



Mellanox ConnectX® Firmware (fw-ConnectX) Release Notes

Rev 2.8.0600

www.mellanox.com

NOTE:

THIS INFORMATION IS PROVIDED BY MELLANOX FOR INFORMATIONAL PURPOSES ONLY AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, 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 OUT OF THE USE OF THIS HARDWARE, 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, Inc. All Rights Reserved.
Mellanox®, BridgeX®, ConnectX®, InfiniBlast®, InfiniBridge®, InfiniHost®, InfiniRISC®, InfiniScale®, InfiniPCI®, PhyX®
and Virtual Protocol Interconnect® are registered trademarks of Mellanox Technologies, Ltd.
CORE-Direct and FabricIT are trademarks of Mellanox Technologies, Ltd.

All other marks and names mentioned herein may be trademarks of their respective companies.

Table of Contents

Table of Contents	3
Chapter 1 Overview	4
Chapter 2 Revision Compatibility	6
Chapter 3 Changes and Major New Features	7
3.1 Changes in Rev 2.8.0600 From Rev 2.8.0000	7
3.2 Changes in Rev 2.8.0000 From Rev 2.7.9112	7
3.3 Changes in Rev 2.7.9112 From Rev 2.7.8100	7
3.4 Changes in Rev 2.7.8100 From Rev 2.7.200	7
3.5 Changes in Rev 2.7.700 From Rev 2.7.200	7
3.6 Changes in Rev 2.7.200 From Rev 2.7.000	7
Chapter 4 Bug Fixes	8
Chapter 5 Known Issues	9
Chapter 6 History of Fixed Bugs	11

1 Overview

These are the release notes for the ConnectX® and ConnectX® EN adapters firmware, fw-ConnectX Rev 2.8.0600. 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 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 EN.

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>.

Notes

1. Firmware v2.8.000 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.
3. The following Mellanox adapters are not supported in MLNX_OFED 1.5.2-2.1.0:
 - MHGH29B-XSR - Dual 4X IB DDR Port, PCIe Gen2 x8 2.0 (5GT/s), Short Bracket, RoHS-R6 HCA Card, CX4 Connectors. (PSID: MT_OD10110008)
 - MHGH29B-XTR - Dual 4X IB DDR Port, PCIe Gen2 x8 2.0 (5GT/s), Tall Bracket, RoHS-R6 HCA Card, CX4 Connectors. (PSID: MT_OD20110008)

Table 1 - PCI Device ID

PCI Device ID (Decimal)	Device Part Number	Device Name	Supported Protocols
25408	MT25408A0-FCC-SI	ConnectX, Dual Port 10Gb/s InfiniBand / 10GigE Adapter IC with PCIe 2.0 x8 2.5GT/s Interface	InfiniBand, Ethernet, FCoE, VPI
25418	MT25408A0-FCC-DI	ConnectX, Dual Port 20Gb/s InfiniBand / 10GigE Adapter IC with PCIe 2.0 x8 2.5GT/s Interface	
26418	MT25408A0-FCC-GI	ConnectX, Dual Port 20Gb/s InfiniBand / 10GigE Adapter IC with PCIe 2.0 x8 5.0GT/s Interface	
26428	MT25408A0-FCC-QI	ConnectX, Dual Port 40Gb/s InfiniBand / 10GigE Adapter IC with PCIe 2.0 x8 5.0GT/s Interface	

Table 1 - PCI Device ID

PCI Device ID (Decimal)	Device Part Number	Device Name	Supported Protocols
25448	MT25448A0-FCC-SE	ConnectX EN, Dual Port 10GigE Adapter IC with PCIe 2.0 x8 2.5GT/s Interface	Ethernet
26448	MT26448A0-FCC-TE	ConnectX EN, Dual Port 10GigE Adapter IC with PCIe 2.0 x8 2.5GT/s Interface	
25458	MT25458A0-FCC-SE	ConnectX EN, Dual Port 10BASE-T Adapter IC with PCIe 2.0 x8 2.5GT/s Interface	
26458	MT25458A0-FCC-TE	ConnectX EN, Dual Port 10BASE-T Adapter IC with PCIe 2.0 x8 5.0GT/s Interface	

The document consists of the following sections:

- “Revision Compatibility” on page 6
- “Changes and Major New Features” on page 7
- “Bug Fixes” on page 8
- “Known Issues” on page 9
- “History of Fixed Issues” on page 11

2 Revision Compatibility

Firmware fw-ConnectX Rev 2.8.0600 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 2.8.0600 From Rev 2.8.0000

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

3.2 Changes in Rev 2.8.0000 From Rev 2.7.9112

Added support for the following features:

- Signal integrity improvements in IBTA 1.2.1 mode (Attenuation Based Algorithm)
- Eye opener MAD
- Cable info MAD
- CQ.OI to the MODIFY_CQ command to enable CORE-Direct
- SET_NODE command
- PCIe device Serial Number capability
- LoosyVI MAD

3.3 Changes in Rev 2.7.9112 From Rev 2.7.8100

- Support for link speed/width changing via SET_PORT

3.4 Changes in Rev 2.7.8100 From Rev 2.7.200

- Support HCA class code configuration
- Support Link and Width change using the SET_PORT command
- Local invalidation enabled on Fibre Channel QPs

3.5 Changes in Rev 2.7.700 From Rev 2.7.200

- Added RDMA over Converged Ethernet (RoCE) support

3.6 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

4 Bug Fixes

Table 2 - Firmware v2.8.0600 Bug Fixes

	Issue	Description	Discovered in	Fixed in
1.	PCIe compatibility issue with a specific PCIe root complex	Fixed	2.8.0000	2.8.0600
2.	HCA may hang upon CLOSE_PORT command during data stress (ID 91657)	Fixed	2.7.000	2.8.0000
3.	Ethernet: Potential packet drops in Pause mode	Fixed	2.7.000	2.8.0000
4.	QUERY_DEV_CAP.log_max_bf_regs_per_page fixed from 8 to 3	Fixed	2.7.000	2.8.0000
5.	PCIe slow handling of configuration cycles may cause NMI	Fixed	2.7.000	2.7.9112
6.	Potential delay in responding to configuration cycles upon data stress	Fixed	2.7.000	2.7.8100
7.	High Bit Error Rate on active cables	Fixed	2.7.000	2.7.8100
8.	PCI Express compliancy: The ConnectX device may transmit before the PCI Express receiver detect is enabled	Fixed	2.7.000	2.7.8100
9.	Slow VPD performance might cause NMI	Fixed	2.7.000	2.7.8100
10.	Possible wrong link state during link bring up	Fixed	2.7.000	2.7.700
11.	High Bit Error Rate on active cables	Fixed	2.7.000	2.7.700
12.	Potential delay in responding to configuration cycles upon data stress	Fixed	2.7.000	2.7.700
13.	RoCE issue: Wrong GID in VPI	Fixed	2.7.700	2.8.0000
14.	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 3 - 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				
8.	InfiniBand Static rate is not supported				

Table 3 - Known Issues (Continued)

Index	Issue	Description	Current Implemented Workaround in FW	Possible Workaround	Scheduled Release (fix)
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

6 History of Fixed Bugs

Table 4 - History of Bug Fixes (Sheet 1 of 5)

	Issue	Description	Discovered in	Fixed in
1.	InfiniBand auto-negotiation issues	Fixed	2.6.900	2.7.000
2.	MOD_STAT_CFG.port_en command does not apply to Ethernet	Fixed	2.6.000	2.6.900
3.	Mellanox auto-negotiation is not optimized for environments with significant crosstalk	Fixed	2.6.000	2.6.900
4.	CQ/EQ doorbell is lost in Configuration cycles	Fixed	2.6.000	2.6.900
5.	Congestion issues	Fixed	2.6.000	2.6.900
6.	SW2HW_MPT command may fail when MPT is in FREE state (used for FRWR)	Fixed	2.6.000	2.6.900
7.	Adapter may get stuck upon multicast traffic stress	Fixed	2.6.000	2.6.900
8.	IB Spec Release 1.2 link data-rate issues	Fixed	2.6.000	2.6.900
9.	Adapter may get stuck if QPC extended auxiliary table base address is not aligned to 1MB	Fixed	2.6.000	2.6.900
10.	Port mod_stat_cfg.disable does not take effect if the device is configured as an Ethernet device	Fixed	2.6.000	2.6.300
11.	HCA might hang when entering PCIe L1 mode	Fixed	2.6.000	2.6.200
12.	Link may come up as SDR instead of DDR on some systems	Fixed	2.5.000	2.6.000
13.	Possible live lock in QP upon retransmission stress	Fixed – (ID: 49870, 52066)	2.5.000	2.6.000
14.	Wrong link state reported during link speed negotiation	Fixed – (ID: 49951)	2.5.000	2.6.000
15.	CQs may be generated after CQ overrun error	Fixed – (ID: 49982)	2.5.000	2.6.000
16.	Slow handling of configuration cycles	Fixed – (ID: 49807)	2.5.000	2.6.000
17.	Wrong fields in CQE-w-Error on XRC QP	Fixed – (ID: 49742)	2.5.000	2.6.000
18.	Wrong handling of Remote Invalidate Error	Fixed	2.5.000	2.6.000
19.	Multiple RNR Nack may cause slowdown	Fixed – (ID: 49559)	2.5.000	2.6.000

Table 4 - History of Bug Fixes (Sheet 2 of 5)

	Issue	Description	Discovered in	Fixed in
20.	QUERY_DEV_CAP.apm bit was fixed at 0 even though APM was active	Fixed – (ID: 49548)	2.5.000	2.6.000
21.	PCIe physical errors upon entering L1 state	Fixed – (ID: 52025)	2.5.000	2.6.000
22.	PCI_CFG.interrupt_disable has no impact	Fixed – (ID: 53350)	2.5.000	2.6.000
23.	Non-default setting of VLCap via .ini does not take effect	Fixed	2.5.000	2.6.000
24.	SET_PORT may lead to non-optimal RX buffer reallocation if opvl was less than vlcap	Fixed	2.5.000	2.6.000
25.	Modified PLL parameter settings in .ini	Some PLL parameter settings were changed to allow a longer period for PLL stabilization	2.5.000	2.6.000
26.	Adapter may generate PCIe transactions with wrong function ID	Fixed	2.5.000	2.6.000
27.	Adapter may generate PCIe ERR_NON_FATAL in case of an unsupported request	Fixed	2.5.000	2.6.000
28.	<u>Ethernet only</u> : Different VLAN priorities for WQE and QP may cause wrong SchedQ allocation	Fixed	2.5.000	2.6.000
29.	Wrong VLAN Priority in PPP mode	The ConnectX device may send pause frames for the wrong priority in PPP mode (ID: 49533)	2.5.000	2.5.900
30.	Wrong aliasing in address steering mode	This issues applies only to the VMware® Infrastructure 3 v3.5 operating system in netq mode	2.5.000	2.5.900
31.	Possible packet dropping though the pause policy is set	Fixed	2.5.000	2.5.900
32.	Wrong PCI Class Code for Ethernet Network Controller	This issue applies to Ethernet devices only. Fixed.	2.5.000	2.5.900
33.	Bringing up/down of an adapter port may lead to PHY errors on the second adapter port	Applies only to dual-port adapter cards. Fixed (ID: 51356)	2.5.000	2.5.900
34.	PPP does not work for an MTU of 9600	Fixed	2.5.000	2.5.900
35.	PCI Express compliancy issues	<ul style="list-style-type: none"> Fixed L1 and L0s power states compliancy issues Fixed PCIE-CV test completion_timeout failure Fixed interoperability issue with all available PCIe Gen. 2.0 servers (Ref. ID: 43852)	2.3.000	2.5.000
36.	INTA may be lost under stress	Fixed. (Ref. ID: 44473)	2.3.000	2.5.000
37.	Modifying SRQ number using RTS2RTS	Modifying SRQ number using RTS2RTS does not guarantee that no new CQEs will be generated using the old SRQ number. Fixed.	2.3.000	2.5.000

Table 4 - History of Bug Fixes (Sheet 3 of 5)

	Issue	Description	Discovered in	Fixed in
38.	QP may get stuck upon Responder Gather Error	Fixed.	2.3.000	2.5.000
39.	Wrong handling of SL mismatch between WQE and MLX QP	An SL mismatch between WQE and MLX QP may cause the QP to get stuck. Fixed.	2.3.000	2.5.000
40.	UC QP CQE with Error causes corruption	Fixed.	2.3.000	2.5.000
41.	Query_CQ/Query_EQ commands may return the old consumer_index	Fixed.	2.3.000	2.5.000
42.	CQ error may cause corruption	A CQ error such as an overrun may cause CQ corruption, leading to a wrong CQ number in the CQ error event or to an internal FW error. Fixed.	2.3.000	2.5.000
43.	Possible FW internal error on a very noisy link	Fixed. (Ref. ID:41526)	2.3.000	2.5.000
44.	QueryDebugMSG command returns wrong status	Fixed. (Ref. ID: 44744)	2.3.000	2.5.000
45.	Dropping a ReadResponse packet may lead to 'retry exceeded'	Fixed.	2.3.000	2.5.000
46.	CQ moderation parameters are wrongly configured	Fixed. (Ref. ID: 45570)	2.3.000	2.5.000
47.	False generation of CQE with error (vendor code 0x6f) upon large stress	Fixed. (Ref. ID: 45317)	2.3.000	2.5.000
48.	Bandwidth degradation if SetPort command is not called	Fixed.	2.3.000	2.5.000
49.	SQERR2RTS command followed by an error causes QP to be unfunctional	Fixed. (Ref. ID: 45828 45848)	2.3.000	2.5.000
50.	QUERY_FW fails after RUN_FW	The command QUERY_FW fails after running the RUN_FW command	2.2.000	2.3.000
51.	HCA stall	The HCA might stall in any of the following scenarios: <ul style="list-style-type: none"> • If running the command SET_DEBUG_MESSAGE (ID:42128) • Under large stress (ID: 43385, 43378) • Upon closing a large number of QPs (ID: 43697) • If the WQE SL is different than the QP Context SL in a UD QP (ID: 41423) • Upon multiple retransmissions 	2.2.000	2.3.000
52.	QUERY_QP errors	Wrong QUERY_QP command in the following cases: <ul style="list-style-type: none"> • Returns wrong values (ID: 42078, 40707) • Enters the error state erroneously (ID: 43110) 	2.2.000	2.3.000
53.	IB & PCI Express links quality	General improvements	2.2.000	2.3.000
54.	Incomplete support for PCI Express 2.0 configuration header	Fixed	2.2.000	2.3.000
55.	Wrong trap generation rate	The HCA might exceed the maximum trap generation rate upon processing different trap types	2.2.000	2.3.000

Table 4 - History of Bug Fixes (Sheet 4 of 5)

	Issue	Description	Discovered in	Fixed in
56.	Client Reregister event not generated	The HCA might fail to generate a Client Reregister event under large stress. (ID: 42232)	2.2.000	2.3.000
57.	Possible ICM corruption	Possible ICM (Interconnect Context Memory) corruption upon large stress (ID: 42529)	2.2.000	2.3.000
58.	Performance	HCA performance improvements for the following cases: <ul style="list-style-type: none"> • Upon receiving multiple ACK packets • Upon multiple QPs in error state (ID:43377) • Upon multiple RNR NACKs 	2.2.000	2.3.000
59.	Wrong wqe_index in Receive CQE with Error	This can occur when running stress IPoIB CM tests. (ID: 43076)	2.2.000	2.3.000
60.	Possible multicast corruption	Fixed (ID: 43301)	2.2.000	2.3.000
61.	Wrong limit on number of supported EQ UARs	The HCA now supports the requested number of EQ UARs specified in INIT_HCA	2.2.000	2.3.000
62.	SchedQueue corruption	Fixed (ID: 43289)	2.2.000	2.3.000
63.	Wrong SL2VL mapping upon set_sl2vl	Fixed	2.2.000	2.3.000
64.	False MAD packet drops	The HCA might drop MAD packet erroneously under large stress	2.2.000	2.3.000
65.	PCI Express 2.0 x1 link fails to rise	Fixed	2.2.000	2.3.000
66.	Command timeouts	The HCA times out commands while closing multiple QPs	2.2.000	2.3.000
67.	False internal error generation	Fixed	2.2.000	2.3.000
68.	Transport timeouts	Multiple RNR NACKs may lead to transport timeouts (ID: 44160)	2.2.000	2.3.000
69.	Opcode/Input Modifier verification	Command Opcode/Input Modifier values are now checked for correctness. If a wrong value is provided, the command status indicates the error.	2.2.000	2.3.000
70.	Wrong sl and/or port number returned	The QUERY_QP command may return a wrong sl value and/or a wrong port number (ID: 40707)	2.1.000	2.2.000
71.	HCA stall	The HCA might stall upon stress involving RNR Nacks and RDMA reads (ID: 41918)	2.1.000	2.2.000
72.	QP corruption	QP corruption may occur following a CQ_overrun	2.1.000	2.2.000
73.	Sched Queue corruption	Sched Queue corruption may occur upon multiple re-transmissions	2.1.000	2.2.000
74.	False SRQ WQE limit event	A false SRQ WQE limit event is generated due to a race condition	2.0.164	2.1.000
75.	Wrong Dt value returned	The QUERY_FW command may return a wrong Dt value	2.0.164	2.1.000
76.	HCA hangs	The device hangs in one of the following cases: <ul style="list-style-type: none"> • upon retry – due to local_ack_timeout • upon retry – due to RNR Nack • upon ringing a CQ doorbell for an invalid QP • upon stress conditions (IDs: 41543,732/6,755,778) 	2.0.164	2.1.000
77.	High ACK latency	Delays in ACK may cause multiple local ACK timeouts	2.0.164	2.1.000

Table 4 - History of Bug Fixes (Sheet 5 of 5)

	Issue	Description	Discovered in	Fixed in
78.	HCA performance	HCA performance may be impacted in the following conditions: <ul style="list-style-type: none">• QPs in error state• Slow QP context handling	2.0.164	2.1.000
79.	IB link stability issues		2.0.164	2.1.000
80.	High QP closing duration	Closing QPs with outstanding posted WQs may take a long time due to slow CQE with error generation	2.0.164	2.1.000