



Mellanox ConnectX[®]-3 Firmware Release Notes for

Rev 2.42.5000



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 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 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

Release Update History	5
Chapter 1 Overview	6
1.1 Supported Devices	6
1.2 Tested BIOS	6
1.3 Supported Switches	6
1.4 Tools, Switch Firmware and Driver Software	6
1.5 Supported FlexBoot, UEFI and CLP	7
1.6 Revision Compatibility	7
Chapter 2 Firmware Rev 2.42.5000 Changes and New Features	8
Chapter 3 Known Issues	9
Chapter 4 Bug Fixes History	12
Chapter 5 Firmware Changes and New Feature History	28
Chapter 6 FlexBoot Changes and New Features	34
6.1 Flexboot Known Issues	39
Chapter 7 CLP Changes and New Features	40

List of Tables

Table 1:	Release Update History	5
Table 2:	Supported PSIDs	6
Table 3:	Tested BIOS	6
Table 4:	Tools, Switch Firmware and Driver Software	6
Table 5:	Supported FlexBoot, UEFI and CLP	7
Table 6:	Firmware Rev 2.42.5000 Changes and New Features	8
Table 7:	Known Issues	9
Table 8:	Fixed Bugs List	12
Table 9:	Firmware Changes and New Feature History	28
Table 10:	FlexBoot Changes and New Feature	34
Table 11:	Known Issues	39
Table 12:	CLP Changes and New Feature	40



Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 2.42.5000	September 7, 2017	Initial version of this firmware version

1 Overview

These are the release notes for firmware Rev 2.42.5000 of
Firmware revision Rev 2.42.5000 supports the following features:

- PCIe 3.0 and PCIe 2.0
- HP OCSD thermal sensors reporting
- Flexboot - PXE
- UEFI
- CLP

1.1 Supported Devices

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

1.2 Tested BIOS

This firmware was tested with the following BIOS versions.

1.3 Supported Switches

1.4 Tools, Switch Firmware and Driver Software

Firmware Rev 2.42.5000 is tested with the following tools, SwitchX® firmware, and driver software:

Table 4 - Tools, Switch Firmware and Driver Software

	Supported Version
MLNX_OFED	4.1-1.0.2.0/3.4-1.0.0.0
MLNX_EN (MLNX_OFED based code)	4.1-1.0.2.0/3.4-1.0.0.0
WinOF	5.35 FUR3/5.22
VMware	2.4.0 3.15.5.5 1.8.2.5
MFT	4.7.1
MLNX-OS	<ul style="list-style-type: none"> • SwitchX: 3.6.4006 • Switch-IB: 3.6.4006
SwitchX/SwitchX-2 Firmware	9.4.3580
Switch-IB Firmware	11.1300.0126
InfiniScale 4 Firmware	7.4.3000

Table 4 - Tools, Switch Firmware and Driver Software

	Supported Version
Linux Inbox Drivers	<ul style="list-style-type: none"> • RH6.6 • RH6.7 • RH6.8 • RH7.0 • RH7.1 • RH7.2 • SLES11 SP3 • SLES11 SP4 • SLES12 SP0 • SLES12 SP1 • Fedora23 • Ubuntu 14.04 • Ubuntu 14.10 • Ubuntu 15.04 • Ubuntu 15.10 • Ubuntu 16.04
Windows Inbox Driver	<ul style="list-style-type: none"> • Windows Server 2012 • Windows Server 2012 R2

1.5 Supported FlexBoot, UEFI and CLP

Firmware Rev 2.42.5000 supports the following FlexBoot, UEFI and CLP:

Table 5 - Supported FlexBoot, UEFI and CLP

	Supported Version
FlexBoot	3.4.752
CLP	8025
UEFI	

1.6 Revision Compatibility

Firmware fw-ConnectX3 Rev 2.42.5000 complies with the following programmer’s reference manual:

- *Mellanox Adapters Programmer’s Reference Manual (PRM), Rev 2.1 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 Firmware Rev 2.42.5000 Changes and New Features

Table 6 - Firmware Rev 2.42.5000 Changes and New Features

Category	Description
Packet Time-stamping	Added support for new TLV: CX3_GLOBAL_CONF to enable/disable time-stamp on incoming packets through mlxconfig configuration.
MAC Configuration	Added support for user MAC configuration.
mstdump	Added support for automatically collecting mstdump before driver reset.
IRISC stuck watchdog	Added a mechanism to detect DEAD_IRISC (plastic) from TPT (iron) and raise an assert.
Debug ability	Improved the debug ability for command timeout cases
MTU size	Added a new field to "set port" command which notifies the firmware what is the user_mtu size.
Bug fixes	See Section 4, "Bug Fixes History," on page 12

3 Known Issues

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

Table 7 - Known Issues

Index	Issue	Description	Current Implemented Workaround in FW
1.	mlxconfig	RM#1119109: Enabling/disabling cq_timestamp using mlxconfig is not supported.	N/A.
2.	LEDs	RM#1121959: In a card with 2 separate LEDs scheme (a Phy LED and a logic LED) only the Phy LED will lit. Meaning, the orange LES will not be active while the ETH link is in an idle mode.	N/A.
3.	mlxconfig/SR-IOV	RM#976761: In SR-IOV setup, using mlxconfig when the PF is passed through to a VM requires a reboot of the Hypervisor.	N/A.
4.	Downgrade to previous GA requires server reboot.	Downgrading from v2.30.8000 or later to an earlier version than 2.30.8000 requires server reboot.	Reboot the server.
5.	GUID ConnectX®-3 Ethernet adapter cards	On ConnectX-3 Ethernet adapter cards, there is a mismatch between the GUID value returned by firmware management tools and that returned by fabric/driver utilities that read the GUID via device firmware (e.g., using ibstat). Mlxburn/flint return 0xffff as GUID while the utilities return a value derived from the MAC address. For all driver/firmware/software purposes, the latter value should be used.	N/A. Please use the GUID value returned by the fabric/driver utilities (not 0xffff).
6.	SBR assertion	SBR should be asserted for a minimum of 50 milliseconds for the ConnectX-3 adapters	N/A
7.	PCIe	On Pilot1 SL230, PCIe link occasionally does not come up at Gen3 speed	Production SL230 should be used for PCIe Gen3 operation
8.	Kernel panic in SR-IOV with RH6.3 Inbox driver and VPI cards	RH6.3 Inbox driver causes kernel panic when SR-IOV is enabled on VPI cards due to driver compatibility issue.	Set the "do_sense=false" parameter in the [IB_TAB] in the INI of the VPI card
9.	Side band Management compatibility with SR-IOV	In advanced steering mode, side band management connectivity may be lost when having more than 8 QP per mcg.	N/A
10.	SR-IOV disabled in the BIOS	When SR-IOV is disabled in the system BIOS, a PCI issue is noticed in Ubuntu v12.04.3 with Linux kernel v3.8 which affects NICs of several manufacturers including Mellanox's, preventing them from operating.	Enable SR-IOV in the BIOS

Table 7 - Known Issues

Index	Issue	Description	Current Implemented Workaround in FW
11.	MFT locking of flash semaphore	MFT tools might leave the flash semaphore locked if the tool operation is forced stopped. The locked semaphore prevents the firmware from accessing the flash and causes firmware hang.	Clear the semaphore using MFT command: 'flint -clear_semaphore'
12.	MC2210411-SR4 module with Cable Info MAD	Cable Info MAD reports a wrong cable info when using the MC2210411-SR4 module	N/A
13.	PCIe failure on temperature shock 10C/min	Gen2 failure at temperature sweep up to 10C/min (for MT27518A1-FDIR-BV only).	N/A
14.	PCIe Gen2 link	PCIe Gen2 link unstable at temperature sweep of 10C/min for MT27518A1-FDIR-BV	N/A
15.	Changing from an LLR to non-LLR requires driver restart	Driver restart required when switching from InfiniBand FDR link with LLR enabled to InfiniBand link w/o LLR (for example: between SwitchX® and GD4036).	N/A
16.	Bloom filter	Bloom filter is currently not supported.	N/A
17.	Firmware downgrade	When downgrading from firmware v2.11.0000 and using MFT 3.0.0-3, the following message is displayed due to the mlxconfig tool: You are trying to override configurable FW by non-configurable FW. If you continue, old FW configurations will be cleared, do you want to continue ? (y/n) [n] : y You are trying to restore default configuration, do you want to continue ? (y/n) [n] : y	N/A
18.	DMFS steering mode with IB in Linux	DMFS should not be enabled when working with InfiniBand on MLNX_OFED-2.0.3	Upgrade to MLNX_OFED-2.1-x.x.x or later
19.	VPD read-only fields	VPD read-only fields are writable.	Do not write to read-only fields if you wish to preserve them
20.	Increasing SymbolErrorCounter	When working in VPI mode with port1 FDR and port2 40G, error counters misbehave and increase rapidly	N/A
21.	128 Byte CQ/EQ stride compatibility with sideband Management	Setting the device to 128Byte CQ/EQ stride will cause misbehavior of sideband management resulting in communication loss.	N/A
22.	128 Byte CQ/EQ stride	CQ and EQ cannot be configured to different stride sizes.	N/A

Table 7 - Known Issues

Index	Issue	Description	Current Implemented Workaround in FW
23.	VPI port protocol change on a port with sideband Management	Changing port protocol from ETH to IB on port with NCSI/IPMI enabled while the port is connected to ETH switch is not supported.	1. Unplug the cable from the switch 2. Restart driver 3. Change the protocol via the appropriate tools.
24.	Link Up time	Adapter card MCX349A-XCCN may experience longer linkup times of a few seconds with specific switches.	N/A
25.	Port identification	Adapter card MCX349A-XCCN does not respond to ethtool “identify” command (ethtool -p/--identify).	N/A
26.	RDP over IPv6	RDP over IPv6 is currently not functional.	N/A
27.	Unicast/Multicast sniffer	Sniffer QP cannot be removed from the regular rule after adding the QP with insertion scheme equals to “push to that rule”	N/A
28.	Boot Entry Vector (BEV)	Since only a single Boot Entry Vector (BEV) per PCI Physical Function is supported, disabling the first port causes the second port to disappear as well.	N/A
29.	Cables	The NIC does not notify the driver of a link-down incident when a cable is unplugged from a NIC port with 56GbE port link.	N/A
30.	Port Link	56GbE link is not raised when using 100GbE optic cables.	N/A
31.	Server reboot	When working with MLNX_OFED v3.3-1.0.0.0, server reboot could get stuck due to a kernel panic in mlx4_en_get_drvinfo() that is called from asynchronous event handler.	N/A
32.	ibdump	832298When running ibdump, loopback traffic is mirroring into the kernel driver.	N/A
33.	MAC address	RM#846523: MAC address that are set from the OS using ifconfig are not reflected in the OCBB buffer.	N/A

4 Bug Fixes History

Table 8 lists the bugs fixed in this release.

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
1.	PortRevPkts counter	1009607: Fixed an issue which prevented the PortRevPkts counter from being cleared after resetting it.	2.40.5030	2.42.5000
2.	FLR, System Time Out, VFs	999432: Fixed an issue which caused a system Time Out on the configuration cycle of the VFs when more than 10 Virtual Functions performed FLR and the completion Time Out value was configured to a range of less than 16 msec.	2.40.5030	2.42.5000
3.	mlxftop	1034523: Fixed an issue that caused the server to hang and result in NMI when running “mlxftop -d mt4103_pci_cr0” while restarting the driver in parallel (from a different thread). In this case, the downstream bridge over the device reported completion timeout error.	2.40.5030	2.42.5000
4.	Flow steering, BMC	659925: Fixed an issue in flow_steering where BMC could not receive a ping over IPV6 after running bmc_reboot.	2.40.7000	2.42.5000
5.	RX packet	825412: Fixed an issue while closing the HCA, where RX packet caused bad access to resources that did not exist, and consequently caused the QPCGW or the irisc to get stuck.	2.32.5100	2.42.5000
6.	masterSMLID, LID	1033071: Fixed an issue where the masterSMLID and the LID was either 0 or 0xFFFF when the port was neither active nor armed.	2.40.7000	2.42.5000
7.	ibdump	1009736: Fixed an issue that prevented ibdump from capturing all MADs packets.	2.40.7000	2.42.5000
8.	Link Down	1000626: Fixed an issue that prevented the link to go up after reboot.	2.40.5030	2.42.5000
9.	PCIe	954259: Fixed a rare issue that cause the PCIe configuration cycle that arrived during the time of sw_reset to generate 2 completions.	2.36.5100	2.42.5000
10.	NC-SI	908959: Fixed an issue that caused NC-SI not to work when adding the <code>disable_static_steering_ini</code> field in the ini file, due to memory allocation issue for this field in the scratchpad.	2.36.5100	2.42.5000

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
11.	Driver Start	890373: Fixed a race between the firmware and the hardware during driver start which blocked outbound completions.	2.40.5000	2.40.5030
12.	Link Down	939162: Fixed an issue which caused the firmware not to send link_down event to the driver when running the close_port command.	2.40.5000	2.40.5030
13.	Auto Sense	861646: Fixed an issue where in rare cases the Auto Sense failed to detect the right protocol.	2.35.5100	2.40.5000
14.	Signal Integrity	780205: Fixed signal integrity issue when connecting a WCS ConnectX4 mezz card to Pikes peak FPGA.	2.35.5100	2.40.5000
15.	DME pages	Added the option to transmit corrupted DME pages for a very short period of time at the beginning of the Auto-Negotiation flow.	2.36.5000	2.40.5000
16.	Counters	877613: Fixed an incorrect report of the PortRcvDataVLExtended/PortXmitDataVLExtended counters by the firmware.	2.35.5000	2.40.5000
17.	Firmware's Packet Injector	870787: Fixed a rare issue which caused firmware's packet injector to cut off packets when the TX was congested.	2.35.5100	2.40.5000
18.	TX requests	702752: Fixed an issue that caused the response to TX requests to take up to 10 milliseconds in IEEE clause 72 Link Training.	2.34.5000	2.40.5000
19.	ECN	70454: Fixed a race between 2 iriscs which caused a QP to get stuck in burst control limit state	2.36.5000	2.40.5000
20.	CQE	748455: When a QP was in error state, the firmware generated too many err CQEs at once, thus causing the cmdif responsiveness to be too slow. To prevent the above, the number of err CQEs was limited to 16 at a time.	2.36.5000	2.40.5000
21.	MAC address	846523: Fixed an issue that caused the MAC address that was set from the OS using ifconfig to be not reflected in the OCBB buffer.	2.36.5000	2.40.5000
22.	ibdump	832298: Fixed an issue where the ibdump got broken when running with loopback traffic.	2.36.5000	2.40.5000
23.	QP to Firmware ownership	745727: Fixed an issue where the firmware took QP to firmware ownership and then released it to the hardware ownership without checking if another firmware flow owns the same QP.	2.36.5000	2.40.5000

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
24.	Cables	806288: Fixed an issue which occurred after disconnecting cable which showed the link type as IB even if the link type of the port is ETH.	2.36.5000	2.40.5000
25.	HCA PoerXmitWait counter	778739: Fixed an issue related to the HCA PoerXmitWait counter on port 2 (connected to port 2 on Switch-IB) where it started counting and reached 0xFF's regardless of connection to switch.	2.36.5000	2.40.5000
26.	ECN	Fixed a completion error issue when ECN was enabled. The ECN usage caused ordering errors in completion queues (CQ).	2.33.5000	2.36.5000
27.		Fixed the length calculation of UDP. The incorrect UDP length in the CNP packet caused miss-calculation of the ICRC.	2.35.5100	2.36.5000
28.	Cable Info MAD	Fixed a wrong returned status in cable info MAD when the cable was not connected.	2.35.5100	2.36.5000
29.	FLR device reset	Fixed failure instances when initiating FLR in the Physical Function.	2.35.5100	2.36.5000
30.	High rate steering mode	Disabled High Rate Steering mode in the INI to enable its compatibility with NC-SI over VLAN.	2.32.5100	2.36.5000
31.	Performance	Fixed performance issues causing slow performance when running in NO-DRAM-NIC mode.	2.34.5000	2.36.5000
32.	RDP over IPv4	Fixed a default hardware configuration issue which caused RDP over IPv4 traffic to be dropped.	2.30.8000	2.35.5100
33.	SR-IOV security	Prevented a Virtual Function from injecting pause frames into the network.	2.30.8000	2.35.5100
34.	NC-SI	MLNX_OEM command GET_TEMP returned a wrong value in the max_temp field.	2.34.5000	2.35.5100
35.	MTU exceptions	Fixed an issue which caused TX traffic to stop when the message MTU size was larger than QP.mtu.	2.32.5100	2.35.5100
36.	NVCONFIG failure	Fixed an issue which caused NVCONFIG to fail when the number of sector was set to 1 and the sector was zeroed.	2.34.5000	2.35.5100
37.	IB/RoCE retransmission	Fixed a race in handling a duplicated “read request from middle”.	2.34.5000	2.35.5100

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
38.	IB traffic issues	Fixed an issue which caused lack of IB traffic on SR-IOV VPI.	2.33.5000	2.35.5100
39.	NVRAM issues	Fixed an issue which caused NVRAM to get stuck when it filled non-valid information in TLV.	2.34.5000	2.35.5100
40.	IB APM	Fixed an issue which caused an internal firmware error when APM changed the QPs port mapping.	2.33.5100	2.35.5100
41.	QP alternate context error	Fixed an issue which caused a firmware internal error when handling QP alternative context.	2.34.5000	2.35.5100
42.	Flow Control security issue	Fixed an issue which caused packet transmission to get stuck when the software tried to send pause frames with dmac equal to one of the device's MAC addresses.	2.32.5100	2.35.5100
43.	Packet Ethertype	Fixed a mistakenly dropped ETH packet with ethertype 0x600 by the NIC.	2.30.8000	2.34.5000
44.	Broadcast traffic lost	Fixed a case preventing broadcast traffic from arriving to their destination after detaching high priority broadcast rule on a port where NC-SI was enabled.	2.33.5100	2.34.5000
45.	Low link speed	Fixed an issue where the port raised as SDR vs. InfiniScale IV QDR Switch	2.33.5100	2.34.5000
46.	RDMA read retransmission	Fixed a rare case of completion Error with Bad Opcode sequence status which occurred when retransmitting read requests.	2.33.5100	2.34.5000
47.	VM QoS	Fixed a case where the actual bandwidth did not match the user settings in VM QoS.	2.33.5100	2.34.5000
48.	Sideband communication loss	Fixed a case where on rare cases, communication to BMC was lost during driver initialization.	2.33.5100	2.34.5000
49.	Link down on cable plugging	Fixed an issue with cable reading, which caused the link not to raise	2.33.5100	2.34.5000
50.	PRM: EQN range	Set the maximum EQN number to 1024.	2.30.8000	2.34.5000
51.	Vital Product Data read failure	Fixed a rare issue with VPD init flow which caused read failures.	2.31.5050	2.34.5000
52.	PRM: Statistic counters not reported	Fixed an issue with RX size counter not being reported.	2.30.8000	2.34.5000
53.	RoCE/InfiniBand reliable connection	The first Read response was not treated as implicit ACK.	2.30.8000	2.33.5100

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
54.	40GbE Link up time	Reduced a long 40GbE link up time with Cisco Nexus3064 and Arista-7050S	2.32.5100	2.33.5100
55.	Promiscuous mode	Fixed promiscuous mode compatibility with A0-DMFS steering.	2.32.5100	2.33.5000
56.		Fixed promiscuous mode compatibility when NC-SI is enabled and configured.	2.32.5100	2.33.5000
57.	NC-SI OEM commands	Fixed sending/receiving OEM temp commands (set/get) with channel ID 0x1f failure.	2.32.5100	2.33.5050
58.	Packet Drops	Fixed an issue which caused packets to drop on a port when changing the interface state of the other port.	2.32.5100	2.33.5050
59.	Side Band Functionality	Fixed long management communication loss and SOL hang during reboot cycles.	2.32.5100	2.33.5050
60.		Fixed wrong processing of inbound traffic towards BMC which caused communication loss.	2.32.5100	2.33.5050
61.		Fixed management link loss upon closing port interface through the driver.	2.32.5100	2.33.5050
62.	NC-SI on SFP+ Adapter Cards	Fixed a false indication in firmware of an expander presence causing delay in EEPROM reading.	2.32.5100	2.33.5050
63.	Port Links	Fixed an issue which caused a link down on a port when the cable was removed from the other port.	2.32.5100	2.33.5050
64.	Inbound Packet Processing	Fixed a rare case where packet with length zero got stuck in hardware queues.	2.32.5100	2.33.5050
65.		Fixed an issue which caused InfiniBand congestion control packet (CNP) to hang in hardware.	2.32.5100	2.33.5050
66.	Asynchronous Event Notification (AEN)	Fixed an issue which caused AEN to be sent after channel reset.	2.32.5100	2.33.5050
67.	Bandwidth Degradation with QoS	Fixed an issue which prevented the restoring of QoS setting to its default consequently causing bandwidth degradation.	2.31.5050	2.33.5050
68.	Port Link Up Time	Fixed an occasional long link up time with 10GbE based devices.	2.32.5100	2.33.5050
69.	SFP Cable Reading	Fixed an issue preventing cable readings from i2c slave address 0x51	2.32.5100	2.33.5050
70.	PCIe Gen3 Equalization	Fixed a wrong parity bit calculation when transmitting PCIe TS1 packets.	2.32.5100	2.33.5050

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
71.	PCIe Power Management	Fixed a possible deadlock in PM turnoff request transmission and ack acceptance flow.	2.32.5100	2.33.5050
72.	PCIe width Degrade	Fixed a rare case with alignments state machines which caused occasional width degradation.	2.32.5100	2.33.5050
73.	Rate Limiters Hang with ECN/QCN Enabled	Fixed an issue where the transmit queues hanged while congestion control was enabled and operational (EQC/QCN)	2.32.5100	2.33.5050
74.	Unexpected Completion Syndrome with Status 0x77	Fixed an unexpected work completion syndrome with vendor syndrome 0x77 received when running RDMA SEN/WRITE traffic with retransmissions.	2.30.8000	2.33.5050
75.	IB Spec MADs	Fixed an issue which caused <code>SetPortInfo</code> to return a good status when receiving invalid <code>LinkSpeedEnabled</code> value.	2.32.5100	2.33.5050
76.	GPIO Mapping	Fixed an issue which caused dual port SFPP module cards to be automatically mapped with expander	2.32.5100	2.33.5050
77.	Steering Mode	Fixed an issue where firmware overrides the steering mode that was chosen by the driver.	2.32.5100	2.33.5050
78.	Port sensing	Fixed invalid return sensing results occurred when the link was up.	2.32.5100	2.33.5050
79.		Fixed an issue causing the sensing result to be delayed when cable was unplugged.	2.32.5100	2.33.5050
80.	Wrong link type display	Fixed an issue causing the link type to be displayed as ETH when set to AUTO.	2.32.5100	2.33.5050
81.	IBDump performance	Fixed performance degradation when running IBDump	2.30.8000	2.32.5100
82.	PCIe link Disable/Enable	Occasionally, a link training timeout occurred in EQ phase0 during disable/enable test.	2.31.5050	2.32.5100
83.	40GbE QoS	Improved strict bandwidth mode functionality	2.30.8000	2.32.5100
84.	Port Counters reporting	Fixed an issue with the <code>PortRcvPkts</code> counter always displaying zero value.	2.31.5050	2.32.5100
85.	GMP MADs in SecureHost	Fixed an issue with processing GMP MADs with SET method in SecureHost mode.	2.31.5050	2.32.5100
86.	NC-SI over IPv6	Fixed an issue causing a wrong usage of MCG size when configuring Global Multicast filter	2.31.5050	2.32.5100
87.	NC-SI link failure	Disabling the first port occasionally causes second port TX failure.	2.31.5050	2.32.5100

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
88.	10GbE link failure	Fixed a mismatch in links status reported. The adapter reports links as down while the switch perceives them as up	2.31.5050	2.32.5100
89.	Link failure	Fixed an occasional 40GbE link failure with SCM5 Switch blade	2.31.5050	2.32.5100
90.	ExtPortInfo MAD	Fixed a wrong FDR10 speed reporting in MAD	2.31.5050	2.32.5100
91.	IB link failure	Fixed an issue preventing the ports to rise up when set to FDR10 vs QDR	2.31.5050	2.32.5100
92.	40GbE link failure	Fixed an occasional link failure vs Arista switch	2.31.5050	2.32.5100
93.	RDMA Write retransmission	Retransmission started from the first PSN of message instead of the last acknowledged PSN	2.30.3200	2.32.5100
94.	Firmware burning	Firmware hangs when receiving <code>GeneralInfoMad</code> during inline firmware burning	2.30.3200	2.32.5100
95.	PCIe PML1	L1 flow adjustments and threshold tuning	2.31.5050	2.32.5100
96.	PCIe reset	Fixed a rare hanging issue during <code>PERST_assertion</code>	2.31.5050	2.32.5100
97.	PCIe Gen3 EQ	Wrong coefficients were reported during phase3	2.31.5050	2.32.5100
98.	Boot	Fixed an issue causing wrong behavior due to reset timing	2.31.5050	2.32.5100
99.	SMBUS	Fixed long timeout issues	2.31.5050	2.32.5100
100.	NVRAM	Fixed NVRAM write issues in driver-less mode	2.31.5050	2.32.5100
101.	40GbE Link support	Fixed 40GbE link support in aux mode	2.31.5050	2.32.5100
102.	NC-SI	Dropped commands with non-existing channel ID	2.31.5050	2.32.5100
103.	PRM PortInfo command	Fixed issues in extended speed reporting	2.31.5050	2.32.5100
104.	Trap 257/8(IB)	Fixed bad QP reporting in trap 257/8	2.30.8000	2.32.5100
105.	Bad Q_KEY errors	Fixed an issue causing false bad q_key error messages	2.30.8000	2.32.5100
106.	PFC	Fixed Pause Frame opcode mismatch	2.30.8000	2.32.5100
107.	Sideband Communication	Fixed communication loss upon PCIe error detection	2.31.5050	2.32.5100

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
108.	NC-SI	Fixed wrong channel value in the SELECT/DESELECT PACKAGE commands	2.30.8000	2.31.5050
109.		Fixed an issue caused response packet to include 4 extra bytes	2.30.8000	2.31.5050
110.		Fixed wrong reason code value returned when using Set Link command with unsupported speed.	2.30.8000	2.31.5050
111.		Added protection from bad MAC address given by BMC	2.30.8000	2.31.5050
112.	PCIe	Removed false TX pulse after PERST_ deassertion	2.30.8000	2.31.5050
113.		Fixed FLR capability bit inconsistency when SR-IOV is enabled.	2.30.3200	2.31.5050
114.		Fixed an issue with the device not reporting PCIe related errors.	2.30.8000	2.31.5050
115.	SDR instead of DDR ConnectX-3 to SX6036	When a link is configured to DDR in a setup of ConnectX-3 to SX6036, SDR link is established instead.	2.30.8000	2.31.5050
116.	VXLAN	VXLAN used the wrong default UDP port. the UDP port number was changed to 4789.	2.30.8000	2.31.5050
117.		Fixed wrong setting of the UDP destination port for VXLAN.	2.30.8000	2.31.5050
118.	Flow Steering	Fixed an internal error caused when moving to the DMS mode with IPMI/NC-SI enabled.	2.30.8000	2.31.5050
119.	FDR speed degradation with 0.5m cables	In a back-to-back setup of FDR cards connected with a 0.5m FDR cable, a link may be established as FDR10 instead of FDR.	2.30.3200	2.31.5050
120.	PCI interrupt	Fixed issues related to working with PCI legacy interrupts.	2.30.8000	2.31.5050
121.	TCP/UDP Checksum	Wrong checksum calculation for short packets which are padded by the software.	2.30.8000	2.31.5050
122.	MFT tool deadlock	Reading PCIe configuration space after using the MFT flint tool caused the device to crash.	2.10.0000	2.31.5050
123.	Side band packet loss	Fixed occasional packet loss over IPMI	2.30.8000	2.31.5050
124.	Eye opening MAD	Fixed wrong values reported in the Eye opening MAD.	2.30.8000	2.31.5050
125.	PCIe Link width	Fixed occasional link width degrades during link negotiation and link transitions from L1 state.	2.30.8000	2.31.5050

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
126.	PCIe signal detect	Fixed adjust signal detect thresholds	2.30.8000	2.31.5050
127.	Error counters	PortExtendedSpeedsCounters MAD counters were mistakenly increased while LLR was active	2.30.8000	2.31.5050
128.	PCIe Gen3 Equalization	Lane reversal was not considered when configured TX parameters	2.30.8000	2.31.5050
129.	Reset On LAN (ROL)	Fixed ROL factory MAC usage when a Flex-Boot address was given.	2.30.8000	2.31.5050
130.	Flow Control	Fixed Pause frames factory MAC usage when a FlexBoot address was given.	2.30.3200	2.31.5050
131.	WOL/ROL	The device did not different between WOL/ROL packets.	2.30.8000	2.31.5050
132.	PortInfo MAD	Fixed a set of extended fields in PortInfo MAD which did not function.	2.10.0000	2.31.5050
133.	LLR cell size	Adjusted LLR cell size according to the MLPN negotiation of ib_128b_llr	2.30.8000	2.31.5050
134.	Link max speed	The max speed restriction was active in full power mode instead of standby mode only.	2.30.8000	2.31.5050
135.	InfiniBand Automatic Path Migration	The InfiniBand Path migration did not work with GRH. http://webdev01:8080/commit/ConnectX.git/a9c37ee4c31038f2c1179d4d9e79c9337e0ab5c7	2.10.0000	2.31.5050
136.	Packet steering	Reading MGM after writing it returned wrong members count.	2.30.8000	2.31.5050
137.	RSS QP context	Fixed corruption of the RSS hash key given by the driver.	2.30.8000	2.31.5050
138.	10Gb/s QoS	Fixed QoS rate limit BW offset.	2.30.3200	2.31.5050
139.	ExtendedPortInfo MAD	Fixed FDR10 speed_en reporting.	2.30.8000	2.31.5050
140.	Management link	Fixed long management link com loss.	2.30.8000	2.31.5050
141.	PRM Query_Port Command	The command results reported both link types active at the same time.	2.30.3200	2.31.5050
142.	Link not raising	Fixed collision between forcing phy type and port sensing.	2.30.8000	2.31.5050
143.	Core clock reporting	Fixed a wrong core clock freq reporting in QUERY_HCA command.	2.30.3200	2.31.5050
144.	56GbE link issues	Fixeds occasional link failure when 56GbE is enabled	2.30.8000	2.31.5050

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
145.	RX calibration	Fixed max eye margins to be per protocol.	2.30.8000	2.31.5050
146.	VPI symbol errors	perfquery reported wrong error symbol on ConnectX [®] -3 VPI mode: IB, ETH.	2.30.8000	2.31.5050
147.	Symbol error on ConnectX-3 dual-port QDR with MC2207312-030 AOCs	On ConnectX-3 dual-port QDR and FDR/FDR10 switch setups, symbol errors may occur with MC2207312-030 AOCs.	2.30.8000	2.31.5050
148.	Symbol error on Falcon QDR against FDR switches with MC2207126-004 copper cables	Symbol errors occur on ConnectX-3 dual-port QDR connected to FDR switches with MC2207126-004 copper cables.	2.30.8000	2.31.5050
149.	PCIe correctable errors in speed change	When PCIe Gen3 is enabled, temporary correctable errors might occur when changing speed between PCIe Gen1 and PCIe Gen2.	2.10.0000	2.30.8000
150.	Incompatibility between Device managed Flow steering and NC-SI	Device managed Flow steering and NC-SI cannot be enabled simultaneously.	2.11.0500	2.30.8000
151.	40GbE is not supported in Auto-Sensing	Auto-Sensing is not supported with 40GbE connections in VPI cards	2.10.0000	2.30.8000
152.	PXE	PXE is currently not supported in 40GbE in VPI cards	2.10.0000	2.30.8000
153.		PXE is currently not supported in QSFP to SFP+ hybrid cable	2.10.0000	2.30.8000
154.	sense_port failure	Ethernet cards failed to work with MLNX_OFED unless the do_sense was disabled in the INI	RH6.4 driver	MLNX_OFED 2.0-3.0.0
155.	Link errors	BER of 10 ⁻¹¹ with 7M copper SFP+ 10GbE cable against Arista switch	2.30.3200	2.30.8000
156.	Linkup Failure vs SwitchX [®] -2 based switch	Port failed to link up in 10GbE if it was previously linked up in 40GbE vs. SwitchX [®] -2 based switch	2.30.3200	2.30.8000
157.	NC-SI 40GbE reporting	Added 40GbE reporting in get_link_status NC-SI command	2.30.3000	2.30.8000
158.	Packets drop in receive when DMFS enabled	Steering entries overlapping caused packets to drop in the receive due to wrong hash size calculation of QP hash folding.	2.30.3000	2.30.8000

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
159.	PCIe speed degrade	Occasionally, PCIe speed degraded during speed change test	2.30.3000	2.30.8000
160.	Link failure vs Cisco	Device failed to raise the link against Cisco b-22 Blade switch	2.30.3000	2.30.8000
161.	False RX drops indication	Fixed an issue causing RX drop counters to falsely increase when using MLNX_OFED 2.0.-3.0.0	2.30.3000	2.30.8000
162.	NC-SI wrong command info	Wrong port information provided in get link status command.	2.30.3000	2.30.8000
163.	Port interfaces stay down	After firmware upgrade, the device failed to raise port interface.	2.30.3000	2.30.8000
164.	SR-IOV guest failure	Guest VM failed to execute firmware commands operations and crashed.	2.30.3000	2.30.8000
165.	SR-IOV command timeouts	Guest MSIX vectors were not assigned properly.	2.30.3000	2.30.8000
166.	PCIe speed degrade/link down	Occasionally, the PCIe link experienced speed degrading or link falling in driver restart/reboot	2.30.3000	2.30.8000
167.	QP Sniffer issue	Fixed a hash fold issue for sniffer QPs	2.30.3000	2.30.8000
168.	Long link up time	A long link up time is experienced in the HCA connected to a 10GbE cable against the MSX1012 switch	2.30.3000	2.30.8000
169.	PCIe TS parity bit	In recovery.EQLZ state TSs used incorrect parity bit calculation	2.30.3000	2.30.8000
170.	PRM Init_port failure	Init port command may fail on a system with NC-SI	2.30.3000	2.30.8000
171.	PortInfo MAD link width support	Wrong link_width_support is reported occasionally in PORT_INFO MAD	2.30.3000	2.30.8000
172.	Wrong Virtual Function completer ID	In SR-IOV a wrong completer ID is added in some of the VF completion packets	2.30.3000	2.30.8000
173.	PCIe PML1 failures	Fixed occasional failures upon entering and exiting L1 state in PCIe Gen1 & 2 speeds	2.30.3000	2.30.8000
174.	ipmitool OOB commands	On rare occasions, ipmitool OOB commands failed upon send payload.	2.30.3000	2.30.8000
175.	Sideband communication	On rare occasions, after stress of power cycles, side- band communication might disconnect.	2.30.3000	2.30.8000
176.	Expansion ROM partition	Expansion ROM partition not found	2.30.3000	2.30.8000

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
177.	AC power cycle issue	In certain servers, AC power cycle may cause BMC connectivity loss.	2.30.3000	2.30.8000
178.	PXE issue	Occasionally, during DC cycle stress, failure occurred in PXE due to race condition	2.30.3000	2.30.8000
179.	SMBUS communication	SMBUS communication lost during AC/DC cycle	2.30.3000	2.30.8000
180.	MTU configuration	Temporal wrong MTU configuration during initialization may cause Serial over LAN disconnection	2.30.3000	2.30.8000
181.	IPMI connectivity	IPMI OOB communication lost during stress	2.30.3000	2.30.8000
182.	iperf stress test	Packet drops during iperf stress w/ different MSS	2.30.3000	2.30.8000
183.	Loss of connection to BMC upon firmware upgrade	Upgrading from firmware v2.10.3898 may cause loss of connection to BMC.	2.30.3000	2.30.8000
184.	Running rmmmod may cause unexpected behavior	Removing the mlx4_en driver using the "rmmmod" command, may cause unexpected behavior	2.30.3000	2.30.8000
185.	IPMI connectivity	On rare occasions, after stress of BMC cold reset, link failure might occur	2.30.3000	2.30.8000
186.	Flexboot 3.4.100-UEFI-4.0.410 BIOS menu	Enter the BIOS menu while using FlexBoot 3.4.100-UEFI-4.0.410 may cause the server to stop responding (hang)	2.30.3000	2.30.8000
187.	ROL failure	ROL failure after disassembling the driver on the SUT	2.30.3000	2.30.8000
188.	IPMI link failure	IPMI link failure after disabling the WoL or disassembling the driver	2.30.3000	2.30.8000
189.	WoL and RoL issues	WoL and RoL issues caused when the IPMI is disabled	2.30.3000	2.30.8000
190.	A link flapping issue	Alignment marker arrival can no longer drop the link.	2.30.3000	2.30.8000
191.	IPMI - SOL traffic performance improvement	SOL with multiple data streams occasionally hang	2.30.3000	2.30.8000
192.	RoCE	RoCE does not function properly after running "ethtool ethX"	2.30.3000	2.30.8000
193.	PCI link errors	PCI link errors false indication. Cleared errors during PCIe link retraining	2.30.3000	2.30.8000
194.	PCIe speed change	Fixed a false indication for incoming PCIe speed change request	2.30.3000	2.30.8000

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
195.	No NC-SI after PXE teardown	Management transport was not supported, after PXE driver teardown (unload). Fixed the flow to issue software reset, after the driver was stopped	2.30.3000	2.30.8000
196.	Ports PLL calibration issue	PLL calibration were affected by operating point configuration	2.30.3000	2.30.8000
197.	Cable Info MAD issue	Wrong cable info was received when using the MC2210411-SR4 module	2.30.3000	2.30.8000
198.	Revision ID	Updated the Revision ID of Node Info and General Info MADs to reflect HW Rev ID instead of 0	2.30.3000	2.30.8000
199.	Port Error counters reset	Port error counters were not cleared upon XAUI/SGMII link up	2.30.3000	2.30.8000
200.	RDMA retransmission on ConnectX®-3 Pro	RDMA retransmission failed in specific scenario of receiving NAKs on ConnectX®-3 Pro due to bad static configuration	2.30.3000	2.30.8000
201.	NMI on PCIe Gen2 server	Fixed a PCIe Gen2 firmware flow to prevent NMI during hibernation on PCIe Gen2 server	2.30.3000	2.30.8000
202.	Access to closed resources	Fixed a possible access to unmapped resource memory	2.30.3000	2.30.8000
203.	Slow link establishment in NC-SI	Slow link establishment in NC-SI caused due to slow cable reading in boot	2.30.3000	2.30.8000
204.	PCIe speed change	Occasionally, a failure occurred in speed when changing to Gen2	2.30.3000	2.30.8000
205.	PXE teardown issue	PXE might halt during teardown	2.30.3000	2.30.8000
206.	InfiniBand loopback	InfiniBand loopback was blocked during link negotiation on the same port	2.11.0500	2.30.8000
207.	Voltage scaling	Fixed process voltage scaling issue	2.30.3000	2.30.8000
208.	DMA address 0x0	Fixed a possible read access to DMA address 0x0	2.11.0500	2.30.8000
209.	cqe issue	Fixed miss cqe issue due to interrupt moderation	2.11.0500	2.30.8000
210.	Cable reading issue	Fixed a rare cable reading issue upon cable insertion	2.30.3000	2.30.8000
211.	LLR Vendor Specific MAD	LLR Dropped cell counter reported CSN error	2.11.0500	2.30.8000

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
212.	PXE boot failure	On rare occasions, PXE boot fails due to a firmware issue interfering with the PXE load from the card's flash. Upon failure, the following message is received: "Payload inaccessible - cannot continue"	2.30.3000	2.30.8000
213.	SR-IOV guest communication channel error	Under certain conditions, SR-IOV guest experienced request timeouts and got stuck.	2.11.0500	2.30.3200
214.	ETH KR Link Training to Data transition failure	A gap (LOS) between training sequences to data transmission causes peer partner to lose lock and fail ETH KR link.	2.30.3000	2.30.3200
215.	No NC-SI after PXE teardown	Management transport was not supported, after PXE driver teardown (unload). Fixed the flow to issue software reset, after the driver was stopped.	2.11.1250	2.30.3000
216.	Revision ID	Updated the Revision ID of Node Info and General Info MADs to reflect hardware Rev ID instead of 0	2.11.1250	2.30.3000
217.	Ports PLL calibration issue	PLL calibration were affected by operating point configuration	2.11.1250	2.30.3000
218.	Port Error counters reset	Port error counters were not cleared upon XAUI/SGMII link up	2.11.1250	2.30.3000
219.	Receiver SerDes tuning	Enhanced the receiver SerDes tuning for 10GE, to support specific 10GbE QSFP to SFP+ splitter cables	2.11.1250	2.30.3000
220.	Access to closed resources	Fixed a possible access to unmapped resource memory	2.11.1250	2.30.3000
221.	Slow link establishment in NC-SI	Slow link establishment in NC-SI caused due to slow cable reading in boot	2.11.1250	2.30.3000
222.	HP Flex Rack 10GbE link failure	HP Flex Rack ALOM failed to raise 10GbE link on Port1 against SwitchX® based InfiniBand switches	2.11.0500	2.30.3000
223.	PCIe speed change	Occasionally, a failure occurred in speed when changing to Gen2	2.11.1250	2.30.3000
224.	PXE teardown issue	PXE might halt during teardown	2.11.1250	2.30.3000
225.	InfiniBand loopback	InfiniBand loopback was blocked during link negotiation on the same port	2.11.1250	2.30.3000
226.	Voltage scaling	Fixed process voltage scaling issue	2.11.1250	2.30.3000
227.	DMA address 0x0	Fixed a possible read access to DMA address 0x0	2.11.1250	2.30.3000

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
228.	cqe issue	Fixed miss cqe issue due to interrupt moderation	2.11.1250	2.30.3000
229.	Cable reading issue	Fixed a rare cable reading issue upon cable insertion	2.11.1250	2.30.3000
230.	PXE boot failure	On rare occasions, PXE boot fails due to a firmware issue interfering with the PXE load from the card's flash. Upon failure, the following message is received: "Payload inaccessible - cannot continue"	2.11.1250	2.30.3000
231.	Momentarily packet drop on one port while the other port goes down	When the same VLAN is configured for both ports and one port goes down, the second port may suffer a momentarily packet drop	2.11.1250	2.30.3000
232.	Advanced Error Reporting	Fixes to Advanced Error Reporting according to the PTC (PCIe compliancy) Test failures	2.11.1250	2.30.3000
233.	10GbE link remains down after changing to a 40GbE link	When changing link connection between 40Gbe to 10Gbe, the port might stay down until the next driver restart.	2.11.1250	2.30.3000
234.	Race in PCIe L1 flow	The device may enter an L1 power state before completing an incoming configuration request if it arrived before the power state change ack for Root Complex	2.11.1250	2.30.3000
235.	Wrong link speed after several cable re-insertions	During repeated cable reinsertion, the link may raise in a lower speed than expected/	2.11.1250	2.30.3000
236.	Changing port protocol from ETH to IB requires driver restart	Bad configuration of <code>ib_protocol</code> when setting the port to InfiniBand after exchanging it from Ethernet/RoCE on the same port	2.11.1250	2.30.3000
237.	Rare NMI error on HP servers when using PXE	On rare occasions, an NMI error is received when stopping PXE boot in the middle of an action on HP servers	2.11.1250	2.30.3000
238.	Race in PCI configuration handling	The system was unresponsive when a race between PCI configuration cycles handling in hardware and firmware occurred.	2.11.1250	2.30.3000
239.	NMI event using the PXE CLI	On rare occasions, an NMI event occurred on HP ALOM while trying to using the PXE CLI	2.11.1250	2.30.3000
240.	RoCE breaks IPv6 traffic	IPv6 packets dropped while RoCE was enabled	2.11.1250	2.30.3000

Table 8 - Fixed Bugs List

	Issue	Description	Discovered in Release	Fixed in Release
241.	Link reset, HCA to SwitchX®	On rare link reset occasions, an HCA to SwitchX® link may remain in ARM state after SwitchX reboot, and not reach the Active state.	2.11.1250	2.30.3000
242.	SET_PORT.mac_table configuration	Fixed SET_PORT.mac_table configuration issue which caused minor packet loss on port A when working in bonding mode and closing port B.	2.11.1022	2.11.1250
243.	Insertion of L4 head rule	Fixed the insertion of L4 head rule in device managed flow steering.	2.11.1022	2.11.1250
244.	Wrong reporting of RSS context	Fixed wrong reporting of RSS context in QUERY_FW of RSS QP.	2.11.1022	2.11.1250
245.	NMI during PXE	Fixed an NMI scenario occurred during PXE boot.	2.11.1020	2.11.1022
246.	GPIO5 mis-configuration	Fixed firmware mis-configuration of GPIO5 which resulted in erroneous server thermal shutdown with previous revision B2/B3 ALOM cards and with previous 2.11.100x firmware versions.		2.11.1020
247.	Infrequent 10GbE or 1GbE link flapping after driver restart	Fixed 10GbE and 1GbE link flapping that occurred in-frequently, specifically once every 50 driver restarts and lasted few seconds.		2.11.1020
248.	Link reset, HCA to SwitchX®	On rare link reset occasions, an HCA to SwitchX® link may remain in ARM state after SwitchX reboot, and not reach the Active state.		2.11.1002
249.	4036 Switch link up time	Fixed the occasional link up time of the 4036 switch to up to 4 minutes.		2.11.1002
250.	ETH port stuck	Fixed ETH port getting stuck issue during link up causing firmware to halt.		2.10.5380
251.	PCI correctable errors	Fixed bad PCI reporting.		2.10.5380
252.	RoCE re-transmission	Fixed an issue that caused the message to not be re-transmitted from the beginning of the message but from PSN NAK.		2.10.5380
253.	Function Level Reset (FLR)	Fixed FLR in no-driver mode.		2.10.5380

5 Firmware Changes and New Feature History

Table 9 - Firmware Changes and New Feature History

Release	Description
2.40.5030	<ul style="list-style-type: none"> • Temperature thresholds: Added temperature thresholds high/low default for MAD sensing and NCSI/IPMI OEM commands. • MTU Header Size: Added a new field to "set port" command which notifies the firmware what is the user_mtu size. • ifconfig: Added a protection mechanism which ensures the firmware drops packets which are received in internal QPs and disables the WQE producer fetching. • Bug fixes: See Section 4, "Bug Fixes History," on page 12
2.40.5000	<ul style="list-style-type: none"> • Bug fixes: See Section 4, "Bug Fixes History," on page 12
2.36.5000	<ul style="list-style-type: none"> • Packet Steering: Enables steering packets to receive queues according to Ethertype matching (See PRM 2.1 for more information). • RX Arbiter: Adds support for additional rate values. • Performance counter for WQE fetch: Counters that count the number of repeated Send WQE cache lookups that resulted in a miss. • 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.
2.35.5100	<ul style="list-style-type: none"> • New performance and back-pressure counters command via PRM (For further information, please refer to the PRM) • Support for Multicast/Unicast sniffer rules (For further information, please refer to the PRM) • Support for VLAN in VLAN encapsulation (For further information, please refer to the PRM) • CQ creation offload by software • Support for rst2rts command • Invalidates a TLV during the firmware boot stage • A new counter for the <code>diag_rprt</code> PRM command to count packet drops due to no-receive buffer • Support for Ethernet TX lifetime cycle control (Head of Queue) • A new register (PPLR) that allows egress and external loopback control (For further information, please refer to the PRM) • A watchdog mechanism to track ingress traffic stalls to prevent flooding the network with Flow Control packets • Inspur LED scheme: A new LED scheme controlled by the INI which causes constant traffic LED indication even without traffic.

Table 9 - Firmware Changes and New Feature History

Release	Description
2.34.5000	<ul style="list-style-type: none"> • Added support for multiple RoCE modes (RoCE v1+v2) on the same port: RoCE mode is per connection now. • Added a new QP command <code>INIT2RTS_QP</code> to enhance QP connection readiness time. • Disabled FCS checks to support switches that replace FCS with Timestamp. • Added RX Port identification for direct rout packets. • Improved RDMA WRITE/SEND performance with retransmissions. • Enabled firmware burning/querying using the PRM <code>ACCESS_REG</code> command. • Added support for VAM. • Enabled bad cable EEPROM reporting to the driver. • Added support for Platform Level Data Model (PLDM) sideband protocol. • Added support for priority based A0-DMFS mode (For further information, please refer to the PRM). • Added support for Unicast/Multicast loopback disablement by the driver. (For further information, please refer to the PRM) • Removed the source IP from the hash calculation (For further information, please refer to the PRM) • Added support for Inline Receive mode up to 2KB.
2.33.5220	<ul style="list-style-type: none"> • FlexBoot: See Section 6, “FlexBoot Changes and New Features,” on page 34

Table 9 - Firmware Changes and New Feature History

Release	Description
2.33.5050	<ul style="list-style-type: none"> • Virtual QoS: <ul style="list-style-type: none"> • Bandwidth allocation support: Including maximum bandwidth and bandwidth share guarantee between VMs for InfiniBand and Ethernet. • Performance/Flow Control: <ul style="list-style-type: none"> • Increased inbound traffic buffer capacity when the PFC on all priorities is enabled. • Non-Volatile Device Configuration: <ul style="list-style-type: none"> • Added support for changing UAR BAR (PCI BAR 2) size. • Cables: <ul style="list-style-type: none"> • Added support for cable sub-power class for Mellanox MFA1A00-EXXX and SMFA1A00-CXXX EDR cables. • Device Managed Steering: <ul style="list-style-type: none"> • Improvements in attachment/detachment flows' rules in both A0-DMFS and DMFS modes. • RoCE Link Aggregation (LAG): <ul style="list-style-type: none"> • Added physical port forcing on specific QPs when virtual mapping is applied • Added support for dynamic enablement of LAG mode • NC-SI: <ul style="list-style-type: none"> • Added support for vendor specific command to report the ports' MAC addresses. • Link Speeds: <ul style="list-style-type: none"> • Enabled 100Mb ability exposure and its enablement via an INI parameter. • Added support for SFP+ with 1GbE when the adapter card is enabled in the EEPROM. • SideBand Management: <ul style="list-style-type: none"> • Optimized the SideBand connectivity loss during driver initialization to minimum. • SMBUS: <ul style="list-style-type: none"> • Added support for SMBUS ARP. • Enabled thermal reporting of TMP421 sensor in OCP cards. • RDMA: <ul style="list-style-type: none"> • RDMA Read retransmission optimizations to improve performance and ensure forward progress while packet drops occur. • Performance: <ul style="list-style-type: none"> • Improved data path WQE prefetch algorithm.

Table 9 - Firmware Changes and New Feature History

Release	Description
	<ul style="list-style-type: none"> • Ethtool Improvements: <ul style="list-style-type: none"> • Added support to query PTYS, PTOS registers through ACCESS_REG PRM command. • Non-volatile Random Access Memory (NVRAM): <ul style="list-style-type: none"> • Added support for CLP access to NVRAM • Packet Steering: <ul style="list-style-type: none"> • Added support for more than 22 QPs per MCG in DMFS. • Added support for high rate steering mode (a.k.a Simplified Steering) • PRM <ul style="list-style-type: none"> • Added support for reading current hardware mode through the QUERY_PORT PRM command • Added CSUM mode reporting in QUERY_DEV_CAP command • Added additional configuration options for UPDATE_QP command • Added support for 128 Byte stride for CQ/EQ • Enabled module EEPROM access using command I/F • Device Reset and Error recovery: <ul style="list-style-type: none"> • Reset Flow improvements and graceful handling of error caused by Virtual Functions • Performance: <ul style="list-style-type: none"> • RX performance optimization for single port cards • Promiscuous mode performance improvements • Secure Host: <ul style="list-style-type: none"> • Added support for Secure Host mode • Non-Volatile device configuration: <ul style="list-style-type: none"> • Added Port protocol configuration option. • GPIO: <ul style="list-style-type: none"> • Added support for GPIO swap • Signal Integrity (SI): <ul style="list-style-type: none"> • 40GbE SI improvements
(cont.)	<ul style="list-style-type: none"> • MAD: <ul style="list-style-type: none"> • Added support for Temp Sensing Vendor specific MAD. • NC-SI: <ul style="list-style-type: none"> • Added Temp Sensing NC-SI cmd. • Added support for AEN. • SMBUS: <ul style="list-style-type: none"> • Added new command to report firmware revision.

Table 9 - Firmware Changes and New Feature History

Release	Description
2.31.5050	<ul style="list-style-type: none"> • MAD: <ul style="list-style-type: none"> • Added support for GeneralInfo SMP MAD. • Updated capability mask in GeneralInfo SMP/GMP MAD • Added support for PortCountersVL MAD • Added support for PortSamplesControl/PortSamplesResults/PortSamplesExtended MADs • Added support query for additional MAC addresses per port (up to 7) through the QUERY_PORT command. • INI: <ul style="list-style-type: none"> • Added Exponential Backoff Timer support. It is enabled via the <code>rtm_ini</code> parameter. The default value is 0. • Flow Steering (DMFS): <ul style="list-style-type: none"> • Added VLAN steering to Device Managed Flow Steering (DMFS). • Non-Volatile configuration tool: <ul style="list-style-type: none"> • Added support for Non-Volatile configuration of TLVs to set device attributes: <ul style="list-style-type: none"> • Query and set of configurations is available through PRM ACCEES_REG command • PRM ACCESS_REG command is now also supported through the <code>tools_hcr</code> command interface • Added support for MTF <code>mlxconfig</code> tool • Management protocols: <ul style="list-style-type: none"> • Added IPv6 support for NC-SI and IPMI Pass-Through. • Added support for the same unicast MAC simultaneously for both IPMI and NC-SI • PCIe: <ul style="list-style-type: none"> • Added enhancements for receiver equalization in Gen3: <ul style="list-style-type: none"> • Enhancements are enabled by the INI. The default value is disabled. Please contact Mellanox support if required to enable it. • PCIe power optimizations for 8X/4X links. • Side Band protocols: <ul style="list-style-type: none"> • SMBUS optimizations • Physical and Virtual Functions reset flows: <ul style="list-style-type: none"> • Added new Physical and Virtual Functions reset flows support. • PXE support: <ul style="list-style-type: none"> • Added support for 64Bit BIOS mode.
2.31.5050 (cont.)	<ul style="list-style-type: none"> • PRM: <ul style="list-style-type: none"> • Added IEEE802.3 CL73 autoneg support to the QUERY_PORT command. • Added factory MAC address reporting to the Query_Port command. • Added support for reverting virtual MAC configuration per port and restoring to factory MAC through MOD_STAT_CFG command. • Added support for inline TLV read through MOD_STAT_CFG command. • Added current MTU reporting to the QUERY_PORT command.

Table 9 - Firmware Changes and New Feature History

Release	Description
2.30.8000	<ul style="list-style-type: none"> • DMFS and GRE steering: Rule insertion adjustments • Removed DIF support from reported capabilities in QUERY_DEV_CAP PRM command • Flow control by DSCP priority for IPv4 • DMFS improvements: Insertion scheme enforcement and block loopback for InfiniBand • Added I2C resiliency support • Support for NC-SI over MCTP over SMBus • Added a flash access interface for persistent (non-volatile) configuration support • Added port BW arbitration configuration through the CONFIG_DEV command • Added IP-in-IP TCP checksum offload support • pci Express compliancy Tx and Rx adjustments • 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. • Added support for 40GbE in WoL and pre-OS driver modes To enable this, add/change the following flags in the INI file in the IB and HCA tabs respectively: <ul style="list-style-type: none"> • <code>restrict_max_eth_standby_speed = NO_RESTRICTION</code> • <code>slow_clock_enable = 0</code> • Bug fixes, see Section 4, "Bug Fixes History," on page 12
2.30.3200	<ul style="list-style-type: none"> • Bug fixes, see Section 4, "Bug Fixes History," on page 12

6 FlexBoot Changes and New Features

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

Table 10 - FlexBoot Changes and New Feature

Version	Description
Rev 3.4.752	<ul style="list-style-type: none"> • Added support for Secure Firmware Update • Enabled booting Grub 2.02 over IB • Disabled SOL by default
Rev 3.4.746	<ul style="list-style-type: none"> • Added support for the following SHELL CLI commands: <ul style="list-style-type: none"> • Non-volatile option storage commands • SAN boot commands • Menu commands • Login command • Sync command • DNS resolving command • Time commands • Image crypto digest commands • Loopback testing commands • VLAN commands • PXE commands • Reboot command • For further information, please refer to: http://ipxe.org/cmd • iSCSI re-imaging: enables the user to install a new image on active iSCSI target. • Removed link status line printout at boot time. • Deprecated the option "rom enable" bit. • Enabled interrupt support. • When Network Boot Program (NBP) uses UNDI, the user can configure the awaiting time (up to 30 seconds) that is needed to raise a link. • Set default banner timeout to 4. • Enabled UDP interface usage after UNDI shutdown. • Fixed a BIOS issue in hybrid BIOSes which resulted in legacy driver load failure when the BIOS loaded legacy driver without closing the UEFI driver. • Fixed an issues causing the PXE to boot first regardless of the boot priority if the client received "PXE boot menu" when contacted the DHCP. • Synced the source with iPXE (upstream sync).

Table 10 - FlexBoot Changes and New Feature

Version	Description
Rev 3.4.718	<ul style="list-style-type: none"> • Added IPv6 support (Beta level) • Removed support for the following SHELL CLI commands: <ul style="list-style-type: none"> • Non-volatile option storage commands • SAN boot commands • Menu commands • Login command • Sync command • DNS resolving command • Time commands • Image crypto digest commands • Loopback testing commands • VLAN commands • PXE commands • Reboot command <p>For further information, please refer to: http://ipxe.org/cmd</p>
Rev 3.4.648	<ul style="list-style-type: none"> • Added support for .mrom images larger than 128kB • Added boot over IB with non-default PKey for ConnectX®-3, ConnectX®-3 Pro cards • Synced the source with iPXE (upstream sync) • Moved to flat real mode when calling INT 1a,b101 to avoid BIOSes issues • Fixed HTTP boot over IPoIB
Rev 3.4.521	<ul style="list-style-type: none"> • Added iSCSI CHAP and mutual CHAP configuration • Added the GRH size when allocating receive buffer for IPoIB • Updated VLAN netdevice's settings with all the trunk's iSCSI required settings • Updated the port event handling process • Enabled console output in Debug mode • Disabled the serial output • Disabled the banner in BEV execution • Disabled function 0x04 (in int21) when serial console is disabled • Preserved COM port settings • Fixed HTTP download over IPoIB • Fixed completion with error handling process
Rev 3.4.467	<ul style="list-style-type: none"> • Fixed an issue preventing TFTP filename with absolute path functionality

Table 10 - FlexBoot Changes and New Feature

Version	Description
Rev 3.4.460	<ul style="list-style-type: none"> • Boot Menu support: Added new FlexBoot GUI. The device can now be configured in the POST stage. • Non volatile memory read/write support • Configurable URI boot retry and delay between retries • Configurable iSCSI settings using DHCP/NVM • Added new interface in order to update the registered devices on the PXE stage • Enabled ConnectX Ethernet adapter cards family to work with interrupts • Enabled PXE to work in promiscuous VLAN mode (configurable through the INI) • Synced version with ipxe.org: Now the latest code in iPXE is used • Added boot priority capability: iSCSI vs PXE and fallback incase one fails • Updated the Proxy DHCP request method for non-existing option 54. ProxyDHCP request is sent to port 67 with broadcast IP address if the server identifier in option 54 is zero. Packets with source port different than BOOTPS_PORT and PXE_PORT are filtered by the PROXY • SHELL CLI is currently supported on ConnectX-3 and ConnectX-3 Pro adapter cards only • Both the GUID and the MAC are printed on the screen when the port link layer is set as InfiniBand • PROXYDHCP and PXEBS settings are saved under netdevice settings • rootpath/filename/nextserver are now fetched from the netdevice settings • The cached DHCP packet are received only if working with the same net device. When pxelinux.0 receives the cached DHCP packet from the UNDI API, it constructs a new (fake) packet for the current net device. If the process is stopped and then restarted and booted from the next boot device which serves as the second port in the HCA, a new (fake) DHCP packet is not constructed. The previous packet which includes all the information of the first port (IP, MAC, Net-mask, etc...) is used. If an old (fake) DHCP packet is discovered, its chaddr is compared to the chaddr in the pxe_netdev, if not similar, a new (fake) DHCP packet is created. • PXE shutdown is called if int22 with function 0x000C is called. • The server's IP address in DHCP server replies is now checked before checking the reply type. This will ignore NACK replies from servers which already were ignored by the client. In case of 2 DHCP servers in the same subnet, the client will eventually choose one of them, by sending the DHCP REQUEST with 'DHCP Server Identifier' (option 54) filled with the requested server's IP address. • Changed DHCP discover timeouts to comply with PXE spec

Table 10 - 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 10 - 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 64-bit PCI BAR support • Added an option for running PXE with promiscuous VLAN • Added support for HII iSCSI configuration • 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 iSCSI initiator's name retrieval from HII issue • 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 TLV with length 0 issue • 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
Rev 3.4.151	<ul style="list-style-type: none"> • 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
Rev 3.4.146	<ul style="list-style-type: none"> • Fixed memory corruption issues • Modified TLV flash access • Added additional WQ
Rev 3.4.142	<ul style="list-style-type: none"> • Enabled firmware to handle the link state with the Subnet Manager • Updated the DHCP class code to NONE • Added flash access capability for reading software-to-software configurations • Enabled DHCP validation of MAC address and XID for a unique tuple • Improved randomness algorithm for DHCP XID
Rev 3.4.112	<ul style="list-style-type: none"> • Broadcast responses for firewall support • Enabled request broadcast responses from DHCP server to support firewall.

Table 10 - FlexBoot Changes and New Feature

Version	Description
Rev 3.4.100	<ul style="list-style-type: none"> • OCSD activation initiation change • Moved the OCSD activation initiation from the FlexBoot to the CLP code. This enables the OCSD • activation to no longer be dependent on the FlexBoot being enabled in the servers's BIOS configuration. • Messages' improvement • Made the FlexBoot on-screen notification messages more informative and user friendly. • FlexBoot and CLP merge improvement • Improved the process of merging the FlexBoot and CLP codes together. • PXE and UFI merge capability • Added the ability to merge the PXE image with a UFI image. • Supported servers • Added FlexBoot support capabilities to several new non-HP servers. • Use of newer iPXE version • Moved to use a newer iPXE version as the basis for the Flexboot release. • Fixed "no more network devices" issues during Flexboot.

6.1 Flexboot Known Issues

The following is a list of general limitations and known issues of the various components of this FlexBoot release.

Table 11 - Known Issues

Internal Ref.	Description
673114/821899	Description: FlexBoot banner might not be shown in some BIOSes.
	WA: N/A
	Keywords: BIOS
572684	Description: FlexBoot Boot Menu will not be visible in serial output.
	WA: N/A
	Keywords: User Interface
792432	Description: Booting PXE using Grub2.X over HP G9/G8 servers results in system hang.
	WA: N/A
	Keywords: PXE boot, Grub2.X, HP G9/G8

7 CLP Changes and New Features

Table 12 - CLP Changes and New Feature

Version	Description
Rev 8025	<ul style="list-style-type: none"> Reverted the flash configuration from v8033 to v8025 to maintain ROM code as small as possible. Reverting the code, does not affect the added flash configuration read/write capability added in v8033. Fixed the following CLP v8024 issue: Occasionally, an uninitialized variable on systems supporting OCSD causes system reboot.
Rev 8024	<ul style="list-style-type: none"> 644161-B21/B22 NFF Mezz 40GbE enabling CLP fix specifically related to 40GbE enabling on NFF Mezz card 644161-B21/B22.
Rev 8023	<ul style="list-style-type: none"> CLP version display Changed the CLP version format to be autonomic and be presented separately from the FlexBoot PXE version. 40GbE addition to the FDR NFF Mezz card Added CLP configuration support of 40GbE mode to the FDR NFF Mezz card (this change is not related to the Watatic ALOM card). OCSD activation initiation change Moved the OCSD activation initiation from the FlexBoot to the CLP code. This enables the OCSD activation to no longer be dependent on the FlexBoot being enabled in the servers' BIOS configuration.

































