# Mellanox ConnectX-3 Firmware Release Notes

Rev 2.35.5100

www.mellanox.com

Mellanox Technologies
350 Oakmead Parkway Suite 100
Sunnyvale, CA 94085
U.S.A.
www.mellanox.com
Tel: (408) 970-3400
Fax: (408) 970-3403

# Table of Contents

# List of Tables

# Release Update History

*Table 1 - Release Update History*

| Release | Date | Description |
|---|---|---|
| 2.35.5100 | November 25th, 2015 | Added the following OPNs:<br>• MCX341A-XCG<br>• MCX342A-XCC |
| | October 25th, 2015 | Added Known Issue #25 to Section 3, "Known Issues", on page 21 |
| | October 06th, 2015 | Updated Section 4, "Bug Fixes History", on page 24. added SR-IOV security bug fix. |
| | September 10th, 2015 | Initial version |

# 1 Overview

These are the release notes for the ConnectX®-3 adapters firmware Rev 2.35.5100. This firmware supports the following protocols:

- InfiniBand – SDR, DDR, QDR, FDR10, FDR
- Ethernet - 1GigE, 10GigE, 40GigE and 56GigE[1]
- 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 (Sheet 1 of 2)*

| Device Part Number | PSID | Device Name |
|---|---|---|
| MCX311A-XCAT | MT_1170110023 | ConnectX®-3 EN network interface card, 10GbE, single-port SFP+, PCIe3.0 x4 8GT/s, tall bracket, RoHS R6 |
| MCX312A-XCBT | MT_1080110023 | ConnectX®-3 EN network interface card, 10GbE, dual-port SFP+, PCIe3.0 x8 8GT/s, tall bracket, RoHS R6 |
|  | MT_1080120023 |  |
| MCX312B-XCBT | MT_1200110023 | ConnectX®-3 EN network interface card, 10GbE, dual-port SFP+, PCIe3.0 x8 8GT/s, tall bracket, RoHS R6 |
|  | MT_1200210023 |  |
| MCX313A-BCBT | MT_1060110023 | ConnectX®-3 EN network interface card, 40/56GbE, single-port QSFP, PCIe3.0 x8 8GT/s, tall bracket, RoHS R6 |
|  | MT_1060140023 |  |
| MCX314A-BCBT | MT_1090110023 | ConnectX®-3 EN network interface card, 40/56GbE, dual-port QSFP, PCIe3.0 x8 8GT/s, tall bracket, RoHS R6 |
| MCX353A-FCBT | MT_1100120019 | ConnectX®-3 VPI adapter card, single-port QSFP, FDR IB (56Gb/s) and 40/56GbE, PCIe3.0 x8 8GT/s, tall bracket, RoHS R6 |
| MCX353A-QCBT | MT_1060110018 | ConnectX®-3 VPI adapter card; single-port QSFP; QDR IB (40Gb/s) and 10GigE; PCIe3.0 x8 8GT/s; RoHS R6 |
| MCX353A-TCBT | MT_1100110028 | ConnectX®-3 VPI adapter card, single-port QSFP, FDR10 IB (40Gb/s) and 10GbE, PCIe3.0 x8 8GT/s, tall bracket, RoHS R6 |
| MCX354A-FCBT | MT_1090110019 | ConnectX®-3 VPI adapter card, dual-port QSFP, FDR IB (56Gb/s) and 40/56GbE, PCIe3.0 x8 8GT/s, tall bracket, RoHS R6 |
|  | MT_1090120019 |  |

---

1. 56 GbE is a Mellanox propriety link speed and can be achieved while connecting a Mellanox adapter cards to Mellanox SX10XX switch series or connecting a Mellanox adapter card to another Mellanox adapter card.

*Table 2 - Supported PSIDs (Sheet 2 of 2)*

| Device Part Number | PSID | Device Name |
|---|---|---|
| MCX354A-QCBT | MT_1090110018 | ConnectX®-3 VPI adapter card, dual-port QSFP, QDR IB (40Gb/s) and 10GbE, PCIe3.0 x8 8GT/s, tall bracket, RoHS R6 |
| MCX354A-TCBT | MT_1090110028 | ConnectX®-3 VPI adapter card, dual-port QSFP, QDR IB (40Gb/s) and 10GbE, PCIe3.0 x8 8GT/s, tall bracket, RoHS R6 |
| MCX342A-XCCN | MT_1680110023 | ConnectX®-3 EN network interface card for OCP, 10GbE, dual-port SFP+, PCIe3.0 x8, IPMI disabled, no bracket, RoHS R6 |
| MCX341A-XCGN | MT_1270114023 | ConnectX®-3 EN network interface card for OCP, 10GbE, single-port SFP+, PCIe3.0 x8, IPMI and NC-SI support, no bracket, RoHS R6 |

## 1.2     Supported Cables and Modules

Please refer to the LinkX™ Cables and Transceivers web page
(www.mellanox.com -> Products -> Cables and Transceivers) for the list of supported cables.

### 1.2.1    Tested Cables and Modules

*Table 3 - Tested Cables and Modules  (Sheet 1 of 9)*

| Speed | OPN # | Description | Vendor |
|---|---|---|---|
| NA | MAM1Q00A-QSA | MELLANOX QSFP TO SFP+ ADAPTER | MELLANOX |
| 1 GB/S | 453151-B21 | HP BLc VC 1Gb SX SFP Opt Kit | HP |
| 1 GB/S | 453154-B21 | HP BLc VC 1Gb RJ-45 SFP Opt Kit | HP |
| 1 GB/S | MC3208011-SX | MELLANOX OPTICAL MODULE ETH 1GBE 1GB/S SFP LC-LC SX 850NM UP TO 500M | MELLANOX |
| 1 GB/S | MC3208411-T | MELLANOX MODULE ETH 1GBE 1GB/S SFP BASE-T UP TO 100M | MELLANOX |
| 10GB/S | CAB-SFP-SFP-1M | passive copper cable,  SFP+, 10 Gb/s,  1m | Arista |
| 10GB/S | CAB-SFP-SFP-2M | passive copper cable,  SFP+, 10 Gb/s,  2m | Arista |
| 10GB/S | CAB-SFP-SFP-3M | passive copper cable,  SFP+, 10 Gb/s,  3m | Arista |
| 10GB/S | CAB-SFP-SFP-5M | passive copper cable,  SFP+, 10 Gb/s,  5m | Arista |
| 10GB/S | XDL-TWX0101 | Brocade passive copper cable, SFP+, 10 Gb/s, 1m | Brocade |
| 10GB/S | XDL-TWX0301 | Brocade passive copper cable, SFP+, 10 Gb/s, 3m | Brocade |
| 10GB/S | XDL-TWX0501 | Brocade passive copper cable, SFP+, 10 Gb/s, 5m | Brocade |

**Table 3 - Tested Cables and Modules  (Sheet 2 of 9)**

| Speed | OPN # | Description | Vendor |
|---|---|---|---|
| 10GB/S | Cisco SFP-H10GB-CU1M | Cisco SFP+ cable | Cisco |
| 10GB/S | Cisco SFP-H10GB-CU3M | Cisco SFP+ cable | Cisco |
| 10GB/S | Cisco SFP-H10GB-CU5M | Cisco SFP+ cable | Cisco |
| 10GB/S | MC2309124-007 | QSFP-4SFP10G | Cisco |
| 10GB/S | SFP-10G-SR | CISCO 10GBASE-SR SFP Module | Cisco |
| 10GB/S | MC2309124-007 | QSFP-4SFP10G | Cisco |
| 10GB/S | 0NWGTV | SFP+ to S FP+ copper cable 1M | Dell |
| 10GB/S | 0C4D08 | SFP+ to  SFP+ copper cable 1M | Dell |
| 10GB/S | 0V250M | SFP+ to  SFP+ copper cable 1M | Dell |
| 10GB/S | 0NMMT9 | SFP+ to  SFP+ copper cable 1M | Dell |
| 10GB/S | 053HVN | SFP+ to  SFP+ copper cable 3M | Dell |
| 10GB/S | 05CWK6 | SFP+ to  SFP+ copper cable 3M | Dell |
| 10GB/S | 00F1VT9 | SFP+ to  SFP+ copper cable 3M | Dell |
| 10GB/S | 00358VV | SFP+ to  SFP+ copper cable 5M | Dell |
| 10GB/S | 05CN56 | SFP+ to  SFP+ copper cable 5M | Dell |
| 10GB/S | 0V492M | SFP+ to  SFP+ copper cable 5M | Dell |
| 10GB/S | 0W25W9 | SFP+ to  SFP+ copper cable 5M | Dell |
| 10GB/S | 0J90VN | 40GbE QSFP+ to QSFP+ copper cable 5M | Dell |
| 10GB/S | TCPM2 | QSFP+ to 4xSFP+ copper cable 1M | Dell |
| 10GB/S | 27GG5 | QSFP+ to 4xSFP+ copper cable 3M | Dell |
| 10GB/S | P8T4W | QSFP+ to 4xSFP+ copper cable 5m | Dell |
| 10GB/S | 0WTRD1 | Dell 10Gb SR SFP+ Opt | Dell |
| 10GB/S | C4D08 | Force 10passive copper cable, SFP+, 10 Gb/s, 1m | Force10 |
| 10GB/S | 53HVN | Force 10passive copper cable, SFP+, 10 Gb/s, 3m | Force10 |
| 10GB/S | 5CN56 | Force 10passive copper cable, SFP+, 10 Gb/s, 5m | Force10 |
| 10GB/S | J9281B | HP X242 10G SFP+ to SFP+ 1m Direct Attach Copper Cable | HP |
| 10GB/S | J9283B | HP X242 10G SFP+ SFP+ 3m DAC Cable | HP |
| 10GB/S | J9285B | HP X242 10G SFP+ SFP+ 7m DAC Cable | HP |

*Table 3 - Tested Cables and Modules  (Sheet 3 of 9)*

| Speed | OPN # | Description | Vendor |
|-------|-------|-------------|--------|
| 10GB/S | JD096B | HP X240 10G SFP+ SFP+ 1.2m DAC Cable | HP |
| 10GB/S | JD095B | HP X240 10G SFP+ SFP+ 0.65m DAC Cable | HP |
| 10GB/S | JD097B | HP X240 10G SFP+ SFP+ 3m DAD Cable | HP |
| 10GB/S | JD096C | HP X240 10G SFP+ SFP+ 1.2m DAC Cable | HP |
| 10GB/S | JD095C | HP X240 10G SFP+ SFP+ 0.65m DAC Cable | HP |
| 10GB/S | JD097C | HP X240 10G SFP+ SFP+ 3m DAD Cable | HP |
| 10GB/S | 487649-B21 | HP BLc SFP+ .5m 10GbE Copper Cable | HP |
| 10GB/S | 487652-B21 | HP BLc SFP+ 1m 10GbE Copper Cable | HP |
| 10GB/S | 487655-B21 | HP BLc SFP+ 3m 10GbE Copper Cable | HP |
| 10GB/S | 537963-B21 | HP BLc SFP+ 5m 10GbE Copper Cable | HP |
| 10GB/S | 487658-B21 | HP BLc SFP+ 7m 10GbE Copper Cable | HP |
| 10GB/S | AP784A | HP 3m C-series Passive Copper SFP+ Cable | HP |
| 10GB/S | AP785A | HP 5m C-series Passive Copper SFP+ Cable | HP |
| 10GB/S | AP818A | HP 1m B-series Active Copper SFP+ Cable | HP |
| 10GB/S | AP819A | HP 3m B-series Active Copper SFP+ Cable | HP |
| 10GB/S | 455883-B21 | HP BLc 10Gb SR SFP+ Opt | HP |
| 10GB/S | 455886-B21 | HP BLc 10Gb LR SFP+ Opt | HP |
| 10GB/S | J9150A | HP X132 10G SFP+ LC SR Transceiver | HP |
| 10GB/S | J9151A | HP X132 10G SFP+ LC LR Transceiver | HP |
| 10GB/S | AJ839A | HP 50m Multi-mode OM3 LC/LC FC Cable | HP |
| 10GB/S | AJ838A | HP 30m Multi-mode OM3 LC/LC FC Cable | HP |
| 10GB/S | AJ837A | HP 15m Multi-mode OM3 LC/LC FC Cable | HP |
| 10GB/S | AJ836A | HP  5m Multi-mode OM3 LC/LC FC Cable | HP |
| 10GB/S | AJ834A | HP  1m Multi-mode OM3 LC/LC FC Cable | HP |
| 10GB/S | AJ833A | HP  0.5m Multi-mode OM3 LC/LC FC Cable | HP |
| 10GB/S | JG329A | HP X240 40G QSFP+ to 4x10G SFP+ 1m | HP |
| 10GB/S | JG330A | HP X240 40G QSFP+ to 4x10G SFP+ 3m | HP |
| 10GB/S | JG331A | HP X240 40G QSFP+ to 4x10G SFP+ 5m | HP |
| 10GB/S | JD095C | HP X240 10G SFP+ SFP+ 0.65m DAC Cable | HP |
| 10GB/S | 90Y9425-N28500A | IBM-Amphenol SFP+ 1m | IBM |

*Table 3 - Tested Cables and Modules  (Sheet 4 of 9)*

| Speed | OPN # | Description | Vendor |
|-------|-------|-------------|--------|
| 10GB/S | 46K6184-L36836B | IBM-Amphenol SFP+ 5m | IBM |
| 10GB/S | 46K6183-L36836B | IBM-Amphenol SFP+ 3m | IBM |
| 10GB/S | 44X1371-N31295E | IBM-Amphenol SFP+ 7m | IBM |
| 10GB/S | 95Y1634-N31295E | IBM SFP+ to  SFP+ copper cable 5M | IBM |
| 10GB/S | 44x1368-N31295E | IBM SFP+ to  SFP+ copper cable 0.5M | IBM |
| 10GB/S | 46K6182-L36836B | IBM SFP+ to  SFP+ copper cable 1M | IBM |
| 10GB/S | 46K6183-L36836B | IBM SFP+ to  SFP+ copper cable 3M | IBM |
| 10GB/S | 46K6184-L36836B | IBM SFP+ to  SFP+ copper cable 5M | IBM |
| 10GB/S | QFX-SFP-DAC-1M | SFP+ 10 Gigabit Ethernet Direct Attach Copper (twinax copper cable) 1 m | Juniper |
| 10GB/S | QFX-SFP-DAC-3M | SFP+ 10 Gigabit Ethernet Direct Attach Copper (twinax copper cable) 3 m | Juniper |
| 10GB/S | QFX-SFP-DAC-5M | SFP+ 10 Gigabit Ethernet Direct Attach Copper (twinax copper cable) 5  m | Juniper |
| 10GB/S | 740-021308 | Juniper 10GE  SFP+ module | Juniper |
| 10GB/S | MC2309124-004 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 4M | MELLANOX |
| 10GB/S | MC2309124-005 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 5M | MELLANOX |
| 10GB/S | MC2309124-006 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 6M | MELLANOX |
| 10GB/S | MC2309124-007 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 7M | MELLANOX |
| 10GB/S | MC2309130-001 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 1M | MELLANOX |
| 10GB/S | MC2309130-002 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 2M | MELLANOX |
| 10GB/S | MC2309130-003 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 3M | MELLANOX |
| 10GB/S | MC2309130-00A | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S QSFP TO SFP+ 0.5M | MELLANOX |
| 10GB/S | MC2609125-004 | MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 4M | MELLANOX |

**Table 3 - Tested Cables and Modules  (Sheet 5 of 9)**

| Speed | OPN # | Description | Vendor |
|---|---|---|---|
| 10GB/S | MC2609125-005 | MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 5M | MELLANOX |
| 10GB/S | MC2609130-001 | MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 1M | MELLANOX |
| 10GB/S | MC2609130-002 | MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 2M | MELLANOX |
| 10GB/S | MC2609130-003 | MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 3M | MELLANOX |
| 10GB/S | MC2609130-0A1 | MELLANOX PASSIVE COPPER HYBRID CABLE ETH 40GBE TO 4X10GBE QSFP TO 4X SFP+ 1.5M | MELLANOX |
| 10GB/S | MC3309124-004 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 4M | MELLANOX |
| 10GB/S | MC3309124-005 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 5M | MELLANOX |
| 10GB/S | MC3309124-006 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 6M | MELLANOX |
| 10GB/S | MC3309124-007 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 7M | MELLANOX |
| 10GB/S | MC3309124-0A3 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 3.5M | MELLANOX |
| 10GB/S | MC3309130-001 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 1M | MELLANOX |
| 10GB/S | MC3309130-002 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 2M | MELLANOX |
| 10GB/S | MC3309130-003 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 3M | MELLANOX |
| 10GB/S | MC3309130-00A | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 0.5M | MELLANOX |
| 10GB/S | MC3309130-0A1 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 1.5M | MELLANOX |
| 10GB/S | MC3309130-0A2 | MELLANOX PASSIVE COPPER CABLE ETH 10GBE 10GB/S SFP+ 2.5M | MELLANOX |
| 10GB/S | MFM1T02A-LR | Mellanox optical module, 1310nm, LR up to 10km | MELLANOX |

*Table 3 - Tested Cables and Modules  (Sheet 6 of 9)*

| Speed | OPN # | Description | Vendor |
|---|---|---|---|
| 10GB/S | MFM1T02A-SR | Mellanox optical module, 850nm, SR up to 300m | MELLANOX |
| 40GB/S | 40GbE QSFP+ to QSFP | QSFP+ copper cable 3M | Arista |
| 40GB/S | 40GbE QSFP+ to QSFP | QSFP+ copper cable 5M | Arista |
| 40GB/S | QSFP-H40G-CU1M | Cisco QSFP  40GbE cable | Cisco |
| 40GB/S | QSFP-H40G-CU3M | Cisco QSFP  40GbE cable | Cisco |
| 40GB/S | QSFP-H40G-CU5M | Cisco QSFP  40GbE cable | Cisco |
| 40GB/S | QSFP-40G-SR4 | CISCO 40G QSFP Module | Cisco |
| 40GB/S | 05NP8R | 40GbE QSFP+ to QSFP+ copper cable 1M | Dell |
| 40GB/S | 00FC6KV | 40GbE QSFP+ to QSFP+ copper cable 3M | Dell |
| 40GB/S | 0RF2MY | Dell 40GbB QSFP module | Dell |
| 40GB/S | 10093084-200AHFLF | FCI QSFP 0.75m 40GbE cable | FCI |
| 40GB/S | 10093084-2005HFLF | FCI QSFP 0.5m 40GbE cable | FCI |
| 40GB/S | 10093084-2010HFLF | FCI QSFP 1m 40GbE cable | FCI |
| 40GB/S | NWGTV | Force 10passive copper cable,  QSFP, 40 Gb/s, 1m | Force10 |
| 40GB/S | V492M | Force 10passive copper cable,  QSFP, 40 Gb/s, 5m | Force10 |
| 40GB/S | GP-QSFP-40GE-1SR | Force10 - Dell 40GbB QSFP module | Force10 |
| 40GB/s | JG325B | HP X140 40G QSFP+ MPO SR4 Transceiver | HP |
| 40GB/S | JG325A | HP X140 40G QSFP+ MPO SR4 Transceiver | HP |
| 40GB/S | JG326A | HP X240 40G QSFP+ QSFP+ 1m DAC Cable | HP |
| 40GB/S | JG327A | HP X240 40G QSFP+ QSFP+ 3m DAC Cable | HP |
| 40GB/S | JG328A | HP X240 40G QSFP+ QSFP+ 5m DAC Cable | HP |
| 40GB/S | 00D5802-N13445C | IBM 40GbE QSFP+ to QSFP+ copper cable 1M | IBM |
| 40GB/S | BN-QS-QS-CBL-3M | IBM 40GbE QSFP+ to QSFP+ copper cable 3M | IBM |
| 40GB/S | BN-QS-QS-CBL-5M | IBM 40GbE QSFP+ to QSFP+ copper cable 5M | IBM |
| 40GB/S | MC2210126-004 | MELLANOX PASSIVE COPPER CABLE ETH 40GBE 40GB/S QSFP 4M | MELLANOX |
| 40GB/S | MC2210126-005 | MELLANOX PASSIVE COPPER CABLE ETH 40GBE 40GB/S QSFP 5M | MELLANOX |

**Table 3 - Tested Cables and Modules  (Sheet 7 of 9)**

| Speed | OPN # | Description | Vendor |
|---|---|---|---|
| 40GB/S | MC2210128-003 | MELLANOX PASSIVE COPPER CABLE ETH 40GBE 40GB/S QSFP 3M | MELLANOX |
| 40GB/S | MC2210130-001 | MELLANOX PASSIVE COPPER CABLE ETH 40GBE 40GB/S QSFP 1M | MELLANOX |
| 40GB/S | MC2210130-002 | MELLANOX PASSIVE COPPER CABLE ETH 40GBE 40GB/S QSFP 2M | MELLANOX |
| 40GB/S | MC2210310-XXX | MELLANOX ACTIVE FIBER CABLE ETH 40GBE 40GB/S QSFP from 3M up to 100M | MELLANOX |
| 40GB/S | MC2210411-SR4L | MELLANOX OPTICAL MODULE 40GB/S QSFP MPO 850NM UP TO 30M | MELLANOX |
| DDR | MC1204128-001 | MELLANOX PASSIVE COPPER HYBRID CABLE IB DDR 20GB/S QSFP TO CX4 1M | MELLANOX |
| DDR | MC1204128-003 | MELLANOX PASSIVE COPPER HYBRID CABLE IB DDR 20GB/S QSFP TO CX4 3M | MELLANOX |
| DDR | MC1204128-005 | MELLANOX PASSIVE COPPER HYBRID CABLE IB DDR 20GB/S QSFP TO CX4 5M | MELLANOX |
| DDR | MC1204130-002 | MELLANOX PASSIVE COPPER HYBRID CABLE IB DDR 20GB/S QSFP TO CX4 2M | MELLANOX |
| FDR | 038-004-066-01 | EMC FDR QSFP+ to QSFP+ copper cable 2M | EMC FDR QSFP |
| FDR | 038-004-067-01 | EMC FDR QSFP+ to QSFP+ copper cable 3M | EMC FDR QSFP |
| FDR | 038-900-027-01 | EMC FDR QSFP+ to QSFP+ copper cable 5M | EMC FDR QSFP |
| FDR | 038-900-030-01 | EMC FDR QSFP+ to QSFP+ copper cable 8M | EMC FDR QSFP |
| FDR | 038-004-236-01 | FDR QSFP+ to QSFP+ copper cable 0.5m | EMC FDR QSFP |
| FDR | 038-004-065-01 | EMC FDR QSFP+ to QSFP+ copper cable 1M | EMC FDR QSFP |
| FDR | 038-004-069-01 | EMC FDR QSFP+ to QSFP+ copper cable 5M | EMC FDR QSFP |
| FDR | MC2207126-004 | MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 4M | MELLANOX |
| FDR | MC2207128-003 | MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 3M | MELLANOX |
| FDR | MC2207128-0A2 | MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 2.5M | MELLANOX |

*Table 3 - Tested Cables and Modules  (Sheet 8 of 9)*

| Speed | OPN # | Description | Vendor |
|-------|-------|-------------|--------|
| FDR | MC2207130-001 | MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 1M | MELLANOX |
| FDR | MC2207130-002 | MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 2M | MELLANOX |
| FDR | MC2207130-00A | MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 0.5M | MELLANOX |
| FDR | MC2207130-0A1 | MELLANOX PASSIVE COPPER CABLE VPI UP TO 56GB/S QSFP 1.5M | MELLANOX |
| FDR | MC2207310-XXX-E | MELLANOX ACTIVE FIBER CABLE VPI UP TO 56GB/S QSFP from 3M up to 100M | MELLANOX |
| FDR | MC2207310-XXX-T | MELLANOX ACTIVE FIBER CABLE VPI UP TO 56GB/S QSFP from 3M up to 100M | MELLANOX |
| FDR | MC2207312-XXX | MELLANOX ACTIVE FIBER CABLE VPI UP TO 56GB/S QSFP  from 3M up to 300M | MELLANOX |
| FDR | MC220731V-XXX | MELLANOX ACTIVE FIBER CABLE VPI UP TO 56GB/S QSFP from 3M up to 100M | MELLANOX |
| FDR | MC2207411-SR4L | MELLANOX OPTICAL MODULE IB FDR 56GB/S QSFP MPO 850NM UP TO 30M | MELLANOX |
| FDR10 | MC2206128-004 | MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 4M | MELLANOX |
| FDR10 | MC2206128-005 | MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 5M | MELLANOX |
| FDR10 | MC2206130-001 | MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 1M | MELLANOX |
| FDR10 | MC2206130-002 | MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 2M | MELLANOX |
| FDR10 | MC2206130-003 | MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 3M | MELLANOX |
| FDR10 | MC2206130-00A | MELLANOX PASSIVE COPPER CABLE VPI UP TO 40GB/S QSFP 0.5M | MELLANOX |
| FDR10 | MC2206310-XXX-E | MELLANOX ACTIVE FIBER CABLE IB QDR/ FDR10 40GB/S QSFP from 3M up to 100M | MELLANOX |
| FDR10 | MC2206310-XXX-T | MELLANOX ACTIVE FIBER CABLE IB QDR/ FDR10 40GB/S QSFP from 3M up to 100M | MELLANOX |
| FDR10 | MC2206310-XXX-F | MELLANOX ACTIVE FIBER CABLE IB QDR/ FDR10 40GB/S QSFP from 3M up to 100M | MELLANOX |
| FDR10 | MC2206310-300-L | MELLANOX ACTIVE FIBER CABLE IB QDR/ FDR10 40GB/S QSFP 300M | MELLANOX |

*Table 3 - Tested Cables and Modules  (Sheet 9 of 9)*

| Speed | OPN # | Description | Vendor |
|---|---|---|---|
| FDR10 | MC2210411-SR4 | MELLANOX OPTICAL MODULE 40GB/S QSFP MPO 850NM UP TO 100M | MELLANOX |
| FDR10 | MC2210411-SR4E | MELLANOX OPTICAL MODULE 40GB/S QSFP MPO 850NM UP TO 300M | MELLANOX |
| FDR10 | MFS4R12CB-XXX | MELLANOX ACTIVE FIBER CABLE VPI UP TO 40GB/S QSFP from  3M up to 100M | MELLANOX |
| QDR | MC2206125-007 | MELLANOX PASSIVE COPPER CABLE IB QDR 40GB/S QSFP 7M | MELLANOX |
| QDR | MC2206126-006 | MELLANOX PASSIVE COPPER CABLE IB QDR 40GB/S QSFP 6M | MELLANOX |

## 1.3    Tested Switches

*Table 4 - Tested Switches  (Sheet 1 of 4)*

| Speed | OPN # /Name | Description | Vendor |
|---|---|---|---|
| 10/40Gb/s | 7050Q | 16-port  40Gb Switch | Arista |
| 10/40Gb/s | 7050S | 48-port 10Gb/40Gb Switch | Arista |
| 40GbE | 7050QX | 32-port 40Gb Switch | Arista |
| 10Gb/s | Brocade 8000 | Brocade 10GB ETH switch | Brocade |
| 10/40Gb/s | 3064 | 48-port 10Gb/40Gb Switch | Cisco |
| 10Gb/s | 5548 | Cisco 10GB ETH switch | Cisco |
| 40GbE | 3016 | Cisco 40GB ETH switch | Cisco |
| 40GbE | 3132Q | Cisco 40GB ETH switch | Cisco |
| 10/40Gb/s | S5000 | 10GbE switch | Dell |
| 10Gb/s | 8024F | 10GbE switch | Dell |
| 10Gb/s | 8132F | 10GbE switch | Dell |
| 10Gb/s | 8164F | 10GbE switch | Dell |
| 1Gb/s | 7024F | 1/10GbE switch | Dell |
| 40GbE | S6000 | 32-port 40Gb Switch | Dell |
| 10/40Gb/s | S4810P-AC | 48-port 10Gb/40Gb Switch | Force 10 |
| 10Gb/s | Fujitsu 10GbE | Ethernet Switch 24 ports, 20xCX4 and 4xQSFP | Fujitsu |
| 10/40Gb/s | ASF5900 | HP 10GB ETH switch | HP |

*Table 4 - Tested Switches  (Sheet 2 of 4)*

| Speed | OPN # /Name | Description | Vendor |
|-------|-------------|-------------|--------|
| 10Gb/s | HP ProCurve 6600-24XG | 24-port 10GbE switch | HP |
| 10GbE | 516733-B21 | HP ProCurve 6120XG 10GbE Ethernet Blade Switch | HP |
| 10GbE | 6125XLG Blade Switch | HP 6125XLG 10/40G Ethernet Blade Switch | HP |
| 10GbE | 538113-B21 | HP 10GbE Pass-Through Module (PTM) | HP |
| 1Gb/s | 2810-24G. | HP 1GB ETH switch | HP |
| 1GbE | 3020X | Cisco Catalyst 3020X 1GbE switch blade | HP |
| 1GbE | 3020 | Cisco Catalyst 3020  1GbE switch blade | HP |
| 1GbE | 438030-B21 | HP 1GbE switch module - GbE2c Layer 2/3 Ethernet Blade Switch | HP |
| 1GbE | 6120G | HP ProCurve 6120G/XG 1GbE switch blade | HP |
| 40GbE | 689638-B21 | Mellanox SX1018HP Enet Switch 40G Ethernet | HP |
| IB DDR | 410398-B21 | HP BLc 4X DDR IB Switch | HP |
| IB FDR | 775144-001 | SwitchX-2 based 18-port QSFP FDR 1U unmanaged InfiniBand switch; R6; compatible to HP Apollo racks | HP |
| IB FDR | 648311-B21 | HP BLc 4X FDR IB Switch | HP |
| IB QDR | 489184-B21 | HP BLc 4X QDR IB Switch | HP |
| 10/40Gb/s | IBM  G8264 | IBM  10/40GB ETH switch | IBM |
| 10GbE | B22 | IBM B22  10 Gigabit Scalable Switch Module | IBM |
| 10GbE | EN4093 | IBM PureFlex System Fabric EN4093 10 Gigabit Scalable Switch Module | IBM |
| 40GbE | IBM G8316 | IBM 40GB  RackSwitch G8316 | IBM |
| 40GbE | 90Y3477 | BM Flex System EN6131 40Gb Ethernet Switch | IBM |
| IB FDR | 90Y3452 | IBM Flex System IB6131 Infiniband Switch | IBM |
| 10/40Gb/s | Juniper EX3500 | Juniper 10/40GB ETH switch | Juniper |
| 10Gb/s | Juniper EX2500 | Juniper 10GB ETH switch | Juniper |
| 10Gb/s | Juniper EX4550 | Juniper 10GB ETH switch | Juniper |
| 40GbE | JuniperQFX5100 | Juniper40GB ETH switch | Juniper |
| FDR | MSX6036F-1SFR | SwitchX based FDR InfiniBand Switch; 36 QSFP; Managed | Mellanox |
| FDR | SRDFSH36F-1BF | SwitchX based FDR InfiniBand Switch; 36 QSFP; Managed | Mellanox |

*Table 4 - Tested Switches  (Sheet 3 of 4)*

| Speed | OPN # /Name | Description | Vendor |
|---|---|---|---|
| QDR | MIS5030Q-2SFC | InfiniScale IV IS5030 36 ports QDR InfiniBand Switch | Mellanox |
| 10/40Gb/s | MSX1024B-1BFS | SwitchX®-2 based 48-port SFP+ 10GbE, 12 port QSFP 40GbE, 1U Ethernet switch | Mellanox |
| 10Gb/s | MSX1016X-1BFR | SwitchX™ based 64-port SFP+ 10GigE, 1U Ethernet switch | Mellanox |
| 40Gb/s | MSX1036B-1BFR | SwitchX™ based 36-port QSFP 40GigE 1U Ethernet | Mellanox |
| EDR | MSB7790-EB2F | Switch-IB(TM) based EDR InfiniBand Switch, 36 QSFP ports, non-blocking switching capacity of 7.2Tbps, | Mellanox |
| IB DDR | Mellanox M2401G | Infiniscale III 24-Port 20Gb/s InfiniBand Switch for Dell M1000E Blade System | Mellanox |
| IB DDR | F-X430044 | 24-port DDR-Switch | Mellanox |
| IB EDR | SB7700 | Scorpion 7700 - Mellanox 36-port EDR 100Gb/s InfiniBand Switch | Mellanox |
| IB FDR | MSX6710-FB2F2 | SwitchX®-2 based FDR InfiniBand 1U Switch, 36 QSFP+ ports, 2 Power Supplies (AC), x86 dual core, short depth, P2C airflow, Rail Kit, RoHS6 | Mellanox |
| IB FDR | MSX6036F-1BFR | SwitchX™ based FDR InfiniBand Switch, 36 QSFP ports, 1 Power Supply, Short depth, Managed, PSU side to Connector side airflow, Rail Kit and RoHS6 | Mellanox |
| IB FDR | Mellanox M4001F | SwitchX® 56Gb/s 16+16 port InfiniBand switch blade for the Dell M1000e Blade System | Mellanox |
| IB FDR | SX6710 | Neptune 83 - Mellanox 36-port FDR 56Gb/s Infini-Band Switch | Mellanox |
| IB FDR10 | MSX6025T-1SFR | SwitchX™ based FDR10 Infiniband Switch, 36 QSFP ports, 1 Power Supply, Standard depth, Unmanaged, PSU side to Connector side airflow, Rail Kit and RoHS6 | Mellanox |
| IB FDR10 | Mellanox M4001T | SwitchX® 40Gb/s 16+16 port InfiniBand switch blade for the Dell M1000e Blade System | Mellanox |
| IB QDR | MIS5025Q-1SFC | InfiniScale® IV QDR InfiniBand Switch, 36 QSFP ports, 1 Power Supply, Unmanaged,PSU side to connector side airflow, Standard depth, Rail Kit and RoHS5 | Mellanox |
| IB QDR | MIS5024Q-1BFR | InfiniScale® IV QDR InfiniBand Switch, 36 QSFP ports, 1 power supply, Unmanaged, PSU side to Connector side airflow, no FRUs, with rack rails, Short Depth Form Factor | Mellanox |

*Table 4 - Tested Switches  (Sheet 4 of 4)*

| Speed | OPN # /Name | Description | Vendor |
|-------|-------------|-------------|--------|
| IB QDR | QDR-Switch 4036 | InfiniScale® IV QDR Mellanox Grid Director 4036 36-Port QDR ,InfiniBand Switch - Part ID: VLT-30011 | Mellanox |
| IB QDR | Mellanox M3601Q | 40Gb/s 16+16 port InfiniBand switch blade for the Dell M1000e Blade System | Mellanox |
| IB SDR | F-X430060 | 24-port SDR-Switch | Mellanox |
| IB DDR | 9024 | 24-port DDR-Switch | Qlogic |
| IB QDR | 12300 | 36-Port 40Gb QDR Infiniband Switch, Management Module, Dual Power | Qlogic |
| 1/10Gb/s | Summit X650 | Extreme 10GB ETH switch | Summit |

## 1.4 Tools, Switch Firmware and Driver Software

Firmware Rev 2.35.5100 is tested with the following tools, SwitchX® firmware, and driver software:

*Table 5 - Tools, Switch Firmware and Driver Software*

|  | Supported Version |
|--|-------------------|
| MLNX_OFED | 3.1-1.0.0/3.0-2.0.4/2.4-1.0.0 |
| MLNX_EN (MLNX_OFED based code) | 3.1-1.0.0/3.0-2.0.4/2.4-1.0.0 |
| WinOF | 4.95/4.90 |
| VMware | 1.9.9.4/1.8.2.4 |
| MFT | 4.1.0/4.0.0 |
| SwitchX/SwitchX-2 Firmware | v9.3.2000 |
| InfiniScale 4 Firmware | v7.4.3000 |

## 1.5 Supported FlexBoot

Firmware Rev 2.35.5100 supports the following FlexBoot version:

*Table 6 - Supported FlexBoot*

|  | Supported Version |
|--|-------------------|
| FlexBoot | 3.4.648 |

## 1.6 Revision Compatibility

Firmware Rev 2.35.5100 complies with the following programmer's reference manual:

- *ConnectX Programmer's Reference Manual (PRM), Rev 2.0 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.*

## 1.7    Firmware Burning Notes

- Firmware Family Version (FFV)

  As of firmware v2.30.8000, all firmware images have the FFV field populated. The FFV value is identical to the firmware version but in a different format.

  FFV format example:

  ```
  FW version:          2.30.8000
  FFV:                 02.30.80.00
  ```

- Updating EXP_ROM

  Updating only the EXP_ROM (FlexBoot) for firmware images which contain FFV requires an additional MFT flag: "-allow_rom_change"

  The following is an example for removing the EXP_ROM from the binary image using Flint (a Mellanox device located at PCI bus function 05:00.0):

  ```
  $ flint -d 05:00.0 -allow_rom_change drom
  ```

# 2      Firmware Rev 2.35.5100 Changes and New Features

*Table 7 - Firmware Rev 2.35.5100 Changes and New Features*

| Category | Description |
|---|---|
| PRM: Performance | New performance and back-pressure counters command via PRM (For further information, please refer to the PRM) |
| PRM: Packet Steering | Support for Multicast/Unicast sniffer rules (For further information, please refer to the PRM) |
| PRM: QinQ | Support for VLAN in VLAN encapsulation (For further information, please refer to the PRM) |
| PRM: QP and CQ initialization optimizations | • CQ creation offload by software<br>• Support for rst2rts command |
| PRM: Non-Volatile configuration | Invalidates a TLV during the firmware boot stage |
| PRM: Diagnostic report | A new counter for the diag_rprt PRM command to count packet drops due to no-receive buffer |
| PRM: Ethernet Head of Queue | Support for Ethernet TX lifetime cycle control (Head of Queue) |
| PRM: Access_Reg | A new register (PPLR) that allows egress and external loopback control (For further information, please refer to the PRM) |
| PRM: NIC resiliency | A watchdog mechanism to track ingress traffic stalls to prevent flooding the network with Flow Control packets |
| LED scheme | Inspur LED scheme: A new LED scheme controlled by the INI which causes constant traffic LED indication even without traffic. |

# 3    Known Issues

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

*Table 8 - Known Issues*

| Index | Issue | Description | Workaround | Scheduled Release (fix) |
|---|---|---|---|---|
| 1. | 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. | Future Release |
| 2. | 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 0xfffff). | N/A |
| 3. | SBR assertion | SBR should be asserted for a minimum of 50 milliseconds for the ConnectX-3 adapters | N/A | N/A |
| 4. | PCIe | On Pilot1 SL230, PCIe link occasionally does not come up at Gen3 speed | Production SL230 should be used for PCIe Gen3 operation | N/A |
| 5. | 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 | N/A |
| 6. | 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 | Future Release |
| 7. | 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 | Future Release |

***Table 8 - Known Issues***

| Index | Issue | Description | Workaround | Scheduled Release (fix) |
|-------|-------|-------------|------------|-------------------------|
| 8. | 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'` | Future Release |
| 9. | MC2210411-SR4 module with Cable Info MAD | Cable Info MAD reports a wrong cable info when using the MC2210411-SR4 module | N/A | Future Release |
| 10. | PCIe failure on temperature shock 10C/min | Gen2 failure at temperature sweep up to 10C/min (for MT27518A1-FDIR-BV only). | N/A | Future Release |
| 11. | PCIe Gen2 link | PCIe Gen2 link unstable at temperature sweep of 10C/min for MT27518A1-FDIR-BV | N/A | Future Release |
| 12. | 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 | Future Release |
| 13. | Bloom filter | Bloom filter is currently not supported. | N/A | Future Release |
| 14. | 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:<br>`You are trying to override configurable FW by non-configurable FW.`<br>`If you continue, old FW configurations will be cleared, do you want to continue ? (y/n) [n] : y`<br>`You are trying to restore default configuration, do you want to continue ? (y/n) [n] : y` | N/A | Future Release |
| 15. | 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 | Future Release |
| 16. | VPD read-only fields | VPD read-only fields are writable. | Do not write to read-only fields if you wish to preserve them | Future Release |

**Table 8 - Known Issues**

| Index | Issue | Description | Workaround | Scheduled Release (fix) |
|-------|-------|-------------|------------|-------------------------|
| 17. | Increasing Sym-bolErrorCounter | When working in VPI mode with port1 FDR and port2 40G, error counters misbehave and increase rapidly | N/A | Future Release |
| 18. | 128 Byte CQ/EQ stride compati-bility with side-band Management | Setting the device to 128Byte CQ/EQ stride will cause misbehavior of side-band management resulting in commu-nication loss. | N/A | Future Release |
| 19. | 128 Byte CQ/EQ stride | CQ and EQ cannot be configured to different stride sizes. | N/A | Future Release |
| 20. | VPI port proto-col change on a port with side-band Manage-ment | 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. | N/A |
| 21. | Link Up time | Habanero (P/N MCX349A-XCCN) may experience longer linkup times of a few seconds with specific switches. | N/A | Future Release |
| 22. | Port identifica-tion | Habanero (P/N MCX349A-XCCN) does not respond to ethtool "identify" command (ethtool -p/--identify). | N/A | Future Release |
| 23. | RDP over IPv6 | RDP over IPv6 is currently not func-tional. | N/A | Future Release |
| 24. | Unicast/Multi-cast 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 | Future Release |
| 25. | Boot Entry Vec-tor (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 | Future Release |

# 4    Bug Fixes History

Table 9 lists the bugs fixed in this release.

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|---|
| 1. | SR-IOV security | Prevented a Virtual Function from injecting pause frames into the network. | 2.30.8000 | 2.35.5100 |
| 1. | PRM:SET_PORT.mac_table | Fixed a wrong hash index calculation during inbound traffic. | 2.33.5100 | 2.35.5100 |
| 2. | RDP over IPv4 | Fixed a default hardware configuration issue which caused RDP over IPv4 traffic to be dropped. | 2.30.8000 | 2.35.5100 |
| 3. | NC-SI | MLNX_OEM command `GET_TEMP` returned a wrong value in the `max_temp` field. | 2.34.5000 | 2.35.5100 |
| 4. | 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 |
| 5. | 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 |
| 6. | IB/RoCE retransmission | Fixed a race in handling a duplicated "read request from middle". | 2.34.5000 | 2.35.5100 |
| 7. | IB traffic issues | Fixed an issue which caused lack of IB traffic on SR-IOV VPI. | 2.33.5000 | 2.35.5100 |
| 8. | 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 |
| 9. | IB APM | Fixed an issue which caused an internal firmware error when APM changed the QPs port mapping. | 2.33.5100 | 2.35.5100 |
| 10. | QP alternate context error | Fixed an issue which caused a firmware internal error when handling QP alternative context. | 2.34.5000 | 2.35.5100 |
| 11. | 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 |
| 12. | Wrong temperature reporting | Wrong temperature reporting when server is in Auxiliary mode after Moonshot AC power cycle. | 2.34.5000 | 2.35.5100 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|---|
| 13. | Link down on Habanero | Fixed a link down issue with 100MbaseT speed. | 2.33.5100 | 2.34.5000 |
| 14. | Packet Ethertype | Fixed a mistakenly dropped ETH packet with ethertype 0x600 by the NIC. | 2.30.8000 | 2.34.5000 |
| 15. | 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 |
| 16. | Low link speed | Fixed an issue where the port raised as SDR vs. InfiniScale IV QDR Switch | 2.33.5100 | 2.34.5000 |
| 17. | 40GbE Link down | Fixed a failure to read cable parameter which caused link failure on 40GbE dual port OCP devices. | 2.33.5100 | 2.34.5000 |
| 18. | 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 |
| 19. | VM QoS | Fixed a case where the actual bandwidth did not match the user settings in VM QoS. | 2.33.5100 | 2.34.5000 |
| 20. | Sideband communication loss | Fixed a case where on rare cases, communication to BMC was lost during driver initialization. | 2.33.5100 | 2.34.5000 |
| 21. | LED behavior | Fixed a wrong LED behavior when the driver is disabled in the following adapter cards: MCX346A-BCQN, MCX345A-BCQN. | 2.33.5100 | 2.34.5000 |
| 22. | Link down on cable plugging | Fixed an issue with cable reading, which caused the link not to raise | 2.33.5100 | 2.34.5000 |
| 23. | PRM: EQN range | Set the maximum EQN number to 1024. | 2.30.8000 | 2.34.5000 |
| 24. | Vital Product Data read failure | Fixed a rare issue with VPD init flow which caused read failures. | 2.31.5050 | 2.34.5000 |
| 25. | PRM: Statistic counters not reported | Fixed an issue with RX size counter not being reported. | 2.30.8000 | 2.34.5000 |
| 26. | RoCE/InfiniBand reliable connection | The first Read response was not treated as implicit ACK. | 2.30.8000 | 2.33.5100 |
| 27. | 40GbE Link up time | Reduced a long 40GbE link up time with Cisco Nexus3064 and Arista-7050S | 2.32.5100 | 2.33.5100 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|-------|-------|-------------|----------------------|------------------|
| 28. | Promiscuous mode | Fixed promiscuous mode compatibility with A0-DMFS steering. | 2.32.5100 | 2.33.5000 |
| 29. | | Fixed promiscuous mode compatibility when NC-SI is enabled and configured. | 2.32.5100 | 2.33.5000 |
| 30. | NC-SI OEM commands | Fixed sending/receiving OEM temp commands (set/get) with channel ID 0x1f failure. | 2.32.5100 | 2.33.5000 |
| 31. | 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 | 2.33.5000 |
| 32. | Side Band Functionality | Fixed long management communication loss and SOL hang during reboot cycles. | 2.32.5100 | 2.33.5000 |
| 33. | | Fixed wrong processing of inbound traffic towards BMC which caused communication loss. | 2.32.5100 | 2.33.5000 |
| 34. | | Fixed management link loss upon closing port interface through the driver. | 2.32.5100 | 2.33.5000 |
| 35. | NC-SI on SFP+ Adapter Cards | Fixed a false indication in firmware of an expander presence causing delay in EEPROM reading. | 2.32.5100 | 2.33.5000 |
| 36. | 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 | 2.33.5000 |
| 37. | Inbound Packet Processing | Fixed a rare case where packet with length zero got stuck in hardware queues. | 2.32.5100 | 2.33.5000 |
| 38. | | Fixed an issue which caused InfiniBand congestion control packet (CNP) to hang in hardware. | 2.32.5100 | 2.33.5000 |
| 39. | Asynchronous Event Notification (AEN) | Fixed an issue which caused AEN to be sent after channel reset. | 2.32.5100 | 2.33.5000 |
| 40. | Bandwidth Degradation with QoS | Fixed an issue which prevented the restoring of QoS setting to its default consequently causing bandwidth degradation. | 2.31.5050 | 2.33.5000 |
| 41. | Port Link Up Time | Fixed an occasional long link up time with 10GbE based devices. | 2.32.5100 | 2.33.5000 |
| 42. | SFP Cable Reading | Fixed an issue preventing cable readings from i2c slave address 0x51 | 2.32.5100 | 2.33.5000 |
| 43. | PCIe Gen3 Equalization | Fixed a wrong parity bit calculation when transmitting PCIe TS1 packets. | 2.32.5100 | 2.33.5000 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|-------|-------|-------------|-----------------------|------------------|
| 44. | PCIe Power Management | Fixed a possible deadlock in PM turnoff request transmission and ack acceptance flow. | 2.32.5100 | 2.33.5000 |
| 45. | PCIe width Degrade | Fixed a rare case with alignments state machines which caused occasional width degradation. | 2.32.5100 | 2.33.5000 |
| 46. | 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 | 2.33.5000 |
| 47. | 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 | 2.33.5000 |
| 48. | IB Spec MADs | Fixed an issue which caused `SetPortInfo` to return a good status when receiving invalid `LinkSpeedEnabled` value. | 2.32.5100 | 2.33.5000 |
| 49. | GPIO Mapping | Fixed an issue which caused dual port SFPP module cards to be automatically mapped with expander | 2.32.5100 | 2.33.5000 |
| 50. | Steering Mode | Fixed an issue where firmware overrides the steering mode that was chosen by the driver. | 2.32.5100 | 2.33.5000 |
| 51. | Port sensing | Fixed invalid return sensing results occurred when the link was up. | 2.32.5100 | 2.33.5000 |
| 52. | | Fixed an issue causing the sensing result to be delayed when cable was unplugged. | 2.32.5100 | 2.33.5000 |
| 53. | Wrong link type display | Fixed an issue causing the link type to be displayed as ETH when set to AUTO. | 2.32.5100 | 2.33.5000 |
| 54. | IBDump performance | Fixed performance degradation when running IBDump | 2.30.8000 | 2.32.5100 |
| 55. | PCIe link Disable/Enable | Occasionally, a link training timeout occurred in EQ phase0 during disable/enable test. | 2.31.5050 | 2.32.5100 |
| 56. | 40GbE QoS | Improved strict bandwidth mode functionality | 2.30.8000 | 2.32.5100 |
| 57. | Port Counters reporting | Fixed an issue with the PortRcvPkts counter always displaying zero value. | 2.31.5050 | 2.32.5100 |
| 58. | GMP MADs in SecureHost | Fixed an issue with processing GMP MADs with SET method in SecureHost mode. | 2.31.5050 | 2.32.5100 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|---|
| 59. | NC-SI over IPv6 | Fixed an issue causing a wrong usage of MCG size when configuring Global Multi-cast filter | 2.31.5050 | 2.32.5100 |
| 60. | NC SI link failure | Disabling the first port occasionally causes second port TX failure. | 2.31.5050 | 2.32.5100 |
| 61. | 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 |
| 62. | Link failure | Fixed an occasional 40GbE link failure with SCM5 Switch blade | 2.31.5050 | 2.32.5100 |
| 63. | ExtPortInfo MAD | Fixed a wrong FDR10 speed reporting in MAD | 2.31.5050 | 2.32.5100 |
| 64. | IB link failure | Fixed an issue preventing theports to to rise up when set to FDR10 vs QDR | 2.31.5050 | 2.32.5100 |
| 65. | 40GbE link failure | Fixed an occasional link failure vs Arista switch | 2.31.5050 | 2.32.5100 |
| 66. | RDMA Write retrans-mission | Retransmission started from the first PSN of message instead of the last acknowledged PSN | 2.30.3200 | 2.32.5100 |
| 67. | Firmware burning | Firmware hangs when receiving `GeneralInfoMad` during inline firmware burning | 2.30.3200 | 2.32.5100 |
| 68. | PCIe PML1 | L1 flow adjustments and threshold tuning | 2.31.5050 | 2.32.5100 |
| 69. | PCIe reset | Fixed a rare hanging issue during PERST_ assertion | 2.31.5050 | 2.32.5100 |
| 70. | PCIe Gen3 EQ | Wrong coefficients were reported during phase3 | 2.31.5050 | 2.32.5100 |
| 71. | Boot | Fixed an issue causing wrong behavior due to reset timing | 2.31.5050 | 2.32.5100 |
| 72. | SMBUS | Fixed long timeout issues | 2.31.5050 | 2.32.5100 |
| 73. | NVRAM | Fixed NVRAM write issues in driver-less mode | 2.31.5050 | 2.32.5100 |
| 74. | 40GbE Link support | Fixed 40GbE link support in aux mode | 2.31.5050 | 2.32.5100 |
| 75. | NC-SI | Dropped commands with non-existing channel ID | 2.31.5050 | 2.32.5100 |
| 76. | PRM PortInfo com-mand | Fixed issues in extended speed reporting | 2.31.5050 | 2.32.5100 |
| 77. | Trap 257/8(IB) | Fixed bad QP reporting in trap 257/8 | 2.30.8000 | 2.32.5100 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|---|
| 78. | Bad Q_KEY errors | Fixed an issue causing false bad q_key error messages | 2.30.8000 | 2.32.5100 |
| 79. | PFC | Fixed Pause Frame opcode mismatch | 2.30.8000 | 2.32.5100 |
| 80. | Sideband Communication | Fixed communication loss upon PCIe error detection | 2.31.5050 | 2.32.5100 |
| 81. | NC-SI | Fixed wrong channel value in the SELECT/DESELECT PACKAGE commands | 2.30.8000 | 2.31.5050 |
| 82. | | Fixed an issue caused response packet to include 4 extra bytes | 2.30.8000 | 2.31.5050 |
| 83. | | Fixed wrong reason code value returned when using Set Link command with unsupported speed. | 2.30.8000 | 2.31.5050 |
| 84. | | Added protection from bad MAC address given by BMC | 2.30.8000 | 2.31.5050 |
| 85. | False Link Indication | Fixed an issue causing the device to report false link up when no cable is connected. | 2.30.8000 | 2.31.5050 |
| 86. | PCIe | Removed false TX pulse after PERST_ de-assertion | 2.30.8000 | 2.31.5050 |
| 87. | | Fixed FLR capability bit inconsistency when SR-IOV is enabled. | 2.30.3200 | 2.31.5050 |
| 88. | | Fixed an issue with the device not reporting PCIe related errors. | 2.30.8000 | 2.31.5050 |
| 89. | 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 |
| 90. | VXLAN | VXLAN used the wrong default UDP port. the UDP port number was changed to 4789. | 2.30.8000 | 2.31.5050 |
| 91. | | Fixed wrong setting of the UDP destination port for VXLAN. | 2.30.8000 | 2.31.5050 |
| 92. | Flow Steering | Fixed an internal error caused when moving to the DMS mode with IPMI/NC-SI enabled. | 2.30.8000 | 2.31.5050 |
| 93. | 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 |
| 94. | PCI interrupt | Fixed issues related to working with PCI legacy interrupts. | 2.30.8000 | 2.31.5050 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|---|
| 95. | TCP/UDP Checksum | Fixed wrong checksum calculation for short packets which were padded by the software. | 2.30.8000 | 2.31.5050 |
| 96. | MFT tool deadlock | Reading PCIe configuration space after using the MFT flint tool caused the device to crash. | 2.10.0000 | 2.31.5050 |
| 97. | Side band packet loss | Fixed occasional packet loss over IPMI | 2.30.8000 | 2.31.5050 |
| 98. | Eye opening MAD | Fixed wrong values reported in the Eye opening MAD. | 2.30.8000 | 2.31.5050 |
| 99. | PCIe Link width | Fixed occasional link width degrades during link negotiation and link transitions from L1 state. | 2.30.8000 | 2.31.5050 |
| 100. | PCIe signal detect | Fixed adjust signal detect thresholds | 2.30.8000 | 2.31.5050 |
| 101. | Error counters | PortExtendedSpeedsCounters MAD counters were mistakenly increased while LLR was active | 2.30.8000 | 2.31.5050 |
| 102. | PCIe Gen3 Equalization | Lane reversal was not considered when configured TX parameters | 2.30.8000 | 2.31.5050 |
| 103. | Reset On LAN (ROL) | Fixed ROL factory MAC usage when a FlexBoot address was given. | 2.30.8000 | 2.31.5050 |
| 104. | Flow Control | Fixed Pause frames factory MAC usage when a FlexBoot address was given. | 2.30.3200 | 2.31.5050 |
| 105. | WOL/ROL | The device did not different between WOL/ROL packets. | 2.30.8000 | 2.31.5050 |
| 106. | PortInfo MAD | Fixed a set of extended fields in PortInfo MAD which did not function. | 2.10.0000 | 2.31.5050 |
| 107. | LLR cell size | Adjusted LLR cell size according to the MLPN negotiation of ib_128b_llr | 2.30.8000 | 2.31.5050 |
| 108. | Link max speed | The max speed restriction was active in full power mode instead of standby mode only. | 2.30.8000 | 2.31.5050 |
| 109. | InfiniBand Automatic Path Migration | The InfiniBand Path migration did not work with GRH. http://webdev01:8080/commit/ConnectX.git/ a9c37ee4c31038f2c1179d4d9e79c9337e0a b5c7 | 2.10.0000 | 2.31.5050 |
| 110. | Packet steering | Reading MGM after writing it returned wrong members count. | 2.30.8000 | 2.31.5050 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|-------|-------|-------------|----------------------|------------------|
| 111. | RSS QP context | Fixed corruption of the RSS hash key given by the driver. | 2.30.8000 | 2.31.5050 |
| 112. | 10Gb/s QoS | Fixed QoS rate limit bandwidth offset. | 2.30.3200 | 2.31.5050 |
| 113. | ExtendedPortInfo MAD | Fixed FDR10 speed_en reporting. | 2.30.8000 | 2.31.5050 |
| 114. | Management link | Fixed long management link com loss. | 2.30.8000 | 2.31.5050 |
| 115. | PRM Query_Port Command | The command results reported both link types active at the same time. | 2.30.3200 | 2.31.5050 |
| 116. | Link not raising | Fixed collision between forcing phy type and port sensing. | 2.30.8000 | 2.31.5050 |
| 117. | Core clock reporting | Fixed a wrong core clock freq reporting in QUERY_HCA command. | 2.30.3200 | 2.31.5050 |
| 118. | 56GbE link issues | Fixes occasional link failure when 56GbE is enabled | 2.30.8000 | 2.31.5050 |
| 119. | RX calibration | Fixed max eye margins to be per protocol. | 2.30.8000 | 2.31.5050 |
| 120. | VPI symbol errors | perfquery reported wrong error symbol on ConnectX©-3 VPI mode: IB, ETH. | 2.30.8000 | 2.31.5050 |
| 121. | 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 |
| 122. | 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 |
| 123. | 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 |
| 124. | 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 |
| 125. | 40GbE is not supported in Auto-Sensing | Auto-Sensing is not supported with 40GbE connections in VPI cards | 2.10.0000 | 2.30.8000 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|-------|-------|-------------|----------------------|------------------|
| 126. | InfiniBand port_rcv_pkts counter | InfiniBand port_rcv_pkts counter over counts when LLR is enabled on the port. The port_rcv_pkts counter continues to count packets even when no traffic is received. | 2.10.0000 | 2.30.8000 |
| 127. | PXE | PXE is currently not supported in 40GbE in VPI cards | 2.10.0000 | 2.30.8000 |
| 128. | | PXE is currently not supported in QSFP to SFP+ hybrid cable | 2.10.0000 | 2.30.8000 |
| 129. | 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 |
| 130. | Link errors | BER of $10^{-11}$ with 7M copper SFP+ 10GbE cable against Arista switch | 2.30.3200 | 2.30.8000 |
| 131. | 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 |
| 132. | NC-SI 40GbE reporting | Added 40GbE reporting in get_link_status NC-SI command | 2.30.3000 | 2.30.8000 |
| 133. | 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 |
| 134. | PCIe speed degrade | Occasionally, PCIe speed degraded during speed change test | 2.30.3000 | 2.30.8000 |
| 135. | Link failure vs Cisco | Device failed to raise the link against Cisco b-22 Blade switch | 2.30.3000 | 2.30.8000 |
| 136. | 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 |
| 137. | NC-SI wrong command info | Wrong port information provided in get link status command. | 2.30.3000 | 2.30.8000 |
| 138. | Port interfaces stay down | After firmware upgrade, the device failed to raise port interface. | 2.30.3000 | 2.30.8000 |
| 139. | SR-IOV guest failure | Guest VM failed to execute firmware commands operations and crashed. | 2.30.3000 | 2.30.8000 |
| 140. | ConnectX®-3 Pro: MCG write timeout | PRM WRITE_MCG command caused the device to hang. | 2.30.3000 | 2.30.8000 |

**Table 9 - Fixed Bugs List**

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|---|
| 141. | ConnectX®-3 Pro: Driver start failure | Driver could not start when NIC was configured for NC-SI SNP. | 2.30.3000 | 2.30.8000 |
| 142. | SR-IOV command timeouts | Guest MSIX vectors were not assigned properly. | 2.30.3000 | 2.30.8000 |
| 143. | 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 |
| 144. | QP Sniffer issue | Fixed a hash fold issue for sniffer QPs | 2.30.3000 | 2.30.8000 |
| 145. | 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 |
| 146. | PCIe TS parity bit | In recovery.EQLZ state TSs used incorrect parity bit calculation | 2.30.3000 | 2.30.8000 |
| 147. | PRM Init_port failure | Init port command may fail on a system with NC-SI | 2.30.3000 | 2.30.8000 |
| 148. | PortInfo MAD link width support | Wrong link_width_support is reported occasionally in PORT_INFO MAD | 2.30.3000 | 2.30.8000 |
| 149. | 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 |
| 150. | PCIe PML1 failures | Fixed occasional failures upon entering and exiting L1 state in PCIe Gen1 & 2 speeds | 2.30.3000 | 2.30.8000 |
| 151. | IB: APM failures in ConnectX-3 Pro | Occasional FSM transition timeouts are seen on APM requests | 2.30.3000 | 2.30.8000 |
| 152. | ipmitool OOB commands | On rare occasions, ipmitool OOB commands failed upon send payload. | 2.30.3000 | 2.30.8000 |
| 153. | Sideband communication | On rare occasions, after stress of power cycles, side-band communication might disconnect. | 2.30.3000 | 2.30.8000 |
| 154. | Expansion ROM partition | Expansion ROM partition not found | 2.30.3000 | 2.30.8000 |
| 155. | AC power cycle issue | In certain servers, AC power cycle may cause BMC connectivity loss. | 2.30.3000 | 2.30.8000 |
| 156. | PXE issue | Occasionally, during DC cycle stress, failure occurred in PXE due to race condition | 2.30.3000 | 2.30.8000 |
| 157. | SMBUS communication | SMBUS communication lost during AC/DC cycle | 2.30.3000 | 2.30.8000 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|---|
| 158. | MTU configuration | Temporal wrong MTU configuration during initialization may cause Serial over LAN disconnection | 2.30.3000 | 2.30.8000 |
| 159. | IPMI connectivity | IPMI OOB communication lost during stress | 2.30.3000 | 2.30.8000 |
| 160. | iperf stress test | Packet drops during iperf stress w/ different MSS | 2.30.3000 | 2.30.8000 |
| 161. | 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 |
| 162. | Running rmmod may cause unexpected behavior | Removing the mlx4_en driver using the "rmmod" command, may cause unexpected behavior | 2.30.3000 | 2.30.8000 |
| 163. | IPMI connectivity | On rare occasions, after stress of BMC cold reset, link failure might occur | 2.30.3000 | 2.30.8000 |
| 164. | Flexboot 3.4.100-UEFI-4.0.410 BIOS menu | Enter the BIOS menu while using Flex-Boot 3.4.100- UEFI-4.0.410 may cause the server to stop responding (hang) | 2.30.3000 | 2.30.8000 |
| 165. | ROL failure | ROL failure after disassembling the driver on the SUT | 2.30.3000 | 2.30.8000 |
| 166. | IPMI link failure | IPMI link failure after disabling the WoL or disassembling the driver | 2.30.3000 | 2.30.8000 |
| 167. | WoL and RoL issues | WoL and RoL issues caused when the IPMI is disabled | 2.30.3000 | 2.30.8000 |
| 168. | A link flapping issue | Alignment marker arrival can no longer drop the link. | 2.30.3000 | 2.30.8000 |
| 169. | IPMI - SOL traffic performance improve- ment | SOL with multiple data streams occasionally hang | 2.30.3000 | 2.30.8000 |
| 170. | RoCE | RoCE does not function properly after running "ethtool ethX" | 2.30.3000 | 2.30.8000 |
| 171. | PCI link errors | PCI link errors false indication. Cleared errors during PCIe link retraining | 2.30.3000 | 2.30.8000 |
| 172. | PCIe speed change | Fixed a false indication for incoming PCIe speed change request | 2.30.3000 | 2.30.8000 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|-------|-------|-------------|----------------------|------------------|
| 173. | 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 |
| 174. | Ports PLL calibration issue | PLL calibration were affected by operating point configuration | 2.30.3000 | 2.30.8000 |
| 175. | Cable Info MAD issue | Wrong cable info was received when using the MC2210411-SR4 module | 2.30.3000 | 2.30.8000 |
| 176. | 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 |
| 177. | Port Error counters reset | Port error counters were not cleared upon XAUI/SGMII link up | 2.30.3000 | 2.30.8000 |
| 178. | 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 |
| 179. | 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 |
| 180. | Access to closed resources | Fixed a possible access to unmapped resource memory | 2.30.3000 | 2.30.8000 |
| 181. | 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 |
| 182. | PCIe speed change | Occasionally, a failure occurred in speed when changing to Gen2 | 2.30.3000 | 2.30.8000 |
| 183. | PXE teardown issue | PXE might halt during teardown | 2.30.3000 | 2.30.8000 |
| 184. | InfiniBand loopback | InfiniBand loopback was blocked during link negotiation on the same port | 2.11.0500 | 2.30.8000 |
| 185. | Voltage scaling | Fixed process voltage scaling issue | 2.30.3000 | 2.30.8000 |
| 186. | DMA address 0x0 | Fixed a possible read access to DMA address 0x0 | 2.11.0500 | 2.30.8000 |
| 187. | cqe issue | Fixed miss cqe issue due to interrupt moderation | 2.11.0500 | 2.30.8000 |
| 188. | Cable reading issue | Fixed a rare cable reading issue upon cable insertion | 2.30.3000 | 2.30.8000 |
| 189. | Linkup issue | Fixed a linkup issue against MSX60XX FDR switch | 2.30.3000 | 2.30.8000 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|---|
| 190. | LLR Vendor Specific MAD | LLR Dropped cell counter reported CSN error | 2.11.0500 | 2.30.8000 |
| 191. | PXE boot failure | On rare occasions, PXE boot fails due to a firmware issue interfering with the PXE load from the card's flash.<br>Upon failure, the following message is received: "Payload inaccessible - cannot continue" | 2.30.3000 | 2.30.8000 |
| 192. | MLNX_OFED v2.0-2.0.3 reports ETH RX errors | MLNX_OFED v2.0-2.0.3 reported RX errors when the driver operated in the ETH only mode. | MLNX-_OFED v2.0.2.0.3 | 2.30.3200 |
| 193. | SR-IOV guest communication channel error | Under certain conditions, SR-IOV guest experienced request timeouts and got stuck. | 2.11.0500 | 2.30.3200 |
| 194. | No NC-SI after PXE teardown | Management transport was not supported, after PXE driver teardown (unload).<br>Fixed the flow to issue software reset, after the driver was stopped. | 2.11.0500 | 2.30.3000 |
| 195. | Revision ID | Updated the Revision ID of Node Info and General Info MADs to reflect hardware Rev ID instead of 0 | 2.11.0500 | 2.30.3000 |
| 196. | Ports PLL calibration issue | PLL calibration were affected by operating point configuration | 2.11.0500 | 2.30.3000 |
| 197. | Port Error counters reset | Port error counters were not cleared upon XAUI/SGMII link up | 2.11.0500 | 2.30.3000 |
| 198. | Receiver SerDes tuning | Enhanced the receiver SerDes tuning for 10GE, to support specific 10GbE QSFP to SFP+ splitter cables | 2.11.0500 | 2.30.3000 |
| 199. | Access to closed resources | Fixed a possible access to unmapped resource memory | 2.11.0500 | 2.30.3000 |
| 200. | Slow link establishment in NC-SI | Slow link establishment in NC-SI caused due to slow cable reading in boot | 2.11.0500 | 2.30.3000 |
| 201. | PCIe speed change | Occasionally, a failure occurred in speed when changing to Gen2 | 2.11.0500 | 2.30.3000 |
| 202. | PXE teardown issue | PXE might halt during teardown | 2.11.0500 | 2.30.3000 |
| 203. | InfiniBand loopback | InfiniBand loopback was blocked during link negotiation on the same port | 2.11.0500 | 2.30.3000 |
| 204. | Voltage scaling | Fixed process voltage scaling issue | 2.11.0500 | 2.30.3000 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|-------|-------|-------------|-----------------------|------------------|
| 205. | DMA address 0x0 | Fixed a possible read access to DMA address 0x0 | 2.11.0500 | 2.30.3000 |
| 206. | cqe issue | Fixed miss cqe issue due to interrupt moderation | 2.11.0500 | 2.30.3000 |
| 207. | Cable reading issue | Fixed a rare cable reading issue upon cable insertion | 2.11.0500 | 2.30.3000 |
| 208. | 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" | 2.11.0500 | 2.30.3000 |
| 209. | 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 | 2.11.0500 | 2.30.3000 |
| 210. | Advanced Error Reporting | Fixes to Advanced Error Reporting according to the PTC (PCIe compliancy) Test failures | 2.11.0500 | 2.30.3000 |
| 211. | 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. | 2.11.0500 | 2.30.3000 |
| 212. | 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 | 2.11.0500 | 2.30.3000 |
| 213. | Wrong link speed after several cable re-insertions | During repeated cable reinsertion, the link may raise in a lower speed than expected/ | 2.11.0500 | 2.30.3000 |
| 214. | Changing port protocol from ETH to IB requires driver restart | Bad configuration of ib_protocol when setting the port to InfiniBand after exchanging it from Ethernet/RoCE on the same port | 2.11.0500 | 2.30.3000 |
| 215. | 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 | 2.11.0500 | 2.30.3000 |
| 216. | Race in PCI configuration handling | The system was unresponsive when a race between PCI configuration cycles handling in hardware and firmware occurred. | 2.11.0500 | 2.30.3000 |
| 217. | RoCE breaks IPv6 traffic | IPv6 packets dropped while RoCE was enabled | 2.11.0500 | 2.30.3000 |

### Table 9 - Fixed Bugs List

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|-------|-------|-------------|----------------------|------------------|
| 218. | 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. | 2.11.0500 | 2.30.3000 |
| 219. | 40GbE switches link down upon repeated insertion and removing of cables | When repeatedly plugging out and plugging in cables to/from a 40GbE switch, the link remains down. | 2.10.0800 | 2.11.0500 |
| 220. | FDR/QDR vs DDR switch | The link is raised as SDR rather than DDR | 2.10.0800 | 2.11.0500 |
| 221. | 0.5M, 1M Copper FDR cable support | 0.5M, 1M Copper FDR cables (Paralink11, Paralink14) are not supported in back to back connection. | 2.10.0800 | 2.11.0500 |
| 222. | Changing FDR cables from Fiber to Copper | Changing FDR cables from Fiber to Copper and vice versa might cause link speed degradation from FDR to FDR10 or FDR10 to QDR. | 2.10.0800 | 2.11.0500 |
| 223. | Changing 40GbE connection | Changing 40GbE connection from back to back to a switch port and vice versa might cause the link to stay down in the ConnectX®-3 side. | 2.10.0800 | 2.11.0500 |
| 224. | SDR vs QDR switch | InfiniBand link vs QDR switch rises as DDR or SDR | 2.10.0800 | 2.11.0500 |
| 225. | QDR link up time | On rare occasions, the QDR link up time between an FDR card and a QDR device was up to 4 minutes | 2.10.0800 | 2.11.0500 |
| 226. | QDR Link stability | QDR Link stability between ConnectX®-3 and InfiniScale® IV | 2.10.0800 | 2.11.0500 |
| 227. | Signal integrity issues | Signal integrity issues in all speeds | 2.10.0800 | 2.11.0500 |
| 228. | PCI correctable error | Fixed bad PCI reporting | 2.10.0800 | 2.11.0500 |
| 229. | RoCE re-transmission | Not re-transmitting from the beginning of the message but from PSN NAK. | 2.10.0800 | 2.11.0500 |
| 230. | RoCE | R-RoCE ignored SMAC check | 2.10.0800 | 2.11.0500 |
| 231. | Function Level Reset (FLR) | FLR to PPF in SRIOV | 2.10.0800 | 2.11.0500 |
| 232. | | FLR in no-driver mode | 2.10.0800 | 2.11.0500 |
| 233. | SR-IOV | Comchannel bug fix | 2.10.0800 | 2.11.0500 |

*Table 9 - Fixed Bugs List*

| Index | Issue | Description | Discovered in Release | Fixed in Release |
|---|---|---|---|---|
| 234. | 40GbE performance | 40GbE full wire speed was not reached with former cards configuration for MCX314 and MCX313 40GbE cards.<br>**Fix:** Increased core frequency to support 40GbE full wire speed. This change requires 1V input.<br>**Affected cards:** MCX313A-BCB, MCX-314A-BCB | 2.10.0700 | 2.10.0800 |
| 235. | | The ConnectX® card was not receiving frames with minimum IPG when configured to 40GbE.<br>Fix: Updated firmware configuration.<br>Affected cards: All cards capable of 40GbE speed | 2.10.0700 | 2.10.0800 |
| 236. | GUIDs for RoCE | Ethernet only cards did not have GUIDs flashed on the device. This caused RoCE to fail once two or more cards were active on the same machine.<br>**Fix:** The firmware generates the GUIDs from the flashed MACs for ETH only cards.<br>**Affected cards:** All ETH only cards. | 2.10.0700 | 2.10.0800 |
| 237. | 40GbE signal integrity | Signal integrity improved in 40GbE speed. | 2.10.0000 | 2.10.0700 |
| 238. | | | | |
| 239. | PCIe Gen3 receiver sensitivity due to BIOS changes | Due to BIOS changes required by latest Intel guidelines for PCIe Gen3, a sensitivity in PCIe eye_centering algorithm caused link degradation on some platforms. | 2.10.0600 | 2.10.0620 |
| 240. | FDR, FDR10, 10GbE XFI/SFI signal integrity | Signal integrity improved for the following speeds:<br>FDR, FDR10, 10GbE XFI/SFI. | 2.10.0000 | 2.10.0600 |
| 241. | | | | |
| 242. | VDD_MODE GPIO changed after reset | In boards with VDD_MODE GPIO, the GPIO was dropped after reset hence, causing voltage to be reduced to 0.9v.<br>**Fix:** The GPIO in boards with VDD_-MODE GPIO is no longer dropped across resets. | 2.10.0000 | 2.10.0600 |
| 243. | PCIe configuration type1 requests were dropped after a PCIe hot reset | After a PCIe hot reset, PCIe configuration type1 requests were dropped instead of being replied as unsupported.<br>**Fix:** A proper notification is sent every time PCIe configuration type1 requests are dropped. | 2.10.0000 | 2.10.0600 |

# 5 Firmware Changes and New Feature History

*Table 10 - Firmware Changes and New Feature History*

| Firmware Version | Description |
|---|---|
| 2.34.5000 | • Added support for multiple RoCE modes (RoCE v1+v2) on the same port: RoCE mode is per connection now.<br>• Added a new QP command "INIT2RTS_QP" to enhance QP connection readiness time.<br>• Disabled FCS checks to support switches that replace FCS with Timestamp.<br>• Added RX Port identification for direct rout packets.<br>• Improved RDMA WRITE/SEND performance with retransmissions.<br>• Enabled firmware burning/querying using the PRM ACCESS_REG command.<br>• Added support for VAM.<br>• Enabled bad cable EEPROM reporting to the driver.<br>• Added support for Platform Level Data Model (PLDM) sideband protocol.<br>• Added support for priority based A0-DMFS mode (For further information, please refer to the PRM).<br>• Added support for Unicast/Multicast loopback disablement by the driver. (For further information, please refer to the PRM)<br>• Removed the source IP from the hash calculation (For further information, please refer to the PRM)<br>• Added support for Inline Receive mode up to 2KB. |
| 2.33.5100 | • Bug fixes - see "Bug Fixes History" on page 24 |
| 2.33.5000 | • Bandwidth allocation support: Including maximum bandwidth and bandwidth share guarantee between VMs for InfiniBand and Ethernet.<br>• Increased inbound traffic buffer capacity when the PFC on all priorities is enabled.<br>• Added support for changing UAR BAR (PCI BAR 2) size.<br>• Added support for cable sub-power class for Mellanox MFA1A00-EXXX and SMFA1A00-CXXX EDR cables.<br>• Improvements in attachment/detachment flows' rules in both A0-DMFS and DMFS modes.<br>• Added physical port forcing on specific QPs when virtual mapping is applied<br>• Added support for dynamic enablement of LAG mode<br>• Added support for vendor specific command to report the ports' MAC addresses.<br>• Enabled 100Mb ability exposure and its enablement via an INI parameter.<br>• Added support for SFP+ with 1GbE when the adapter card is enabled in the EEPROM.<br>• Optimized the SideBand connectivity loss during driver initialization to minimum.<br>• Added support for SMBUS ARP.<br>• Enabled thermal reporting of TMP421 sensor in OCP cards.<br>• RDMA Read retransmission optimizations to improve performance and ensure forward progress while packet drops occur.<br>• Improved data path WQE prefetch algorithm. |

*Table 10 - Firmware Changes and New Feature History*

| Firmware Version | Description |
|---|---|
| 2.32.5100 | • Added support to query PTYS, PTOS registers through ACCESS_REG PRM command.<br>• Added support for CLP access to NVRAM<br>• Added support for more than 22 QPs per MCG in DMFS.<br>• Added support for high rate steering mode (a.k.a Simplified Steering)<br>• Added support for reading current hardware mode through the QUERY_PORT PRM command<br>• Added CSUM mode reporting in QUERY_DEV_CAP command<br>• Added additional configuration options for UPDATE_QP command<br>• Added support for 128 Byte stride for CQ/EQ<br>• Enabled module EEPROM access using command I/F<br>• Reset Flow improvements and graceful handling of error caused by Virtual Functions<br>• RX performance optimization for single port cards<br>• Promiscuous mode performance improvements<br>• Added support for Secure Host mode<br>• Added Port protocol configuration option.<br>• Added support for GPIO swap<br>• 40GbE SI improvements<br>• Added support for Temp Sensing Vendor specific MAD.<br>• Added Temp Sensing NC-SI cmd.<br>• Added support for AEN.<br>• Added new command to report firmware revision.<br>• Added support for QCN<br>• Enabled the driver to use VXLAN offloads on TX side without Device Managed Flow Steering (DMFS)<br>• Enabled non Mellanox cables to rise FDR10 link via new INI parameter. To unlock the cables run:<br>`Fdr10_cable_stamping_override` |
| 2.31.5050 | • Added support for GeneralInfo SMP MAD<br>• Updated capability mask in GeneralInfo SMP/GMP MAD<br>• Added support for PortCountersVL MAD<br>• Added support for PortSamplesControl/PortSamplesResults/PortSamplesExtended MADs<br>• Added Exponential Backoff Timer support. It is enabled via the `rtm_ini` parameter. The default value is 0<br>• Added VLAN steering to Device Managed Flow Steering (DMFS)<br>• Added support for Non-Volatile configuration of TLVs to set device attributes:<br>  • Query and set of configurations is available through PRM ACCEES_REG command<br>  • PRM ACCESS_REG command is now also supported through the tools_hcr command interface<br>• Added support for MTF mlxconfig tool<br>• Added IPv6 support for NC-SI and IPMI Pass-Through<br>• Added support for the same unicast MAC simultaneously for both IPMI and NC-SI<br>• PCIe power optimizations for 8X/4X links<br>• SMBUS optimizations |

*Table 10 - Firmware Changes and New Feature History*

| Firmware Version | Description |
|---|---|
| 2.31.5050 (cont.) | • Added enhancements for receiver equalization in Gen3:<br>  • Enhancements are enabled by the INI. The default value is disabled. Please contact Mellanox support if required to enable it.<br>• Added new Physical and Virtual Functions reset flows support<br>• Added support for 64Bit BIOS mode<br>• Added IEEE802.3 CL73 autoneg support to the QUERY_PORT command.<br>• Added factory MAC address reporting to the Query_Port command.<br>• Added support for reverting virtual MAC configuration per port and restoring to factory MAC through MOD_STAT_CFG command.<br>• Added support for inline TLV read through MOD_STAT_CFG command.<br>• Added current MTU reporting to the QUERY_PORT command.<br>• Added support query for additional MAC addresses per port (up to 7) through the QUERY_PORT command. |
| 2.30.8000 | • Initial GA release of ConnectX-3 Pro<br>• UDP packets with zero checksum<br>• RoCE v2 support, including CONFIG_DEV command support<br>• Enabled SR-IOV by default on all Mellanox ConnectX-3 Pro cards with 8 virtual functions<br>• indiscard packets counter support in DUMP_ETH_STAT command<br>• NVGRE support<br>• VXLAN support<br>• DMFS and GRE steering: Rule insertion adjustments<br>• Removed DIF support from reported capabilities in QUERY_DEV_CAP PRM command<br>• Flow control by DSCP priority for IPv4<br>• DMFS improvements: Insertion scheme enforcement and block loopback for InfiniBand<br>• Added I2C resiliency support<br>• Support for NC-SI over MCTP over SMBus<br>• Added a flash access interface for persistent (non-volatile) configuration support<br>• Added port BW arbitration configuration through the CONFIG_DEV command<br>• Added IP-in-IP TCP checksum offload support<br>• PCI Express compliancy Tx and Rx adjustments<br>• 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.<br>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. |
| 2.30.3200 | • Added support for FDR AOC MC2207312-XXX<br>• Bug Fixes, see "Bug Fixes History" on page 24 |

*Table 10 - Firmware Changes and New Feature History*

| Firmware Version | Description |
|---|---|
| 2.30.3000 | • Added support for the UPDATE_QP PRM command<br>• Added support for resetting the modified MAC addresses in the standby mode by the MOD_STAT_CFG command<br>• Added support for receiving TCP and UDP truncated packets of certain type<br>• Added support for 56GbE in all devices supporting 40GbE<br>  • Establishing 56GbE link with SwitchX® requires 56GbE enablement on the switch<br>  • Establishing 56GbE link on back to back setup requires additions to the INI.<br>  • For further information, please contact Mellanox Support<br>• mlxconfig tool in not supported in this release<br>• RDMA/RoCE read retransmission improvement<br>• PFC thresholds improvements<br>• PCI speed_change flow improvement<br>• Added support for DIAG_RPRT per port<br>• Added PCIe Polling Compliance mode<br>• RoCE default configuration fixes:<br>  • ethertype now updated per port at SET_PORT<br>  • The default value of rroce.ip_next_protocol is 0xfe<br>• Increased the number of extended interface counters (max_if_cnt_extended) to 0x80, as reported in QUERY_DEV_CAP PRM command<br>• Improved link parallel detection calibration of 40GbE<br>• Added support for PFC counters in DUMP_ETH_STAT PRM command<br>• Fixed wrong reporting of RSS context in QUERY_FW of RSS QP<br>• Added Sniffer QP support on Port#1<br>• Device managed Flow Steering performance enhancements<br>• Disabled "pkt.dmac==qp.mac" for RoCE/RoCE over IP<br>• Added missing loopback blocking for device managed Flow Steering<br>• Fixed SET_PORT.mac_table configuration issue which caused minor packet loss on Port A when working in bonding mode and closing Port B.<br>• Fixed issues with NC-SI commands reason codes<br>• Fixed the insertion of L4 head rule in device managed Flow Steering<br>• Added to the INI for 10/40GbE parallel detect Serdes parameters<br>• Added support for "reset upon parity error"<br>• Added support for 40Gb/s MC2210411-SR4 optical module<br>• Fixed interoperability issue with the Intel 12300 switch using firmware version 6.1.0.1.11<br>• Improved QDR link stability when connected to InfiniScale® IV and SwitchX® switches |

*Table 10 - Firmware Changes and New Feature History*

| Firmware Version | Description |
|---|---|
| 2.11.0500 | • Added SR-IOV support<br>• Added VPI auto-sensing support<br>• Mellanox Link Property Negotiation (MLPN)<br>• Enables ports to negotiate link properties between Mellanox devices.<br>• The MLPN is activated by INI.<br>• CR4 + KR4<br>• Auto-negotiation 40GBASE CR4 and KR4 as described in IEEE 802.3.<br>• CR4 + KR4 is enabled by the INI.<br>• 1GbE Clause37<br>• Auto-negotiation 1000BASE-X as described in IEEE 802.3 clause 37.<br>• 1GbE Clause37 is enabled by default.<br>• cable_info MAD extension for more I2C addressed<br>• Enables different address for cable access through I2C.<br>• cable_info MAD extension is enabled by the INI.<br>• secure_host<br>• smp firewall described in the PRM<br>• Activated by the INI (active by default).<br>• cq_2_eq mapping command<br>• Modifies EQ by MODIFY_CQ command, described in the PRM<br>• multi-function reserved lkey<br>• Described in the PRM.<br>• increase CQE timestamp to 48bit<br>• Described in the PRM.<br>• 56Gb Ethernet (proprietary) - Beta level<br>• Activated by the INI (disabled by default).<br>• mlxconfig - Beta level (requires MFT 3.0.0-3 or above)<br>• Modifies the device cfg<br>• FMR for SRIOV - Beta level<br>• Described in the PRM.<br>• Power reduction in PCI Gen3<br>• Fixed general_info MAD<br>• "Bug Fixes History" on page 24 |
| 2.10.0800 | • Bug fixes - see "Bug Fixes History" on page 24 |
| 2.10.0700 | • Bug fixes - see "Bug Fixes History" on page 24 |
| 2.10.0000 | • InfiniBand:<br>  • FDR<br>  • FDR10<br>  • QDR<br>  • SDR<br>  • DDR<br>• Ethernet:<br>  • 1GigE<br>  • 10GigE XAUI<br>  • 10GigE XFI/SFI<br>  • 40GigE<br>• PCI Express 3.0, with backwards compatibility with v2.0 and v1.1<br>• Huge pages<br>• ConnectX®-3 firmware includes all ConnectX®-2 cards supported features |

# 6      Flexboot Changes and New Features

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

*Table 11 - FlexBoot Changes and New Feature*

| Version | Description |
|---------|-------------|
| Rev 3.4.648 | • Added support for .mrom images larger than 128kB<br>• Added boot over IB with non-default PKey for ConnectX®-3, ConnectX®-3 Pro cards<br>• Added support for ConnectX-4 and ConnectX-4 Lx<br>• Synced the source with iPXE (upstream sync)<br>• Moved to flat real mode when calling INT 1a,b101 to avoid BIOSes issues<br>• Fixed chainloading undionly.kpxe over Connect-IB functionality<br>• Fixed HTTP boot over IPoIB |
| Rev 3.4.521 | • Added iSCSI CHAP and mutual CHAP configuration<br>• Added the GRH size when allocating receive buffer for IPoIB<br>• Updated VLAN netdevice's settings with all the trunk's iSCSI required settings<br>• Updated the port event handling process<br>• Enabled console output in Debug mode<br>• Disabled the serial output<br>• Disabled the banner in BEV execution<br>• Disabled function 0x04 (in int21) when serial console is disabled<br>• Preserved COM port settings<br>• Fixed HTTP download over IPoIB<br>• Fixed completion with error handling process |

*Table 11 - FlexBoot Changes and New Feature*

| Version | Description |
|---|---|
| Rev 3.4.460 | • Boot Menu support: Added new FlexBoot GUI. The device can now be configured in the POST stage.<br>• Non volatile memory read/write support<br>• Configurable URI boot retry and delay between retries<br>• Configurable iSCSI settings using DHCP/NVM<br>• Added new interface in order to update the registered devices on the PXE stage<br>• Enabled ConnectX Ethernet adapter cards family to work with interrupts<br>• Enabled PXE to work in promiscuous VLAN mode (configurable through the INI)<br>• Synced version with ipxe.org: Now the latest code in iPXE is used<br>• Added boot priority capability: iSCSI vs PXE and fallback incase one fails<br>• 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.<br>Packets with source port different than BOOTPS_PORT and PXE_PORT are filtered by the PROXY<br>• SHELL CLI is currently supported on ConnectX-3 and ConnectX-3 Pro adapter cards only<br>• 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.<br>• Both the GUID and the MAC are printed on the screen when the port link layer is set as InfiniBand<br>• PROXYDHCP and PXEBS settings are saved under netdevice settings<br>• rootpath/filename/nextserver are now fetched from the netdevice settings<br>• The cached DHCP packet are received only if working with the same net device.<br>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, Netmask, etc...) is used.<br>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.<br>• PXE shutdown is called if int22 with function 0x000C is called.<br>• Changed DHCP discover timeouts to comply with PXE spec |

*Table 11 - FlexBoot Changes and New Feature*

| Version | Description |
|---------|-------------|
| Rev 3.4.306 | • Added validation script for the released ROMs<br>• Added the option to always keep SAN hook to enable WIN install on iSCSI target<br>• Added compilation flag around the flash readout.<br>• Added URI Boot retry. Default retries = 0.<br>• Added Unmap MPT command in teardown.<br>• Added support for HII ISCSI configuration.<br>• Added 64-bit PCI BAR support (Large bar).<br>• Added the option added for running PXE with promiscuous VLAN.<br>• Re-added COMBOOT image support by default.<br>• Enabled pages-function handling in Connect-IB initialization stage to work according to the PRM.<br>• Applied additional patches from ipxe.org<br>• Updated the window even if ACK does not acknowledge new data.<br>• Modified the error print to debug print.<br>• Modified the printed string when initializing devices.<br>• Modified the error print. Added additional information to make the output more user-friendly.<br>• Changed the size of the domain name array to 0xfd.<br>• Disabled the waiting period for link up on trunk-net-device when VLAN is enabled on port.<br>• Removed unsupported EQ event in Connect-IB®<br>• Fixed an issue for TLV with length 0.<br>• Fixed an issue related to sync VLAN IRQ operation with trunk IRQ operation.<br>• Fixed an issue which enabled a netdevice (VLAN) to open/close twice.<br>• Fixed an issue which prevented the iSCSI initiator's name from being received from HII.<br>• 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.<br>• Fixed an issue which caused PXE boot failure when using a filename if iSCSI rootpath is set.<br>• Fixed an issue which prevented the device to PXE boot from the 2nd port if first port was already downloaded.<br>• Fixed compilation issue.<br>• Fixed a broken VLAN issue.<br>• Fixed a retry issue when the value is infinite. |

*Table 11 - FlexBoot Changes and New Feature*

| Version | Description |
|---------|-------------|
| Rev 3.4.225 | • Added additional information to the error print output<br>• Added compilation flag around the flash readout<br>• Added URI Boot retry. Default retries = 0<br>• Added Unmap MPT command in teardown<br>• Added 64-bit PCI BAR support<br>• Added an option for running PXE with promiscuous VLAN<br>• Added support for HII iSCSI configuration<br>• Enlarged the mailbox size to 4kb<br>• Enlarged the number of WQE to 64 (from 4)<br>• Enabled multiple DHCP offers to be received before proceeding to request state<br>• Changed the size of the domain name array to 0xfd<br>• Changed error print to debug print<br>• Changed printed string when initializing devices<br>• Kept the SAN connection permanently open to enable Windows install on iSCSI target even when the iSCSI target is empty<br>• Re-added COMBOOT image support by default<br>• Prevented a netdevice (VLAN) from opening/closing twice<br>• Removed unsupported EQ event in Connect-IB®<br>• Disabled the waiting time for link up on trunk net device when VLAN is enabled on a port<br>• Fixed sync VLAN IRQ operation with trunk IRQ operation<br>• Fixed iSCSI initiator's name retrieval from HII issue<br>• 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.<br>• Fixed retry issue when the value is infinite<br>• Fixed a TLV with length 0 issue<br>• Fixed a PXE boot failure issue occurred when using a filename when iSCSI rootpath is set<br>• Fixed "Impossible to PXE boot from 2nd port if first port already down-loaded." issue<br>• Fixed compilation issue<br>• Fixed broken VLAN support issues |
| Rev 3.4.151 | • Enlarged the mailbox size to 4kb<br>• Enlarged the number of WQE to 64 (from 4)<br>• Enabled multiple DHCP offers to be received before proceeding to request state |
| Rev 3.4.146 | • Fixed memory corruption issues<br>• Modified TLV flash access<br>• Added additional WQ |
| Rev 3.4.142 | • Enabled firmware to handle the link state with the Subnet Manager<br>• Updated the DHCP class code to NONE<br>• Added flash access capability for reading software-to-software configurations<br>• Enabled DHCP validation of MAC address and XID for a unique tuple<br>• Improved randomness algorithm for DHCP XID |

*Table 11 - FlexBoot Changes and New Feature*

| Version | Description |
|---------|-------------|
| Rev 3.4.112 | • Broadcast responses for firewall support<br>• Enabled request broadcast responses from DHCP server to support firewall. |
| Rev 3.4.100 | • OCSD activation initiation change<br>• Moved the OCSD activation initiation from the FlexBoot to the CLP code. This enables the OCSD<br>• activation to no longer be dependent on the FlexBoot being enabled in the servers's BIOS configuration.<br>• Messages' improvement<br>• Made the FlexBoot on-screen notification messages more informative and user friendly.<br>• FlexBoot and CLP merge improvement<br>• Improved the process of merging the FlexBoot and CLP codes together.<br>• PXE and UFI merge capability<br>• Added the ability to merge the PXE image with a UFI image.<br>• Supported servers<br>• Added FlexBoot support capabilities to several new non-HP servers.<br>• Use of newer iPXE version<br>• Moved to use a newer iPXE version as the basis for the Flexboot release.<br>• Fixed "no more network devices" issues during Flexboot. |