in | Fixed in |
|-----|--|----------------------|-----------------|
| 38. | PCIe slow handling of configuration cycles may cause NMI | 2.7.000 | 2.7.9112 |
| 39. | Wrong completion ID in PCI Express INTx message | 2.7.000 | 2.7.8100 |
| 40. | PCI Express Selective De-emphasis functionality | 2.7.000 | 2.7.8100 |
| 41. | Potential delay in responding to configuration cycles upon data stress | 2.7.000 | 2.7.8100 |
| 42. | High Bit Error Rate on active cables | 2.7.000 | 2.7.8100 |
| 43. | PCI Express compliancy: The ConnectX-2 device may transmit before the PCI Express receiver detect is enabled | 2.7.000 | 2.7.8100 |
| 44. | Slow VPD performance might cause NMI | 2.7.000 | 2.7.8100 |
| 45. | Possible wrong link state during link bring up | 2.7.000 | 2.7.700 |
| 46. | High Bit Error Rate on active cables | 2.7.000 | 2.7.700 |
| 47. | Potential delay in responding to configuration cycles upon data stress | 2.7.000 | 2.7.700 |
| 48. | Ethernet link issues with back-to-back setting on 5m Twinax cables | 2.7.000 | 2.7.700 |
| 49. | RoCE issue: Wrong GID in VPI | 2.7.700 | 2.8.0000 |
| 50. | RoCE issue: Missing GUIDs in Ethernet only cards | 2.7.700 | 2.8.0000 |