



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

Rev 10.16.1200

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT ("PRODUCT(S)") AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES "ASIS" 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 2017. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, Accelio®, BridgeX®, CloudX logo, CompustorX®, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, EZchip®, EZchip logo, EZappliance®, EZdesign®, EZdriver®, EZsystem®, GPUDirect®, InfiniHost®, InfiniBridge®, InfiniScale®, Kotura®, Kotura logo, Mellanox CloudRack®, Mellanox CloudXMellanox®, Mellanox Federal Systems®, Mellanox HostDirect®, Mellanox Multi-Host®, Mellanox Open Ethernet®, Mellanox OpenCloud®, Mellanox OpenCloud Logo®, Mellanox PeerDirect®, Mellanox ScalableHPC®, Mellanox StorageX®, Mellanox TuneX®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, NPS®, Open Ethernet logo, PhyX®, PlatformX®, PSIPHY®, SiPhy®, StoreX®, SwitchX®, Tiler®, Tiler logo, TestX®, TuneX®, The Generation of Open Ethernet logo, UFM®, Unbreakable Link®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

For the most updated list of Mellanox trademarks, visit <http://www.mellanox.com/page/trademarks>

Table of Contents

Chapter 1 Overview	8
1.1 Supported Devices	8
1.2 Supported Cables and Modules	9
1.2.1 Validated and Supported DDR/SDR Cables	9
1.2.2 Validated and Supported QDR/FDR10 Cables	9
1.2.3 Validated and Supported FDR Cables	10
1.2.4 Validated and Supported EDR Cables	11
1.3 Tested Switches	12
1.3.1 Tested QDR/FDR10 Switches	12
1.3.2 Tested FDR Switches	12
1.3.3 Tested EDR Switches	12
1.4 Tools, Switch Firmware and Driver Software	12
1.5 Revision Compatibility	13
Chapter 2 Changes and New Features in Rev 10.16.1200	14
Chapter 3 Known Issues	15
Chapter 4 Bug Fixes History	18
Chapter 5 Firmware Changes and New Feature History	23
Chapter 6 Unsupported Features and Commands	26
6.1 Unsupported Features	26
6.2 Unsupported Commands	26
Chapter 7 Supported Non-Volatile Configurations	27

List of Tables

Table 1:	Release Update History	7
Table 2:	Supported PSIDs	8
Table 3:	Validated and Supported DDR/SDR Cables	9
Table 4:	Validated and Supported QDR/FDR10 Cable	9
Table 5:	Validated and Supported FDR Cables	10
Table 6:	Validated and Supported EDR Cables	11
Table 7:	Tested QDR/FDR10 Switches	12
Table 8:	Tested FDR Switches	12
Table 9:	Tested EDR Switches	12
Table 10:	Tools, Switch Firmware and Driver Software	12
Table 11:	Changes and New Feature	14
Table 12:	Known Issues	15
Table 13:	Fixed Bugs List	18
Table 14:	Firmware Changes and New Feature History	23
Table 15:	Global Settings	27

Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 10.16.1200	November 30, 2017	Initial version of this firmware version

1 Overview

These are the release notes for the Connect-IB® adapters firmware, fw-ConnectIB Rev 10.16.1200. 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.

Table 2 - Supported PSIDs

Device Part Number	PSID	Device Name	Compiled with FlexBoot	Compiled with UEFI
MCB191A-FCAT	MT_1230110019	Connect-IB® Host Channel Adapter, single-port QSFP, FDR 56Gb/s, PCIe3.0 x8, tall bracket, RoHS R6	No	No
MCB192A-FCAT	MT_1240110019	Connect-IB® Host Channel Adapter, dual-port QSFP, FDR 56Gb/s, PCIe3.0 x8, tall bracket, RoHS R6	No	No
MCB193A-FCAT	MT_1220110019	Connect-IB® Host Channel Adapter, single-port QSFP, FDR 56Gb/s, PCIe3.0 x16, tall bracket, RoHS R6	No	No
MCB194A-FCAT	MT_1210110019	Connect-IB® Host Channel Adapter, dual-port QSFP, FDR 56Gb/s, PCIe3.0 x16, tall bracket, RoHS R6	No	No

1.2 Supported Cables and Modules

Please refer to the LinkX™ Cables and Transceivers web page

(<http://www.mellanox.com/products/interconnect/variable-optical-attenuators.php>) for the list of supported cables.

1.2.1 Validated and Supported DDR/SDR Cables

Table 3 - Validated and Supported DDR/SDR Cables

Speed	Cable OPN #	Description
DDR	MC1204128-005	Mellanox Passive Copper Hybrid Cable IB DDR 20GB/S QSFP TO CX4 5M
DDR	MC1204130-001	Mellanox Passive Copper Hybrid Cable IB DDR 20GB/S QSFP TO CX4 1M
DDR	MC1204130-002	Mellanox Passive Copper Hybrid Cable IB DDR 20GB/S QSFP TO CX4 2M
DDR	MC1204130-003	Mellanox Passive Copper Hybrid Cable IB DDR 20GB/S QSFP TO CX4 3M

1.2.2 Validated and Supported QDR/FDR10 Cables

Table 4 - Validated and Supported QDR/FDR10 Cable

Speed	Cable OPN #	Description
QDR	MC2206125-007	Mellanox Passive Copper Cable 4X QSFP 40GB/S 25AWG 7M
QDR	MC2206126-006	Mellanox Passive Copper Cable 4X QSFP 40GB/S 26AWG 6M
FDR10	MC2206128-004	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 4M
FDR10	MC2206128-005	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 5M
FDR10	MC2206130-001	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 1M
FDR10	MC2206130-002	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 2M
FDR10	MC2206130-003	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 3M
FDR10	MC2206130-00A	Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 0.5M
FDR10	MC2206310-300-L	Mellanox Active Fiber Cable IB QDR/FDR10 40GB/S QSFP 300M
FDR10	MC2206310-XXX-F	Mellanox Active Fiber Cable IB QDR/FDR10 40GB/S QSFP from 3M up to 100M
FDR10	MC2206310-XXX-T	Mellanox Active Fiber Cable IB QDR/FDR10 40GB/S QSFP from 3M up to 100M
FDR10	MC2210411-SR4	Mellanox Optical Module 40GB/S QSFP MPO 850NM UP TO 100M
FDR10	MC2210411-SR4E	Mellanox Optical Module 40GB/S QSFP MPO 850NM UP TO 300M
FDR10	MFS4R12CB-XXX	Mellanox Active Fiber Cable VPI UP TO 40GB/S QSFP from 3M up to 100M

1.2.3 Validated and Supported FDR Cables

Table 5 - Validated and Supported FDR Cables

Speed	Cable OPN #	Description
EMC FDR QSFP	038-004-066-01	EMC FDR QSFP+ to QSFP+ copper cable 2M
EMC FDR QSFP	038-004-067-01	EMC FDR QSFP+ to QSFP+ copper cable 3M
EMC FDR QSFP	038-900-027-01	EMC FDR QSFP+ to QSFP+ copper cable 5M
EMC FDR QSFP	038-900-030-01	EMC FDR QSFP+ to QSFP+ copper cable 8M
FDR	MC2207126-004	Mellanox Passive Copper Cable 4X QSFP 56GB/S 28AWG 4M
FDR	MC2207128-003	Mellanox Passive Copper Cable 4X QSFP 56GB/S 28AWG 3M
FDR	MC2207128-0A2	Mellanox Passive Copper Cable VPI UP TO 56GB/S QSFP 2.5M
FDR	MC2207130-001	Mellanox Passive Copper Cable VPI UP TO 56GB/S QSFP 1M
FDR	MC2207130-002	Mellanox Passive Copper Cable 4X QSFP 56GB/S 30AWG 2M
FDR	MC2207130-00A	Mellanox Passive Copper Cable 4X QSFP 56GB/S 30AWG 0.5M
FDR	MC2207130-0A1	Mellanox Passive Copper Cable VPI UP TO 56GB/S QSFP 1.5M
FDR	MC2207310-030	Mellanox Active Fiber Cable 4X QSFP 56GB/S 30M
FDR	MC2207310-100)	Mellanox Active Fiber Cable 4X QSFP 56GB/S 100M
FDR	MC2207310-XXX-E	Mellanox Active Fiber Cable VPI UP TO 56GB/S QSFP from 3M up to 100M
FDR	MC2207310-XXX-T	Mellanox Active Fiber Cable VPI UP TO 56GB/S QSFP from 3M up to 100M
FDR	MC2207312-003	Mellanox Active Fiber Cable VPI IB FDR (56GB/S) AND ETH 40GBE QSFP 3M
FDR	MC2207312-050	Mellanox Active Fiber Cable VPI UP TO 56GB/S QSFP from 3M up to 300M
FDR	MC2207312-100	Mellanox Active Fiber Cable 4X QSFP 56GB/S 100M
FDR	MC220731V-003	Mellanox Active Fiber IB QSFP 3M
FDR	MC220731V-100	Mellanox Active Fiber IB QSFP 100M

Table 5 - Validated and Supported FDR Cables

Speed	Cable OPN #	Description
FDR	MC2207411-SR4L	Mellanox Optical Module IB FDR 56GB/S QSFP MPO 850NM UP TO 30M

1.2.4 Validated and Supported EDR Cables

Table 6 - Validated and Supported EDR Cables

Speed	Cable OPN #	Description
EDR	MCP1600-E001	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 1M
EDR	MCP1600-E002	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 2M
EDR	MCP1600-E003	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 3M
EDR	MCP1600-E00A	Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 0.5M
EDR	MFA1A00-E003	Mellanox Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 3m
EDR	MFA1A00-E005	Mellanox Active Fiber Cable VPI UP TO 100GB/S QSFP 5M
EDR	MFA1A00-E050	Mellanox Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 50m
EDR	MFA1A00-E100	Mellanox Active Fiber Cable VPI UP TO 100GB/S QSFP 100M
EDR	MFS1200-EXXX	Mellanox® Active Fiber Cable, IB EDR, up to 100Gb/s, QSFP, LSZH, up to 100m
100GB/S	MMA1B00-C100	Mellanox® transceiver, 100GbE, QSFP28, MPO, 850nm, up to 100m
EDR	MMA1B00-E100	Mellanox® transceiver, IB EDR, up to 100Gb/s, QSFP28, MPO, 850nm, SR4, up to 100m

1.3 Tested Switches

1.3.1 Tested QDR/FDR10 Switches

Table 7 - Tested QDR/FDR10 Switches

Speed	Switch Family	OPN # / Name	Description
QDR	N/A	489184-B21	HP BLc 4X QDR IB Switch Module
QDR	N/A	QLogic 12300	36-Port 40Gb QDR Infiniband Switch, Management Module, Dual Power
QDR	InfiniScale® IV	IS5025Q-1SFC	36-port 40Gb/s InfiniBand Switch Systems
QDR	InfiniScale® IV	Switch 4036	Grid Director™ 4036E
FDR10	SwitchX®	SX6025T-1SFR	36-port 56Gb/s InfiniBand/VPI Switch Systems

1.3.2 Tested FDR Switches

Table 8 - Tested FDR Switches

Speed	Switch Family	OPN # / Name	Description
FDR	N/A	648312-B21	HP BLc 4X FDR IB Switch Module
FDR	N/A	648311-B21	HP BLc 4X FDR IB Managed Switch Module
FDR	SwitchX®	SX6036F-1BFR	36-port 56Gb/s InfiniBand/VPI Switch Systems
FDR	SwitchX®	SX6018F-1SFR	18-port 56Gb/s InfiniBand/VPI Switch Systems

1.3.3 Tested EDR Switches

Table 9 - Tested EDR Switches

Speed	Switch Family	OPN # / Name	Description
EDR	Switch-IB	SB7790-EB2F	36-port EDR 100Gb/s InfiniBand Switch Systems

1.4 Tools, Switch Firmware and Driver Software

Firmware Rev 10.16.1200 is tested with the following tools and driver software:

Table 10 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_OFED	4.2-1.1.x.x.

Table 10 - Tools, Switch Firmware and Driver Software

	Supported Version
MFT	4.8.1/4.8.0
MLNX-OS	<ul style="list-style-type: none"> • SwitchX: 3.6.5000 • Switch-IB: 3.6.5000
SwitchX®/SwitchX®-2 Firmware	9.3.7240
SwitchX-IB™ Firmware	11.1500.0106
SwitchX-IB 2 Firmware	15.1500.0106
InfiniScale® IV Firmware	7.4.3000
Linux Inbox Drivers	<ul style="list-style-type: none"> • RHEL6.8 • RHEL6.9 • RHEL7.2 • RHEL7.3 • Ubuntu 12.04 • Ubuntu 14.04 • SLES11.2 • SLES11.3 • SLES12.0

1.5 Revision Compatibility

Firmware fw-ConnectIB Rev 10.16.1200 complies with the following programmer's reference manual:

- *Mellanox Adapters Programmer's Reference Manual (PRM), Rev 0.42 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 Changes and New Features in Rev 10.16.1200

Table 11 - Changes and New Feature

Feature/Change	Description
User Access Region (UAR) 4KB Granularity Allocation	UAR page size currently is set to 4KB and not according to what the system page size determines.
Memory Usage Reduction	Reduced HCA memory consumption, which allowed PXE driver to consume 5MB less memory.

3 Known Issues

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

Table 12 - Known Issues (Sheet 1 of 3)

Internal Ref.	Description
1218195	Description: On rare occasions, the VF fails to initiate on the 2nd PCIe Physical Function
	Workaround: N/A
	Keywords: VF, PCIe PF
687113	Description: Some Port Control Register do not return to the default value after the last port owner host restarts the driver.
	Workaround: Reboot or reset the driver. <code>reboot / mlxfwreset</code>
	Keywords: PRM
688670	Description: Configuring the SM with VL weight 0 on some VL, and running traffic on it, causes the driver to hang during unload.
	Workaround: Cold power-cycle the server
	Keywords: QoS
680719	Description: When all SLs are mapped to non-VL0, the firmware might hang.
	Workaround: Fix the SL configuration and power cycle the system.
	Keywords: Virtualization
691754	Description: <code>end_padding_mode</code> is required in <code>CREATE_QP</code> and not in <code>INIT_2_RTR</code> command as defined in the PRM
	Workaround: N/A
	Keywords: <code>end_padding_mode</code> , PRM
747967	Description: Burning in firmware on the same device in parallel from multiple interfaces (e.g. PCIe and MTUSB) is not supported.
	Workaround: N/A
	Keywords: PCIe, MTUSB, burning in parallel
776830	Description: Performing warm reboot during firmware image burning for VPI/IB devices configured with IB port protocol, might cause the device to disappear from the PCIe.
	Workaround: Cold reboot the device instead
	Keywords: Warm/cold reboot
-	Description: Setting the port to 'sleep' state is not supported.
	Workaround: N/A
	Keywords: MADs
-	Description: Link width x1 might get Replay Timer Timeout, on speed change.
	Workaround: N/A
	Keywords: PCIe

Table 12 - Known Issues (Sheet 2 of 3)

Internal Ref.	Description
-	<p>Description: L1 power state enter requests are ignored by the device.</p> <p>Workaround: N/A</p> <p>Keywords: PCIe</p>
-	<p>[For customers developing custom low level drivers]</p> <p>Description: The device does not recover if the requested number of pages are not supplied during device initialization.</p> <p>Workaround: N/A</p> <p>Keywords: Device Initialization</p>
-	<p>Description: On rare occasions, SL to VL modification with functioning QPs results in traffic hangs.</p> <p>Workaround: Configure SL2VL during SM initialization</p> <p>Keywords: Quality of Service</p>
-	<p>Description: Vport transmit packets are not blocked if vport policy is Down.</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>
-	<p>Description: DC transport is not supported when SR-IOV is enabled.</p> <p>Workaround: N/A</p> <p>Keywords: DC transport</p>
-	<p>Description: PF direct pass-through is not supported (since PF FLR is not supported)</p> <p>Workaround: N/A</p> <p>Keywords: PF direct pass-through</p>
-	<p>Description: Older MFT versions (4.0.0 and 3.8.0) may indicate that the latest GA firmware is old or that it cannot be compared with the existing firmware. A message similar to the below will be displayed upon firmware upgrade stage:</p> <pre># flint -d <mst device> -i <image> burn Current FW version on flash: 12.1100.6630 New FW version: 12.0012.0572 Note: The new FW version is not newer than the current FW version on flash. Do you want to continue ? (y/n) [n] : y</pre> <p>Workaround: Choose one of the options below to upgrade firmware:</p> <ul style="list-style-type: none"> • Upgrade to the latest MFT version (4.1.0) • Type "y" after the note flint provides <p>Run flint with the "-force" flag</p> <p>Keywords: Firmware Upgrade/MFT</p>

Table 12 - Known Issues (Sheet 3 of 3)

Internal Ref.	Description
-	<p>Description: Flashing the firmware requires server reboot. Firmware cannot be flashed twice without server reboot after first flashing</p> <p>Workaround: Reboot the server after firmware flashing</p> <p>Keywords: Upgrading/Downgrading</p>
-	<p>[For customers developing custom low level drivers]</p> <p>Description: VFs internal FLR is not supported in PF teardown HCA command.</p> <p>Workaround: Before unloading the PF driver, PF driver must disable all its active VFs by performing the following:</p> <ol style="list-style-type: none"> 1. Run the <code>disable_hca</code> command on all the <code>function_ids</code> 2. Wait until firmware returns all VFs allocated pages. <p>Keywords: Virtualization, FLR</p>
-	<p>[For customers developing custom low level drivers]</p> <p>Description: VNodeInfo and VPortGuidInfo virtualization Attributes MADs are not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>
-	<p>[For customers developing custom low level drivers]</p> <p>Description: The value of <code>log_max_ra_res_qp</code> in <code>set_hca_cap</code> command should be the same in all functions.</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>
-	<p>Description: Function (PF/VF) TX port counters are not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>
-	<p>Description: PF driver must work with pages event queue.</p> <p>Workaround: N/A</p> <p>Keywords: Virtualization</p>

4 Bug Fixes History

Table 13 lists the bugs fixed in this release.

Table 13 - Fixed Bugs List

#	Category	Description	Discovered in Release	Fixed in Release
1.	HCA hang	#867896: Fixed an issue that cause the HCA to hang due to a fatal hardware error during resources creation.	10.16.1020	10.16.1200
2.	ECN	#943484: Fixed an issue that prevented the software from setting ECN parameters (min_rate, max_rate, rate_to_set_on_first_cnp) to values >32768	10.16.1020	10.16.1200
3.	ibdump	#1002884: Fixed an issue that prevented ibdump from functioning properly on the cards second port	10.16.1020	10.16.1200
4.	Firmware hang	#1126486: Fixed a rare issue that caused the firmware to hang.	10.16.1020	10.16.1200
5.	iSER	#1124226: Fixed an issue that caused QP connection timeout due to firmware not being able to handle duplicate packets with AckReq bit set. The fix stopped ignoring duplicate AckReq packets to avoid timeout on the sender side.	10.16.1020	10.16.1200
6.	Driver restart	#937018/1030160: Fixed a rare issue/race condition that might cause a double command completion or other corrupted completion.	10.16.1020	10.16.1200
7.	Counters	#1002190: Fixed an issue related to the <code>PortRcvDataVLExtended/PortXmitDataVLExtended</code> parameter that caused the counters' value to be reported in octets instead of dwrods.	10.16.1020	10.16.1200
8.	RX pipe	#1008546: Fixed a rare issue where the RX Pipe got stuck. This occurred only when using an IB transport QP connected to SRQ with SRQ Limit event.	10.16.1020	10.16.1200
9.	Function Level Reset (FLR)	#1023501: Fixed an issue that caused the RX pipe to get stuck in case FLR was performed in parallel with the traffic that was send either to an Ethernet RQ or an IB transport connected to SRQ.	10.16.1020	10.16.1200
10.	Bad DMA memory access	Fixed a rare issue that caused the HCA to access the "bad DMA memory access".	10.16.1020	10.16.1200
11.	PCIe	RM#996058: Fixed a PCI issue which resulted in fatal assert and caused a driver load failure.	10.16.1006	10.16.1200

Table 13 - Fixed Bugs List

#	Category	Description	Discovered in Release	Fixed in Release
12.	SR-IOV	#866181: Fixed an issue which caused system fail when enabled SR-IOV	10.14.2036	10.16.1200
13.	SRQ limit event	#787039: Fixed a rare issue which caused the RX to hang when triggered the SRQ limit event.	10.12.1100	10.16.1200
14.	RX traffic	#801374: Fixed an issue which occasionally caused the RX traffic to hang in DC when received a PCI error on WQE fetch.	10.16.1006	10.16.1200
15.	mlxconfig configuration	#785128: Fixed an issue which caused the mlxconfig configuration of VF_LOG_BAR_SIZE to be ignored and to be set to 5 (32MB).	10.16.1006	10.16.1200
16.	EEH error	#846520: Fixed an EEH error from PCI which caused firmware to hang.	10.16.1006	10.16.1200
17.	VL, SM	#689503: Fixed an issues which occasionally caused the driver to hang during unload on some VLs when configuring the SM with a VL weight 0 and running traffic on it.	10.16.1006	10.16.1200
18.	DC transport	#882234: Fixed a rare case which caused an assert reported to the driver when the DC transport was enabled in the following cases: retransmission occurred and the RX received the same packet twice	10.16.1006	10.16.1200
19.	VF vports/VF GUIDs	#897449: Fixed an issue which caused the HCA to hang when enabled /disabled the VFs vports when the VFs GUIDs configuration were overloaded in the steering table.	10.16.1006	10.16.1200
20.	mlxconfig	#785128: mlxconfig configuration of VF_LOG_BAR_SIZE and PF_LOG_BAR_SIZE are ignored and set to 5 (32MB).	10.16.1006	10.16.1020
21.	Non-volatile configuration, TLV	#690890: Updating a non-volatile configuration of port type TLV more than 50 times might cause system to hang.	10.14.1100	10.16.1006
22.	SR-IOV	#693832: In an SR-IOV setup, traffic should contain GRH (GID index), traffic without GRH will be forwarded to vport0 (“Host0”). OpenSM should be configured as follow (opensm.conf): <ul style="list-style-type: none"> • virt_enable should be 2 • Enable Qos: qos TRUE 	10.14.1100	10.16.1006
23.	Virtualization	#597718: Privileged Vport egress traffic is not blocked when Vport is not active.	10.14.1100	10.16.1006

Table 13 - Fixed Bugs List

#	Category	Description	Discovered in Release	Fixed in Release
24.	Link	#563626: ibstat reports the link speed as FDR instead of FDR10.	10.12.1100	10.16.1006
25.	Physical Port TLVs	#780651: Fixed an issue causing single port devices to query and write Physical Port TLVs to Port 2.	10.14.2036	10.16.1006
26.	qkey/pkey violation counter	#687096: Fixed an issue which caused the device to hang when resetting qkey/pkey violation counter via port_info mad.	10.14.1100	10.14.2036
27.	RDMA READ bandwidth	#665089: Improved RDMA READ bandwidth under packet lost scenario.	10.14.1100	10.14.2036
28.	Multihost loopback	#649696/690681: If the PF driver or the tool (e.g. ethtool) use PAOS DOWN command (e.g. by ifconfig down or ip link set down), loopback traffic is blocked for all functions on this port (PF<->VFs / VF<->VF) In Multihost loopback, the traffic will be blocked once the firmware receives the PAOS down command from all PFs. However, the loopback traffic will not be blocked when the port is down due to the physical link (for example: cable plugged out, switch port down).	10.14.1100	10.14.2036
29.	QP permission	#667224: Fixed an issue which prevented QP permission for reserve lkey to be passed to the memop machine.	10.12.0780	10.14.1100
30.	DC transport	#660719: Fixed a MLX QP SL mismatch handling which occurred when the SL in the WQE was different than the SL in the QP.	10.12.0780	10.14.1100
31.		#660939: Fixed wrongly implementation of SM SL2VL configuration.	10.12.0780	10.14.1100
32.		#660708: Fixed a DC re-connect flow which in some cases sent bad completion.	10.12.0780	10.14.1100
33.		#557979: Fixed a DC performance issue; separated DCRs SQ from the DCI SQs.	10.12.0780	10.14.1100
34.	MADs	#631226: Fixed an issue causing the firmware to hang when running ibdiagnet. The received DiagData MAD included the following values: <ul style="list-style-type: none"> • Clear_all = 1 • PageNum = 0 • Port_select = 0 To prevent the firmware from hanging, a port check was added to Set() as well.	10.12.0780	10.14.1100

Table 13 - Fixed Bugs List

#	Category	Description	Discovered in Release	Fixed in Release
35.	Diagnostic Tools	#552462: Fixed an issue which caused hardware fatal error when running ibdump.	12.0100.6440	10.12.0780
36.	Port Link	Fixed an FDR10 incorrect speed indication reported due to the usage of a translation function from the hardware speed to the PRM speed twice.	12.0100.6440	10.12.0780
37.	Phy Management	#592712: Fixed a Phy manager PCS event handling when the port's next state was disable.	12.1100.6630	10.12.0780
38.	MADs	#561387: Fixed an issue that caused invalid data returned by EyeOpening MAD.	12.0100.6440	10.12.0780
39.	Virtualization	#552227: Reduced the VF ICM footprint for VFs.	10.0100.6440	10.1100.6630
40.	Miscellaneous	#515812: Increased the number of regular memory region from 2 ²¹ to 2 ²² .	10.10.5052	10.0100.6440
41.	DC transport	#517702: Fixed health buffer false alarm "iris stuck" indication when using DC transport.	10.10.5020	10.10.5054
42.		#422852/518173: Fixed rare unexpected completion with error on DC QPs.	10.10.5020	10.10.5054
43.	Link Speed	#500090/493611: On rare occasions, after PXE boot, the port speed came up as SDR instead of a higher speed.	10.10.5020	10.10.5054
44.	Thermal sensors	#501903: On very rare occasions, firmware wrongly reported board over-temperature warning.	10.10.5020	10.10.5054
45.	DCT	destroy-DCT command handling may experience delays while the DCT port is down.	10.10.4064	10.10.5020
46.	MAD	Fixed an issue causing diagnostic counters VS-MAD page offset to start at a wrong address.	10.10.4020	10.10.5020
47.	DC transport	Fixed stability issue in the event of no-local-DC-resources	10.10.4020	10.10.4064
48.		Fixed improper handling of multiple DCT errors	10.10.4050	10.10.4064
49.		Fixed bad handling of DC RNR state	10.10.4050	10.10.4064
50.	DC transport	Reduced DCT destroy firmware handling time	10.10.4020	10.10.4050
51.	Ports	Fixed link flapping issue which occurred when LLR was active.	10.10.3000	10.10.4050
52.	PCIe class code of the device	Deprecated code 0x0c0600 was changed to 0x020700 (InfiniBand network adapter)	10.10.3000	10.10.4020

Table 13 - Fixed Bugs List

#	Category	Description	Discovered in Release	Fixed in Release
53.	Atomic support	Atomic response endianness is always a big endian	10.10.1000	10.10.3000
54.	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
55.	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
56.		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
57.	DC	On rare occasions DC Initiator completions might be lost.	10.10.2000	10.10.3000
58.	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
59.	Quality of Service	VL arbitration configuration does not ensure minimum bandwidth for VL as configured.	10.10.1000	10.10.2000
60.	False alarm report	On very rare occasions, a false firmware "hanged" report is printed in the <code>dmesg</code> .	10.10.1000	10.10.2000
61.	CQ buffer resize	CQ buffer resize not supported	10.10.1000	10.10.2000
62.	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 Firmware Changes and New Feature History

Table 14 - Firmware Changes and New Feature History

Firmware Version	Description
10.16.1020	<ul style="list-style-type: none"> Firmware v10.16.1020 code is the same as v10.16.1006. Its version number was changed to align to other family of products version numbering that use the same base code.
10.16.1006	<ul style="list-style-type: none"> 63 VFs per port: Increased the number of VFs from 32 to 63 per PF. <p>Note: When increasing the number of VFs, the following limitations must be taken into consideration:</p> <ul style="list-style-type: none"> $server_total_bar_size \geq (num_pfs) * (2log_pf_uar_bar_size + 2log_vf_uar_bar_size * total_vfs)$ $server_total_msix \geq (num_pfs) * (num_pf_msix + num_vfs_msix * total_vfs)$ <ul style="list-style-type: none"> Bug Fixes: See Section 4, “Bug Fixes History,” on page 18
10.14.2026	<ul style="list-style-type: none"> Bug Fixes: See Section 4, “Bug Fixes History,” on page 18
10.14.1100	<ul style="list-style-type: none"> CQE Time Stamping: Keeps track of the creation of a packet. A time-stamping service supports assertions of proof that a datum existed before a particular time. RDMA retransmission counters: Custom port counters provide the user a clear indication about RDMA send/receive statistics and errors. Link params modification via access registers: The change includes the following: <ul style="list-style-type: none"> Changed port configuration which required link re-training (such as speed) PAOS down PAOS up This change, will cause the link to toggle and new configurations to take effect. Checksum Calculation on Image/Device: Flint utility allows performing an MD5 checksum on the non-persistent sections of the firmware image. For further information, please refer to MFT User Manual.
10.12.1100	<ul style="list-style-type: none"> Improved robustness during negotiation of Clause 73 (DME)
10.12.0780	<ul style="list-style-type: none"> Dynamically Connected (DC) transport Non-Volatile Configuration (NVConfig). For the complete list, lease refer to Section 7, “Supported Non-Volatile Configurations,” on page 27.
10.1100.6330	<ul style="list-style-type: none"> Added support for SR-IOV Added support for MADs Virtualization Attributes according to “ib_virt_annex_v17” Updated Virtualization command set according to the new PRM. Enabled SR-IOV, NUM_VFS and INT_LOG_MAX_PAYLOAD_SIZE configuration via the mlxconfig tool

Table 14 - Firmware Changes and New Feature History

Firmware Version	Description
10.0100.6440	<ul style="list-style-type: none"> • T10 DIF pipeline Data Integrity Signature off-loading - Beta level • Automatic Path Migration • Congestion Control • Programmable Port/Node GUID • Thermal monitoring and protection • Port LEDs indications • Physical Port Counter - Beta level • Q Counter - Beta level • V port commands • Config space address in MAD management class 0x09
10.10.5054	<ul style="list-style-type: none"> • Enabled firmware burning (using mstflint) when the driver is down. • Bug Fixes (See Section 4, “Bug Fixes History,” on page 18.)
10.10.5020	<ul style="list-style-type: none"> • Added support for EDR cables (up to FDR speeds). • Added additional InfiniBand spec optional counters • Improved resource allocation for absolute-priority VL's (when VL-high-limit is set to 255 in SM) to reduce latency of high priority traffic. • Bug Fixes See Section 4, “Bug Fixes History,” on page 18.
10.10.4064	<ul style="list-style-type: none"> • Bug Fixes (see Section 4, “Bug Fixes History,” on page 18)
10.10.4050	<ul style="list-style-type: none"> • 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 18)
10.10.4020	<ul style="list-style-type: none"> • 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
10.10.3000	<ul style="list-style-type: none"> • 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

Table 14 - Firmware Changes and New Feature History

Firmware Version	Description
10.10.2000	<ul style="list-style-type: none"> • Dynamically Connected (DC) transport (at Beta level) • CORE-Direct® <ul style="list-style-type: none"> • Provides Collective Off-loading in HCA • Frees CPU to perform computation in parallel with collective operations • T10 DIF Data Integrity Signature off-loading • 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
10.10.1000	<ul style="list-style-type: none"> • 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

6 Unsupported Features and Commands

6.1 Unsupported Features

The following advanced features 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)
- 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
- PCIe Function Level Reset (FLR)
- SM is not supported on VFs

6.2 Unsupported Commands

- QUERY_MAD_DEMUX
- SET_MAD_DEMUX
- PAGE_FAULT_RESUME
- ACTIVATE_TRACER
- DEACTIVATE_TRACER

7 Supported Non-Volatile Configurations

Table 15 - Global Settings

Name	Parameter Index
PCI settings	0x80
PCI setting capabilities	0x81
TPT settings	0x82
TPT capabilities	0x83
Option ROM ini	0x100
Option ROM capabilities	0x101

