



Connect. Accelerate. Outperform.™

Mellanox Connect-IB® Firmware (fw-ConnectIB) Release Notes

Rev 10.10.4050

Last Updated: February 2015

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.
 Hakidma 26
 Ofer Industrial Park
 Yokneam 2069200
 Israel
www.mellanox.com
 Tel: +972 (0)74 723 7200
 Fax: +972 (0)4 959 3245

© Copyright 2015. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, ConnectX®, Connect-IB®, CoolBox®, CORE-Direct®, GPUDirect®, InfiniBridge®, InfiniHost®, InfiniScale®, Kotura®, Kotura logo, MetroX®, MLNX-OS®, PhyX®, ScalableHPC®, SwitchX®, TestX®, UFM®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

ExtendX™, FabricIT™, HPC-X™, Mellanox CloudX™, Mellanox Open Ethernet™, Mellanox PeerDirect™, Mellanox Virtual Modular Switch™, MetroDX™, Switch-IB™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

List of Tables	4
Chapter 1 Overview	6
1.1 Supported Devices	6
1.2 Supported Cables and Modules	6
1.2.1 Tested Cables and Modules	7
1.3 Tested Switches	8
1.4 Tools and Driver Software	9
1.5 Revision Compatibility	9
Chapter 2 Supported Features	10
2.1 Changes in Rev 10.10.4050 From Rev 10.10.4020	10
2.2 Changes in Rev 10.10.4020 From Rev 10.10.3000	10
2.3 Changes in Rev 10.10.3000 From Rev 10.10.2000	10
2.4 Changes in Rev 10.10.2000 From Rev 10.10.1000	10
2.5 New Features in Rev 10.10.1000	11
Chapter 3 Known Issues	12
Chapter 4 Bug Fixes History	13
Chapter 5 Unsupported Features and Commands	15
5.1 Unsupported Features	15
5.2 Unsupported Commands	15
Appendix A Flexboot Changes and New Features	16

List of Tables

Table 1: Release Update History	5
Table 2: Supported PSIDs	6
Table 3: Tested Cables and Modules	7
Table 4: Tested Switches	8
Table 5: Known Issues	12
Table 6: Fixed Bugs List	13
Table 7: FlexBoot Changes and New Feature	16

Release Update History

Table 1 - Release Update History

Release	Date	Description
10.10.4050	February 2015	Added bug fix # 2 to Section 4, "Bug Fixes History" , on page 13
	November 2014	Initial version

1 Overview

These are the release notes for the Connect-IB® adapters firmware, fw-ConnectIB Rev 10.10.4050. This firmware supports the following protocols:

- InfiniBand – SDR, DDR, QDR, FDR10, FDR
- PCI Express 3.0, supporting backwards compatibility for v2.0 and v1.1

1.1 Supported Devices

This firmware supports the devices and protocols listed in Table 2. 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 2 - Supported PSIDs

Device Part Number	PSID	Device Name	Supported Protocols
MCB191A-FCAT	MT_1230110019	Connect-IB® Host Channel Adapter, single-port QSFP, FDR 56Gb/s, PCIe3.0 x8, tall bracket, RoHS R6	InfiniBand
MCB192A-FCAT	MT_1240110019	Connect-IB® Host Channel Adapter, dual-port QSFP, FDR 56Gb/s, PCIe3.0 x8, tall bracket, RoHS R6	InfiniBand
MCB193A-FCAT	MT_1220110019	Connect-IB® Host Channel Adapter, single-port QSFP, FDR 56Gb/s, PCIe3.0 x16, tall bracket, RoHS R6	InfiniBand
MCB193A-FBAT	MT_1220110030	Connect-IB® Host Channel Adapter; single-port QSFP; FDR 56Gb/s; PCIe2.0 x16; RoHS R6	InfiniBand
MCB194A-FCAT	MT_1210110019	Connect-IB® Host Channel Adapter, dual-port QSFP, FDR 56Gb/s, PCIe3.0 x16, tall bracket, RoHS R6	InfiniBand

1.2 Supported Cables and Modules

Please refer to the LinkX™ Cables and Transceivers web page

(www.mellanox.com -> Products -> Cables and Transceivers) for the list of supported cables.

1.2.1 Tested Cables and Modules

Table 3 - Tested Cables and Modules

Speed	OPN #	Description
DDR	MC1204128-005	MT PASSIVE COPPER CABLE 4X CX4 TO QSFP 20GB/S 28AWG 5M
DDR	MC1204130-001	MT PASSIVE COPPER CABLE 4X CX4 TO QSFP 20GB/S 30AWG 1M
DDR	MC1204130-003	MT PASSIVE COPPER CABLE 4X CX4 TO QSFP 20GB/S 30AWG 3M
SDR	MC1104130-001	MT PASSIVE COPPER CABLE 4X CX4 20GB/S 30AWG 1M
SDR	MC1104130-003	MT PASSIVE COPPER CABLE 4X CX4 20GB/S 30AWG 3M
QDR	MC2206310-300	MT ACTIVE FIBER CABLE 4X QSFP 40GB/S 200M
QDR	MC2206125-007	MT PASSIVE COPPER CABLE 4X QSFP 40GB/S 25AWG 7M
QDR	MC2206126-006	MT PASSIVE COPPER CABLE 4X QSFP 40GB/S 26AWG 6M
QDR	MC1204128-005	MT PASSIVE COPPER CABLE 4X CX4 TO QSFP 20GB/S 28AWG 5M
QDR/FDR10/ 40GE	MC2206310-030-T	MT ACTIVE FIBER CABLE 4X QSFP 40GB/S 30M
QDR / FDR10/ 40G	MC2206310-100-T	MT ACTIVE FIBER CABLE 4X QSFP 40GB/S 100M
QDR/FDR10/ 40GE	MC2206310-030-F	MT ACTIVE FIBER CABLE 4X QSFP 40GB/S 30M
QDR / FDR10/ 40G	MC2206310-100-F	MT ACTIVE FIBER CABLE 4X QSFP 40GB/S 100M
QDR / FDR10/ 40GE	MFS4R12CB-100	MT ACTIVE FIBER CABLE IB QDR/FDR10 40GB/S QSFP 100M
QDR / FDR10	MC2206128-005	MT PASSIVE COPPER CABLE 4X QSFP 40GB/S 28AWG 5M
QDR / FDR10	MC2206130-003	MT PASSIVE COPPER CABLE 4X QSFP 40GB/S 30AWG 3M
QDR / FDR10	MC2206130-00A	MT PASSIVE COPPER CABLE 4X QSFP 40GB/S 30AWG 0.5M
QDR / FDR10/ 40GE	MC2210411-SR4 -T	Mellanox Optical Module 40Gb/S QSFP LC-LC 850NM LR4. Up To 100M
QDR / FDR10/ 40GE	MC2210411-SR4 -F	Mellanox Optical Module 40Gb/S QSFP LC-LC 850NM LR4. Up To 100M
QDR	MFM4R12C-QDR	40GB/S INFINIBAND QSFP OPTICAL MODULE IN SINGLE BOX
FDR; 56G VPI	MC2207130-00A	MT PASSIVE COPPER CABLE 4X QSFP 56GB/S 30AWG 0.5M

Table 3 - Tested Cables and Modules

Speed	OPN #	Description
FDR; 56G VPI	MC2207130-002	MT PASSIVE COPPER CABLE 4X QSFP 56GB/S 30AWG 2M
FDR; 56G VPI	MC2207128-003	MT PASSIVE COPPER CABLE 4X QSFP 56GB/S 28AWG 3M
FDR; 56G VPI	MC2207126-004*	MT PASSIVE COPPER CABLE 4X QSFP 56GB/S 28AWG 4M
FDR; 56 VPI	MC2207312-003	MT ACTIVE FIBER CABLE VPI IB FDR (56GB/S) AND ETH 40GBE QSFP 3M
FDR; 56G VPI	MC2207310-030-E	MT ACTIVE FIBER CABLE 4X QSFP 56GB/S 30M
FDR; 56G VPI	MC2207312-100	MT ACTIVE FIBER CABLE 4X QSFP 56GB/S 100M
FDR; 56G VPI	MC2207310-100-E	MT ACTIVE FIBER CABLE 4X QSFP 56GB/S 100M
FDR; 56G VPI	MC2207310-030-F	MT ACTIVE FIBER CABLE 4X QSFP 56GB/S 30M
FDR; 56G VPI	MC2207310-100-F	MT ACTIVE FIBER CABLE 4X QSFP 56GB/S 100M
IB FDR	MC220731V-003	Mellanox Active Fiber IB QSFP 3M
IB FDR	MC220731V-100	Mellanox Active Fiber IB QSFP 100M
PCC 100G ECLIPSE	MCP1600-E001-V	MELLANOX PASSIVE COPPER CABLE VPI 100GB/S QSFP LSZH 1M
PCC 100G ECLIPSE	MCP1600-E002-V	MELLANOX PASSIVE COPPER CABLE VPI 100GB/S QSFP LSZH 2M
PCC 100G ECLIPSE	MCP1600-E003-V	MELLANOX PASSIVE COPPER CABLE VPI 100GB/S QSFP LSZH 3M

1.3 Tested Switches

Table 4 - Tested Switches

Speed	OPN #/Name	Description
IB SDR	F-X430060	24-port SDR-Switch
IB QDR	QLogic 12300	36-Port 40Gb QDR Infiniband Switch, Management Module, Dual Power
IB QDR	MIS5025Q-1SFC	InfiniScale® IV QDR InfiniBand Switch, 36 QSFP ports, 1 Power Supply, Unmanaged, PSU side to connector side airflow, Standard depth, Rail Kit and RoHS5
IB QDR	QDR-Switch 4036	InfiniScale® IV QDR Mellanox Grid Director 4036 36-Port QDR InfiniBand Switch - Part ID: VLT-30011
IB FDR	MSX6036F-1BFR	SwitchX™ based FDR InfiniBand Switch, 36 QSFP ports, 1 Power Supply, Short depth, Managed, PSU side to Connector side airflow, Rail Kit and RoHS6

Table 4 - Tested Switches

Speed	OPN # /Name	Description
IB FDR	MSX6018F-1SFR	SwitchX® based FDR InfiniBand Switch, 18 QSFP ports, 1 Power Supply, Standard depth, Managed, PSU side to Connector side air-flow, Rail Kit and RoHS6
IB FDR10	MSX6025T-1SFR	SwitchX™ based FDR10 Infiniband Switch, 36 QSFP ports, 1 Power Supply, Standard depth, Unmanaged, PSU side to Connector side air-flow, Rail Kit and RoHS6
IB FDR	648311-B21	HP BLc 4X FDR IB Switch
IB QDR	489184-B21	HP BLc 4X QDR IB Switch

1.4 Tools and Driver Software

Firmware Rev 10.10.4050 is tested with the following tools and driver software:

- Driver versions:
 - MLNX_OFED 2.1-1.0.0 and higher (For new features, required driver version is 2.3-1.0.0)
- MFT for Linux version: 3.7.0 and higher
- PXE version 3.4.306 and higher
- UEFI version 10.5.26 and higher

1.5 Revision Compatibility

Firmware fw-ConnectIB Rev 10.10.4050 complies with the following programmer's reference manual:

- *ConnectIB Programmer's Reference Manual (PRM), Rev 2.05 or later*, which has Command Interface Revision 0x5. 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 in Rev 10.10.4050 From Rev 10.10.4020

- Suspend to RAM (S3) support
- Diagnostic counters vendor-specific MAD support, as defined by VS-MAD spec version 1.2
- Bug Fixes (see [Section 4, “Bug Fixes History,”](#) on page 13)

2.2 Changes in Rev 10.10.4020 From Rev 10.10.3000

- On Demand Paging (ODP) - Memory can now be used without pinning memory beforehand. Supported transports are UD, and RC
ODP support is GA for RC RDMA-write, RC send, and UD send. Beta level for RDMA-read/atomics
- Enhanced Atomic Operations to include all PRM atomic operations of 32 bytes or below
- Dynamically Connected (DC) transport improvements

2.3 Changes in Rev 10.10.3000 From Rev 10.10.2000

- Dynamically Connected (DC) transport (at GA level)
- Enabled Atomic Operations. For further information, please refer to the PRM section *“Atomic Capabilities”*
- Added sniffer QP support (Note: Sniffer QP is currently not available in MLNX_OFED v2.2-1.0.0 or the MFT tools package)
- Increased the maximum number of InfiniBand partitions to 0x1000

2.4 Changes in Rev 10.10.2000 From Rev 10.10.1000

- Dynamically Connected (DC) transport (at Beta level)
- CORE-Direct®
 - Provides Collective Off-loading in HCA
 - Frees CPU to perform computation in parallel with collective operations
- T10 DIF Data Integrity Signature off-loading¹
- Removed software limitations that were required for the use of Mellanox-certified FDR InfiniBand cables with Mellanox FDR InfiniBand adapters and switches. Please refer to "Memo: FDR 56Gb/s InfiniBand Cables" that was released on Dec/2013.
Mellanox will offer an EXTENDED diagnostics support plan which will be available for mixed environments only and that will help identify issues they may encounter with the FDR installations.
- User Memory Registration (UMR)
- InfiniBand Automatic Path Migration

1. For further information, please refer to MLNX_OFED Release Notes.

2.5 New Features in Rev 10.10.1000

- Initial Release of Connect-IB™
- Port speed up to FDR
- PCI Express 3.0 x16, with backwards compatibility with v2.0 and v1.1
- Dynamically Connected (DC) transport at Alpha level

3 Known Issues

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

Table 5 - Known Issues

Index	Issue	Description	Workaround	Scheduled Release (fix)
1.	Upgrading/Downgrading	Flashing the firmware requires server reboot. Firmware cannot be flashed twice without server reboot after first flashing	Reboot the server after firmware flashing	Future Release
2.	MADs	Setting the port to 'sleep' state is not supported.	N/A	Future Release
3.	PCIe	Link width x1 might get Replay Timer Timeout, on speed change.	N/A	Future Release
4.		L1 power state enter requests are ignored by the device.	N/A	Future Release
5.	Miscellaneous	[For customers developing custom low level drivers] The device does not recover if the requested number of pages are not supplied during device initialization.	N/A	Future Release
6.	Quality of Service	On rare occasions, SL to VL modification with functioning QPs results in traffic hangs.	Configure SL2VL during SM initialization	Future Release

4 Bug Fixes History

Table 6 lists the bugs fixed in this release.

Table 6 - Fixed Bugs List

#	Category	Description	Discovered in Release	Fixed in Release
1.	DC transport	Reduced DCT destroy firmware handling time	10.10.4020	10.10.4050
2.	Ports	Fixed link flapping issue which occurred when LLR was active.	10.10.3000	10.10.4050
3.	PCIe class code of the device	Deprecated code 0x0c0600 was changed to 0x020700 (InfiniBand network adapter)	10.10.3000	10.10.4020
4.	Atomic support	Atomic response endianness is always a big endian	10.10.1000	10.10.3000
5.	Miscellaneous	[Documentation fix in PRM v2.01, no changes to the firmware code.] Port asynchronous events documentation are different from the PRM. All port events have a type value of 0x9. The following subtype values are used for the following events: <ul style="list-style-type: none"> • link down=0x1 • link up=0x4 • link initialized=0x5 • lid change=0x6 • PKEY change=0x7 • GUID change=0x8 • client reregister=0x9 	10.10.1000	10.10.3000
6.	InfiniBand Transport	Alternate Path Migration (APM) triggers only a single affiliated asynchronous error event in the case of a path migration failure.	10.10.2000	10.10.3000
7.		Using a <code>min_rnr_nak</code> value of 0x5 will cause failures when creating reliable connection (RC) QPs.	10.10.2000	10.10.3000
8.	DC	On rare occasions DC Initiator completions might be lost.	10.10.2000	10.10.3000
9.	Data Integrity Signature	The following signature rules are not supported (Numbering based on "signature rules table" in PRM): <ul style="list-style-type: none"> • Rule #12: T10 DIF • Rule #13: T10 DIF CS • Rule #14 T10 DIF CS 	10.10.2000	10.10.3000
10.	Quality of Service	VL arbitration configuration does not ensure minimum bandwidth for VL as configured.	10.10.1000	10.10.2000
11.	False alarm report	On very rare occasions, a false firmware "hanged" report is printed in the dmesg.	10.10.1000	10.10.2000
12.	CQ buffer resize	CQ buffer resize not supported	10.10.1000	10.10.2000

Table 6 - Fixed Bugs List

#	Category	Description	Discovered in Release	Fixed in Release
13.	Ports	When connecting to InfiniScale family switches and non-Mellanox InfiniBand switches DDR and QDR speeds may show line errors and in some cases might downgrade to SDR speed.	10.10.1000	10.10.2000

5 Unsupported Features and Commands

5.1 Unsupported Features

The following advanced feature as described in Connect-IB PRM Rev 2.01 are unsupported in the current firmware version:

- Service types not supported:
 - SyncUMR
 - Mellanox transport
 - PTP
 - RAW IPv6
 - PTP (ieee 1588)
- SR-IOV
- Connect-IB® currently supports only a single physical function model
- INT-A not supported for EQs only MSI-X
- PCI VPD write flow (RO flow supported)
- Streaming receive queue (STRQ) and collapsed CQ
- Precise clock synchronization over the network (IEEE 1588)
- Data integrity validation of control structures
- NC-SI interface is not enabled in Connect-IB® firmware
- PCIe Function Level Reset (FLR)

5.2 Unsupported Commands

- QUERY_MAD_DEMUX
- SET_MAD_DEMUX
- PAGE_FAULT_RESUME
- ACTIVATE_TRACER
- DEACTIVATE_TRACER

Appendix A: Flexboot Changes and New Features

For further information, please refer to FlexBoot Release Notes
(www.mellanox.com > Software > InfiniBand/VPI Drivers > FlexBoot).

Table 7 - FlexBoot Changes and New Feature

Version	Description
Rev 3.4.306	<ul style="list-style-type: none"> • Added validation script for the released ROMs • Added the option to always keep SAN hook to enable WIN install on iSCSI target • Added compilation flag around the flash readout. • Added URI Boot retry. Default retries = 0. • Added Unmap MPT command in teardown. • Added support for HII iSCSI configuration. • Added 64-bit PCI BAR support (Large bar). • Added the option added for running PXE with promiscuous VLAN. • Re-added COMBOOT image support by default. • Enabled pages-function handling in Connect-IB initialization stage to work according to the PRM. • Applied additional patches from ipxe.org • Updated the window even if ACK does not acknowledge new data. • Modified the error print to debug print. • Modified the printed string when initializing devices. • Modified the error print. Added additional information to make the output more user-friendly. • Changed the size of the domain name array to 0xfd. • Disabled the waiting period for link up on trunk-net-device when VLAN is enabled on port. • Removed unsupported EQ event in Connect-IB® • Fixed an issue for TLV with length 0. • Fixed an issue related to sync VLAN IRQ operation with trunk IRQ operation. • Fixed an issue which enabled a netdevice (VLAN) to open/close twice. • Fixed an issue which prevented the iSCSI initiator's name from being received from HII. • Fixed an issue related to dual port adapters; occasionally, booting from the second port resulted in TFTP download failure when the first port was already linked up with DHCP, and has received a TFTP address. • Fixed an issue which caused PXE boot failure when using a filename if iSCSI rootpath is set. • Fixed an issue which prevented the device to PXE boot from the 2nd port if first port was already downloaded. • Fixed compilation issue. • Fixed a broken VLAN issue. • Fixed a retry issue when the value is infinite.

Table 7 - FlexBoot Changes and New Feature

Version	Description
Rev 3.4.225	<ul style="list-style-type: none"> • Added additional information to the error print output • Added compilation flag around the flash readout • Added URI Boot retry. Default retries = 0 • Added Unmap MPT command in teardown • Added an option for running PXE with promiscuous VLAN • Enlarged the mailbox size to 4kb • Enlarged the number of WQE to 64 (from 4) • Enabled multiple DHCP offers to be received before proceeding to request state • Changed the size of the domain name array to 0xfd • Changed error print to debug print • Changed printed string when initializing devices • Kept the SAN connection permanently open to enable Windows install on iSCSI target even when the iSCSI target is empty • Re-added COMBOOT image support by default • Prevented a netdevice (VLAN) from opening/closing twice • Removed unsupported EQ event in Connect-IB® • Disabled the waiting time for link up on trunk net device when VLAN is enabled on a port • Fixed sync VLAN IRQ operation with trunk IRQ operation • Fixed an issue caused in dual port adapters, when the first port was already linked up with DHCP, and had received a TFTP address. Booting from the second port resulted in TFTP download failure. • Fixed retry issue when the value is infinite • Fixed a PXE boot failure issue occurred when using a filename when iSCSI rootpath is set • Fixed "Impossible to PXE boot from 2nd port if first port already downloaded." issue • Fixed compilation issue • Fixed broken VLAN support issues • Enlarged the mailbox size to 4kb • Enlarged the number of WQE to 64 (from 4) • Enabled multiple DHCP offers to be received before proceeding to request state • Fixed memory corruption issues • Added additional WQ