



# Mellanox ConnectX<sup>®</sup>-5 Firmware Release Notes

---

Rev 16.26.1040

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT ("PRODUCT(S)") AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES "AS-IS" WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER'S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



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

© Copyright 2019. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, Connect-IB®, ConnectX®, CORE-Direct®, GPUDirect®, LinkX®, Mellanox Multi-Host®, Mellanox Socket Direct®, UFM®, and Virtual Protocol Interconnect® are registered trademarks of Mellanox Technologies, Ltd.

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

All other trademarks are property of their respective owners.

# Table of Contents

|   |           |
|---|-----------|
| <b>Release Update History</b> .....                               | <b>6</b>  |
| <b>Chapter 1 Overview</b> .....                                   | <b>7</b>  |
| 1.1 Supported Devices .....                                       | 7         |
| 1.2 Supported Cables and Modules .....                            | 10        |
| 1.2.1 Validated and Supported 1GbE Cables .....                   | 10        |
| 1.2.2 Validated and Supported 10GbE Cables .....                  | 10        |
| 1.2.3 Validated and Supported 25GbE Cables .....                  | 12        |
| 1.2.4 Validated and Supported 40GbE Cables .....                  | 13        |
| 1.2.5 Validated and Supported 50GbE Cables .....                  | 13        |
| 1.2.6 Validated and Supported 100GbE Cables .....                 | 14        |
| 1.2.7 Validated and Supported 200GbE Cables .....                 | 15        |
| 1.2.8 Validated and Supported QDR Cables .....                    | 15        |
| 1.2.9 Validated and Supported FDR10 Cables .....                  | 15        |
| 1.2.10 Validated and Supported FDR Cables .....                   | 16        |
| 1.2.11 Validated and Supported EDR / 100Gb/s Cables .....         | 16        |
| 1.2.12 Validated and Supported HDR / 200Gb/s Cables .....         | 17        |
| 1.3 Tested Switches .....   | 18        |
| 1.3.1 Tested 10GbE Switches .....                                 | 18        |
| 1.3.2 Tested 40GbE Switches .....                                 | 18        |
| 1.3.3 Tested 100GbE Switches .....                                | 18        |
| 1.3.4 Tested QDR Switches .....                                   | 19        |
| 1.3.5 Tested FDR Switches .....                                   | 19        |
| 1.3.6 Tested EDR / 100Gb/s Switches .....                         | 19        |
| 1.4 Tools, Switch Firmware and Driver Software .....              | 20        |
| 1.5 Supported FlexBoot, UEFI .....                                | 21        |
| 1.6 Revision Compatibility .....                                  | 21        |
| <b>Chapter 2 Changes and New Features in Rev 16.26.1040</b> ..... | <b>22</b> |
| <b>Chapter 3 Known Issues</b> .....                               | <b>23</b> |
| <b>Chapter 4 Bug Fixes History</b> .....                          | <b>30</b> |
| <b>Chapter 5 Firmware Changes and New Feature History</b> .....   | <b>35</b> |
| <b>Chapter 6 PreBoot Drivers (FlexBoot/UEFI)</b> .....            | <b>49</b> |
| 6.1 FlexBoot Changes and New Features .....                       | 49        |
| 6.2 UEFI Changes and Major New Features .....                     | 49        |
| <b>Chapter 7 Unsupported Features and Commands</b> .....          | <b>50</b> |
| 7.1 Unsupported Features .....                                    | 50        |
| 7.2 Unsupported Commands .....                                    | 50        |

## Chapter 8 Supported Non-Volatile Configurations . . . . . 51

## List of Tables

|           |   |    |
|-----------|---|----|
| Table 1:  | Release Update History                                    | 6  |
| Table 2:  | Supported Devices   | 7  |
| Table 3:  | Validated and Supported 1GbE Cables                       | 10 |
| Table 4:  | Validated and Supported 10GbE Cables                      | 10 |
| Table 5:  | Validated and Supported 25GbE Cables                      | 12 |
| Table 6:  | Validated and Supported 40GbE Cables                      | 13 |
| Table 7:  | Validated and Supported 50GbE Cables                      | 13 |
| Table 8:  | Validated and Supported 100GbE Cables                     | 14 |
| Table 9:  | Validated and Supported 200GbE Cables                     | 15 |
| Table 10: | Validated and Supported QDR Cables                        | 15 |
| Table 11: | Validated and Supported FDR10 Cables                      | 15 |
| Table 12: | Validated and Supported FDR Cables                        | 16 |
| Table 13: | Validated and Supported EDR / 100Gb/s Cables              | 16 |
| Table 14: | Validated and Supported HDR Cables                        | 17 |
| Table 15: | Tested 10GbE Switches                                     | 18 |
| Table 16: | Tested 40GbE Switches                                     | 18 |
| Table 17: | Tested 100GbE Switches                                    | 18 |
| Table 18: | Tested QDR Switches                                       | 19 |
| Table 19: | Tested FDR Switches                                       | 19 |
| Table 20: | Tested EDR Switches                                       | 19 |
| Table 21: | Tools, Switch Firmware and Driver Software                | 20 |
| Table 22: | Supported FlexBoot, UEFI                                  | 21 |
| Table 23: | Changes and New Features in Rev 16.26.1040                | 22 |
| Table 24: | Ethernet Rate Limit per VF in RoCE Mode Limitations       | 23 |
| Table 25: | Ethernet Rate Limit per VF in InfiniBand Mode Limitations | 23 |
| Table 26: | Known Issues  | 23 |
| Table 27: | Bug Fixes History   | 30 |
| Table 28: | Firmware Changes and New Feature History                  | 35 |
| Table 29: | Unsupported Features                                      | 50 |
| Table 30: | Supported Non-Volatile Configurations                     | 51 |

## Release Update History

*Table 1 - Release Update History*

| Release        | Date               | Description   |
|----------------|--------------------|---|
| Rev 16.26.1040 | September 29, 2019 | Initial version of this firmware release.<br>This version introduces New Features <a href="#">Section 2, “Changes and New Features in Rev 16.26.1040”</a> , on page 22 and Bug Fixes (see <a href="#">Section 4, “Bug Fixes History”</a> , on page 30). |

# 1 Overview

These are the release notes for the ConnectX®-5 adapters firmware Rev 16.26.1040.

This firmware supports the following protocols:

- InfiniBand - SDR, QDR, FDR10, FDR, EDR
- Ethernet - 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 56GbE<sup>1</sup>, 100GbE
- PCI Express 4.0, supporting backwards compatibility for v3.0, v2.0 and v1.1

## 1.1 Supported Devices

This firmware supports the devices and protocols listed in [Table 2](#)

**Table 2 - Supported Devices (Sheet 1 of 4)**

| Device Part Number | PSID          | Device Name  | FlexBoot          | UEFI x86           | UEFI ARM           | Enable/disable exprom Feature |
|--------------------|---------------|--|-------------------|--------------------|--------------------|-------------------------------|
| MCX512A-ACUT       | MT_0000000425 | ConnectX®-5 EN network interface card, 10/25GbE dual-port SFP28, PCIe3.0 x8, UEFI Enabled (x86/ARM), tall bracket  | Present (Enabled) | Present (Enabled)  | Present (Enabled)  | Exists                        |
| MCX512A-ADAT       | MT_0000000361 | ConnectX®-5 Ex EN network interface card, 25GbE dual-port SFP28, PCIe3.0/4.0 x8, tall bracket  | Present (Enabled) | Present (Disabled) | Not Present        | Exists                        |
| MCX562A-ACAB       | MT_0000000241 | ConnectX®-5 EN network interface card for OCP 3.0, with host management, 25GbE Dual-port SFP28, PCIe3.0 x16, Thumbscrew (Pull Tab) bracket                         | Present (Enabled) | Present (Enabled)  | Not Present        | Exists                        |
| MCX4121A-ACUT      | MT_0000000266 | ConnectX®-4 Lx EN network interface card, 25GbE dual-port SFP28, PCIe3.0 x8, UEFI Enabled, tall bracket  | Present (Enabled) | Present (Enabled)  | Present (Disabled) | Exists                        |
| MCX566A-CDA        | MT_0000000242 | ConnectX®-5 Ex EN network interface card for OCP 3.0, with host management, 100GbE Dual-port QSFP28, PCIe4.0 x16, Internal Lock bracket                            | Present (Enabled) | Present (Disabled) | Not Present        | Exists                        |
| MCX545B-CCUN       | MT_0000000419 | ConnectX®-5 EN network interface card for OCP2.0, Type 1, with host management, 100GbE, single-port QSFP28, PCIe3.0 x16, UEFI Enabled, no bracket                  | Present (Enabled) | Present (Enabled)  | Not Present        | Exists                        |
| MCX542B-ACUN       | MT_0000000427 | ConnectX®-5 EN network interface card for OCP2.0, Type 1, with host management, 25GbE dual-port SFP28, PCIe3.0 x8, UEFI Enabled (x86/ARM), no bracket Halogen free | Present (Enabled) | Present (Enabled)  | Present (Enabled)  | Exists                        |

1. 56GbE 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 Devices (Sheet 2 of 4)**

| Device Part Number | PSID          | Device Name   | FlexBoot          | UEFI x86           | UEFI ARM    | Enable/disable exprom Feature |
|--------------------|---------------|---|-------------------|--------------------|-------------|-------------------------------|
| MCX512F-ACHT       | MT_0000000416 | ConnectX®-5 EN network interface card, with host management, 25GbE Dual-port SFP28, PCIe3.0 x16, UEFI Enabled, tall bracket                       | Present (Enabled) | Present (Enabled)  | Not Present | Exists                        |
| MCX545A-CCUN       | MT_0000000418 | ConnectX®-5 EN network interface card for OCP2.0, Type 2, with host management, 100GbE, single-port QSFP28, PCIe3.0 x16, UEFI Enabled, no bracket | Present (Enabled) | Present (Enabled)  | Not Present | Exists                        |
| MCX516A-CCHT       | MT_0000000417 | ConnectX®-5 EN network interface card, with host management 100GbE dual-port QSFP28, PCIe3.0 x16, UEFI Enabled, tall bracket                      | Present (Enabled) | Present (Enabled)  | Not Present | Exists                        |
| MCX566A-CCAI       | MT_0000000348 | ConnectX®-5 EN network interface card for OCP 3.0, with host management, 100GbE Dual-port QSFP28, PCIe3.0 x16, Internal Lock bracket              | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX512A-ACAT       | MT_0000000080 | ConnectX®-5 EN network interface card, 10/25GbE dual-port SFP28, PCIe3.0 x8, tall bracket, ROHS R6  | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX515A-CCAT       | MT_0000000011 | ConnectX-5 EN network interface card, 100GbE single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6   | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX515A-GCAT       | MT_0000000087 | ConnectX®-5 EN network interface card, 50GbE single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6   | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX516A-BDAT       | MT_0000000123 | ConnectX®-5 Ex EN network interface card, 40GbE dual-port QSFP28, PCIe 4.0 x16, tall bracket, ROHS R6   | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX516A-CCAT       | MT_0000000012 | ConnectX-5 EN network interface card, 100GbE dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6   | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX516A-CDAT       | MT_0000000013 | ConnectX-5 Ex EN network interface card, 100GbE dual-port QSFP28, PCIe4.0 x16, tall bracket, ROHS R6  | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX546A-BCAN       | MT_0000000069 | ConnectX®-5 EN network interface card for OCP, 40GbE dual-port QSFP28, PCIe3.0 x16, no bracket, ROHS R6   | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX546A-CDAN       | MT_0000000058 | ConnectX-5 Ex network interface card for OCP; 100GbE dual-port QSFP28; PCIe4.0 x16; no bracket; ROHS R6   | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX545A-CCAN       | MT_0000000157 | ConnectX-5 EN network interface card for OCP 100GbE; single-port QSFP28; PCIe3.0 x16; no bracket; ROHS R6;  | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |



**Table 2 - Supported Devices (Sheet 3 of 4)**

| Device Part Number | PSID          | Device Name   | FlexBoot          | UEFI x86           | UEFI ARM    | Enable/disable exprom Feature |
|--------------------|---------------|---|-------------------|--------------------|-------------|-------------------------------|
| MCX512F-ACAT       | MT_0000000183 | ConnectX®-5 EN network interface card, 25GbE Dual-port SFP28, PCIe3.0 x16, tall bracket, ROHS R6  | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX511F-ACA        | MT_0000000182 | ConnectX-5 EN network interface card; 25GbE single-port SFP28; PCIe4.0 x16; ROHS R6   | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX542B-ACA        | MT_0000000248 | ConnectX-5 EN network interface card for OCP; with host management; 25GbE dual-port SFP28; PCIe3.0 x8; no bracket; ROHS R6 Halogen free                                       | Present (Enabled) | Present (Enabled)  | Not Present | Exists                        |
| MCX542A-ACAN       | MT_0000000167 | ConnectX®-5 EN network interface card for OCP, with host management, 25GbE dual-port SFP28, PCIe3.0 x16, no bracket, ROHS R6 Halogen free                                     | Present (Enabled) | Present (Disabled) | Not Present | Not Present                   |
| MCX516A-GCAT       | MT_0000000090 | ConnectX®-5 EN network interface card, 50GbE dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6   | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX553Q-ECAS       | MT_0000000309 | ConnectX®-5 VPI adapter card with Multi-Host, EDR IB (100Gb/s) and 100GbE, Single-port QSFP28, PCIe3.0 x4 on board, external connectors to 3x auxiliary cards?, Short bracket | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX545A-ECAN       | MT_0000000077 | ConnectX®-5 VPI network interface card for OCP EDR IB (100Gb/s) and 100GbE, single-port QSFP28, PCIe3.0 x16, no bracket, ROHS R6  | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX545A-ECA_BX     | MT_0000000092 | ConnectX-5 VPI network interface card for OCP EDR IB (100Gbs) and 100GbE; single-port QSFP28; PCIe3.0 x16; no bracket; ROHS R6; Halogen free                                  | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX545B-ECAN       | MT_0000000207 | ConnectX-5 VPI network interface card for OCP; with host management; EDR IB (100Gb/s) and 100GbE; single-port QSFP28; PCIe3.0 x16; no bracket; 8mm Heat Sink; ROHS R6         | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX555A-ECAT       | MT_0000000010 | ConnectX-5 VPI adapter card, EDR IB (100Gb/s) and 100GbE, single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6  | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX556A-ECAT       | MT_0000000008 | ConnectX-5 VPI adapter card, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6  | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |

**Table 2 - Supported Devices (Sheet 4 of 4)**

| Device Part Number | PSID          | Device Name   | FlexBoot          | UEFI x86           | UEFI ARM    | Enable/disable exprom Feature |
|--------------------|---------------|---|-------------------|--------------------|-------------|-------------------------------|
| MCX556A-EDAT       | MT_0000000009 | ConnectX-5 Ex VPI adapter card, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, PCIe4.0 x16, tall bracket, ROHS R6   | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX556M-ECAT-S25   | MT_0000000023 | ConnectX@-5 VPI adapter card with Multi-Host Socket Direct supporting dual-socket server, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, 2x PCIe3.0 x8, 25cm harness, tall bracket, ROHS R6 | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |
| MCX546A-EDAN       | MT_0000000135 | ConnectX-5 VPI network interface card for OCP; EDR IB (100Gb/s) and 100GbE dual-port QSFP28; PCIe4.0 x16; no bracket; ROHS R6   | Present (Enabled) | Present (Disabled) | Not Present | Exists                        |

## 1.2 Supported Cables and Modules

Please refer to the LinkX® Cables and Transceivers web page

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

### 1.2.1 Validated and Supported 1GbE Cables

**Table 3 - Validated and Supported 1GbE Cables**

| Speed | Cable OPN #  | Description  |
|-------|--------------|--|
| 1GB/S | MC3208011-SX | Mellanox Optical module, SX, 850nm                         |
| 1GB/S | MC3208411-T  | Mellanox® module, ETH 1GbE, 1Gb/s, SFP, Base-T, up to 100m |

### 1.2.2 Validated and Supported 10GbE Cables

**Table 4 - Validated and Supported 10GbE Cables**

| Speed | Cable OPN #      | Description  |
|-------|------------------|--|
| 10GbE | BN-QS-SP-CBL-5M  | 40G QSFP+ to 4xSFP+ DAC Breakout Direct Attach Cable 5m        |
| 10GbE | BN-QS-SP-CBL-5M  | 40G QSFP+ to 4xSFP+ DAC Breakout Direct Attach Cable 5m        |
| 10GbE | CAB-SFP-SFP-1M   | Arista 10GBASE-CR SFP+ Cable 1 Meter                           |
| 10GbE | CAB-SFP-SFP-3M   | Arista 10GBASE-CR SFP+ Cable 3 Meter                           |
| 10GbE | CAB-SFP-SFP-5M   | Arista 10GBASE-CR SFP+ Cable 5 Meter                           |
| 10GbE | FTLX1471D3BCL-ME | 10GBASE-LR SFP+ 1310nm 10km DOM Transceiver Module             |
| 10GbE | L45593-D178-B50  | QSFP-4SFP10G-CU5M  |
| 10GbE | MC2309124-004    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 4M |
| 10GbE | MC2309124-005    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 5M |

**Table 4 - Validated and Supported 10GbE Cables**

| Speed | Cable OPN #      | Description  |
|-------|------------------|--|
| 10GbE | MC2309130-001    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 1M                           |
| 10GbE | MC2309130-002    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 2M                           |
| 10GbE | MC2309130-003    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 3M                           |
| 10GbE | MC2309130-00A    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 0.5M                         |
| 10GbE | MC2609125-004    | Mellanox Passive Copper Hybrid Cable ETH 40GbE TO 4X10GBE QSFP TO 4X SFP+ 4M             |
| 10GbE | MC2609130-001    | Mellanox Passive Copper Hybrid Cable ETH 40GbE TO 4X10GBE QSFP TO 4X SFP+ 1M             |
| 10GbE | MC2609130-002    | Mellanox Passive Copper Hybrid Cable ETH 40GbE TO 4X10GBE QSFP TO 4X SFP+ 2M             |
| 10GbE | MC2609130-003    | Mellanox Passive Copper Hybrid Cable ETH 40GbE TO 4X10GBE QSFP TO 4X SFP+ 3M             |
| 10GbE | MC2609130-0A1    | Mellanox Passive Copper Hybrid Cable ETH 40GbE TO 4X10GBE QSFP TO 4X SFP+ 1.5M           |
| 10GbE | MC3309124-004    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 4M                                   |
| 10GbE | MC3309124-005    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 5M                                   |
| 10GbE | MC3309124-006    | Mellanox® Passive Copper Cable, ETH 10GbE, 10Gb/s, SFP+, 6m                              |
| 10GbE | MC3309124-007    | Mellanox® Passive Copper Cable, ETH 10GbE, 10Gb/s, SFP+, 7m                              |
| 10GbE | MC3309130-001    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 1M                                   |
| 10GbE | MC3309130-002    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 2M                                   |
| 10GbE | MC3309130-003    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 3M                                   |
| 10GbE | MC3309130-004    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 4M                                   |
| 10GbE | MC3309130-005    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 5M                                   |
| 10GbE | MC3309130-006    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 6M                                   |
| 10GbE | MC3309130-007    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 7M                                   |
| 10GbE | MC3309130-00A    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 0.5M                                 |
| 10GbE | MC3309130-0A1    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 1.5M                                 |
| 10GbE | MC3309130-0A2    | Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 2.5M                                 |
| 10GbE | MFMT102A-LR-F    | Mellanox Optical Module ETH 10GBE 10GB/S SFP+ LC-LC 1310NM LR up to 10KM                 |
| 10GbE | SFP-10G-SR       | Cisco 10GBASE-SR SFP+ transceiver module for MMF, 850-nm wavelength, LC duplex connector |
| 10GbE | SFP-H10GB-CU1M   | Cisco 1-m 10G SFP+ Twinax cable assembly, passive  |
| 10GbE | SFP-H10GB-CU3M   | Cisco 3-m 10G SFP+ Twinax cable assembly, passive  |
| 10GbE | SFP-H10GB-CU5M   | Cisco 5-m 10G SFP+ Twinax cable assembly, passive  |
| 10GbE | FTLX8571D3BCL-ME | 10gb SFP 850nm Optic Transceiver   |

### 1.2.3 Validated and Supported 25GbE Cables



The 25GbE cables can be supported in ConnectX-5 adapter cards only when connected to the MAM1Q00A-QSA28 module.

**Table 5 - Validated and Supported 25GbE Cables**

| Speed | Cable OPN #      | Description   |
|-------|------------------|---|
| 25GbE | MCP2M00-A001     | Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1m                              |
| 25GbE | MCP2M00-A002     | Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2m                              |
| 25GbE | MCP2M00-A003     | Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 3m                              |
| 25GbE | MCP2M00-A003AP   | Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 3m, 26AWG                       |
| 25GbE | MCP2M00-A005E26L | Mellanox Passive Copper Cable, ETH, up to 25GB/S, SFP28, 5M, black, 26AWG, CA-L           |
| 25GbE | MCP2M00-A00A     | Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 0.5m                            |
| 25GbE | MCP2M00-A01A     | Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1.5m                            |
| 25GbE | MCP2M00-A02A     | Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2.5m                            |
| 25GbE | MCP7F00-A001     | Mellanox Passive Copper Hybrid cable ETH 100GbE to 4X25GBS QSFP28 to 4XSFP28 1M           |
| 25GbE | MCP7F00-A002     | Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 4X25GBS QSFP28 TO 4XSFP28 2M           |
| 25GbE | MCP7F00-A003     | Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 4X25GBS QSFP28 TO 4XSFP28 3M           |
| 25GbE | MCP7F00-A003-AM  | Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 3M 30AWG |
| 25GbE | MCP7F00-A005AM   | Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 4X25GBS QSFP28 to 4XSFP28 5M           |
| 25GbE | MCP7F00-A01A     | Mellanox Passive Copper Hybrid Cable ETH 100GbE to 4X25GBS QSFP28 to 4XSFP28 1.5M         |
| 25GbE | MCP7F00-A02A     | Mellanox Passive Copper Hybrid Cable ETH 100GbE to 4X25GBS QSFP28 to 4XSFP28 2.5M         |
| 25GbE | MFA7A50-CXXX     | Mellanox Active Fiber Hybrid Solution ETH 100GbE to 4X25GBE QSFP28 to 4XSFP28 up to 30M   |
| 25GbE | MMA2L20-AR       | Mellanox® optical transceiver, 25GbE, 25Gb/s, SFP28, LC-LC, 1310nm, LR up to 10km         |
| 25GbE | MMA2P00-AS       | Mellanox® transceiver, 25GbE, SFP28, LC-LC, 850nm, SR, up to 100m                         |
| 25GbE | MFM1T02A-SR-P    | Mellanox® Optical Module ETH 10GbE 10GB/S SFP+ LC-LC 850NM SR up to 300M                  |
| 25GbE | SFP-H25G-CU1M    | 25GBASE-CR1 Copper Cable 1-meter  |
| 25GbE | SFP-H25G-CU2M    | 25GBASE-CR1 Copper Cable 2-meter  |

## 1.2.4 Validated and Supported 40GbE Cables

**Table 6 - Validated and Supported 40GbE Cables**

| Speed | Cable OPN #      | Description   |
|-------|------------------|---|
| NA    | MAM1Q00A-QSA     | Mellanox® cable module, ETH 10GbE, 40Gb/s to 10Gb/s, QSFP to SFP+     |
| NA    | MAM1Q00A-QSA28   | Mellanox® cable module, ETH 25GbE, 100Gb/s to 25Gb/s, QSFP28 to SFP28 |
| 40GbE | MC2210126-004    | Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 4m            |
| 40GbE | MC2210126-005    | Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 5m            |
| 40GbE | MC2210128-003    | Mellanox Passive Copper Cable ETH 40GbE 40GbE QSFP 3M                 |
| 40GbE | MC2210130-001    | Mellanox Passive Copper Cable ETH 40GbE 40GbE QSFP 1M                 |
| 40GbE | MC2210130-002    | Mellanox Passive Copper Cable ETH 40GbE 40GbE QSFP 2M                 |
| 40GbE | MC2210130-00A    | Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 0.5m          |
| 40GbE | MC2210130-00B    | Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 0.75m         |
| 40GbE | MC2210310-XXX    | Mellanox Active Fiber Cable ETH 40GbE 40GbE QSFP from 3M up to 100M   |
| 40GbE | MC2210411-SR4    | Mellanox Optical Module 40GbE QSFP MPO 850NM UP TO 100M               |
| 40GbE | MC2210411-SR4E   | Mellanox Optical Module 40GbE QSFP MPO 850NM UP TO 300M               |
| 40GbE | QSFP-40G-SR-BD   | Cisco 40GBASE-SR-BiDi, duplex MMF                                     |
| 40GbE | QSFP-40G-SR4     | Cisco 40GBASE-SR4, 4 lanes, 850 nm MMF                                |
| 40GbE | QSFP-H40G-ACU10M | Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 10-meter, active   |
| 40GbE | QSFP-H40G-AOC10M | Cisco 40GBase-AOC QSFP direct-attach Active Optical Cable, 10-meter   |

## 1.2.5 Validated and Supported 50GbE Cables

**Table 7 - Validated and Supported 50GbE Cables**

| Speed | Cable OPN #  | Description  |
|-------|--------------|--|
| 50GbE | MCP7H00-G001 | Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 2X50GBS QSFP28 TO 2XQSFP28 1M       |
| 50GbE | MCP7H00-G002 | Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 2X50GBS QSFP28 TO 2XQSFP28 2M       |
| 50GbE | MCP7H00-G003 | Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 2X50GBS QSFP28 TO 2XQSFP28 3M       |
| 50GbE | MCP7H00-G01A | Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 2X50GBS QSFP28 TO 2XQSFP28 1.5M     |
| 50GbE | MCP7H00-G02A | Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 2X50GBS QSFP28 TO 2XQSFP28 2.5M     |
| 50GbE | MFA7A20-C020 | Mellanox® active fiber hybrid solution, ETH 100GbE to 2x50GbE, QSFP28 to 2xQSFP28, 20m |

## 1.2.6 Validated and Supported 100GbE Cables

**Table 8 - Validated and Supported 100GbE Cables**

| Speed  | Cable OPN #       | Description  |
|--------|-------------------|--|
| 100GbE | CAB-Q-Q-100GbE-3M | Passive 3 meter , QSFP+ to QSFP+ QSFP100 TWINAX 103.125Gbps-CR4  |
| 100GbE | MCP1600-C001      | Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 1M   |
| 100GbE | MCP1600-C002      | Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 2M   |
| 100GbE | MCP1600-C003      | Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 3M   |
| 100GbE | MCP1600-C005AM    | Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP, 5m, 26AWG  |
| 100GbE | MCP1600-C005E26L  | Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP28, 5m, Black, 26AWG, CA-L                         |
| 100GbE | MCP1600-C00A      | Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 0.5M   |
| 100GbE | MCP1600-C01A      | Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP, LSZH, 1.5m                                       |
| 100GbE | MCP1600-C02A      | Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP, LSZH, 2.5m                                       |
| 100GbE | MCP1600-C03A      | Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP, PVC, 3.5m 26AWG                                  |
| 100GbE | MCP7F00-A005R26L  | Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 5m, Colored, 26AWG, CA-L  |
| 100GbE | MCP7H00-G005R26L  | Mellanox® passive copper hybrid cable, ETH 100GbE to 2x50GbE, QSFP28 to 2xQSFP28, 5m, Colored, 26AWG, CA-L |
| 100GbE | MFA1A00-C003      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 3m   |
| 100GbE | MFA1A00-C005      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 5m   |
| 100GbE | MFA1A00-C010      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 10m  |
| 100GbE | MFA1A00-C015      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 15m  |
| 100GbE | MFA1A00-C020      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 20m  |
| 100GbE | MFA1A00-C030      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 30m  |
| 100GbE | MFA1A00-C050      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 50m  |
| 100GbE | MFA1A00-C100      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 100m   |
| 100GbE | MFS1200-C005      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 5m   |
| 100GbE | MFS1200-C010      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 10m  |
| 100GbE | MFS1200-C015      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 15m  |
| 100GbE | MFS1200-C020      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 20m  |
| 100GbE | MFS1200-C030      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 30m  |
| 100GbE | MFS1200-C050      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 50m  |
| 100GbE | MFS1200-C100      | Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 100m   |
| 100GbE | MMA1B00-C100_B    | Mellanox® transceiver, up to 100GbE, QSFP28, MPO, 850nm, up to 100m OM3                                    |
| 100GbE | MMA1B00-C100D     | Mellanox® Transceiver, 100GbE, QSFP28, MPO, 850nm, up to 100m  |
| 100GbE | MMS1C10-CM        | Mellanox® active optical module, 100Gb/s, QSFP, MPO, 1310nm, PSM4  |
| 100GbE | MMA1L30-CM        | Mellanox® optical module, 100GbE, 100Gb/s, QSFP28, LC-LC, 1310nm, CWDM4, up to 2km                         |
| 100GbE | MMS1C00-C500      | Mellanox® transceiver, 100GbE, QSFP28, MPO, 1550nm PSM4, up to 2km   |
| 100GbE | MMS1C00-C500      | Mellanox® transceiver, 100GbE, QSFP28, MPO, 1550nm PSM4, up to 2km   |
| 100GbE | QSFP-40/100-SRBD  | 100GbE and 40GBASE SR-BiDi QSFP Transceiver, LC, 100m OM4 MMF  |

**Table 8 - Validated and Supported 100GbE Cables**

| Speed  | Cable OPN #      | Description   |
|--------|------------------|---|
| 100GbE | TR-FC13L-N00     | 100G QSFP28 Optical Transceivers, QSFP28 LR4 (10km)   |
| 100GbE | SO-QSFP28-LR4    | QSFP28, 100GBase, 1310nm, SM, DDM, 10km, LC   |
| 100GbE | 10137499-4050LF  | HP 5m splitter 100G to 4X25G  |
| 100GbE | FTLF8519P3BTL-N1 | 1000BASE-SX and 2G Fibre Channel (2GFC) 500m Industrial Temperature SFP Optical Transceiver |

## 1.2.7 Validated and Supported 200GbE Cables

**Table 9 - Validated and Supported 200GbE Cables**

| Speed  | Cable OPN #     | Description  |
|--------|-----------------|--|
| 200GbE | MCP1650-V001E30 | Mellanox® Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 1m, black pultab, 30AWG                     |
| 200GbE | MCP1650-V002E26 | Mellanox® Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 2m, black pultab, 26AWG                     |
| 200GbE | MCP1650-V003E26 | Mellanox® Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 3m, black pultab, 26AWG                     |
| 200GbE | MCP1650-V00AE30 | Mellanox® Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 0.5m, black pultab, 30AWG                   |
| 200GbE | MCP1650-V02AE26 | Mellanox® Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 2.5m, black pultab, 26AWG                   |
| 200GbE | MCP7H50-V003R26 | Mellanox® Passive Copper hybrid cable, 200GbE 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, colored, 3m, 26AWG |

## 1.2.8 Validated and Supported QDR Cables

**Table 10 - Validated and Supported QDR Cables**

| Speed | Cable OPN #   | Description   |
|-------|---------------|---|
| QDR   | MC2206125-007 | Mellanox Passive Copper Cable IB QDR 40GB/S QSFP 7M |
| QDR   | MC2206126-006 | Mellanox Passive Copper Cable IB QDR 40GB/S QSFP 6M |

## 1.2.9 Validated and Supported FDR10 Cables

**Table 11 - Validated and Supported FDR10 Cables**

| Speed | Cable OPN #   | Description  |
|-------|---------------|--|
| FDR10 | MC2206128-004 | Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 4M   |
| FDR10 | MC2206128-005 | Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 5M   |
| FDR10 | MC2206130-001 | Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 1M   |
| FDR10 | MC2206130-002 | Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 2M   |
| FDR10 | MC2206130-003 | Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 3M   |
| FDR10 | MC2206130-00A | Mellanox Passive Copper Cable VPI UP TO 40GB/S QSFP 0.5M |

**Table 11 - Validated and Supported FDR10 Cables**

| Speed | Cable OPN #   | Description   |
|-------|---------------|---|
| FDR10 | MC2206310-XXX | Mellanox Active Fiber Cable IB QDR/FDR10 40GB/S QSFP from 3M up to 100M |
| FDR10 | MFS4R12CB-XXX | Mellanox Active Fiber Cable VPI UP TO 40GB/S QSFP from 3M up to 100M    |

## 1.2.10 Validated and Supported FDR Cables

**Table 12 - Validated and Supported FDR Cables**

| Speed | Cable OPN #    | Description  |
|-------|----------------|--|
| FDR   | MC2207126-004  | Mellanox® Passive Copper Cable, VPI, up to 56Gb/s, QSFP, 4m              |
| FDR   | MC2207128-003  | Mellanox® Passive Copper Cable, VPI, up to 56Gb/s, QSFP, 3m              |
| FDR   | MC2207130-001  | Mellanox® Passive Copper Cable, VPI, up to 56Gb/s, QSFP, 1m              |
| FDR   | MC2207130-00A  | Mellanox® Passive Copper Cable, VPI, up to 56Gb/s, QSFP, 0.5m            |
| FDR   | MC2207310-003  | Mellanox® Active Fiber Cable, VPI, up to 56Gb/s, QSFP, 3m                |
| FDR   | MC2207310-010  | Mellanox® Active Fiber Cable, VPI, up to 56Gb/s, QSFP, 10m               |
| FDR   | MC2207310-015  | Mellanox® Active Fiber Cable, VPI, up to 56Gb/s, QSFP, 15m               |
| FDR   | MC2207310-100  | Mellanox® Active Fiber Cable, VPI, up to 56Gb/s, QSFP, 100m              |
| FDR   | MC2207312-XXX  | Mellanox® Active Fiber Cable, VPI, up to 56Gb/s, QSFP, up to 100m        |
| FDR   | MC220731V-XXX  | Mellanox® Active Fiber cable, VPI, up to 56Gb/s, QSFP, up to 100m        |
| FDR   | MC2207411-SR4L | Mellanox® Optical Module, VPI, up to 56Gb/s, QSFP, MPO, 850nm, up to 30m |

## 1.2.11 Validated and Supported EDR / 100Gb/s Cables

**Table 13 - Validated and Supported EDR / 100Gb/s Cables**

| Speed | Cable OPN #     | Description  |
|-------|-----------------|--|
| EDR   | MCP1600-E001    | Mellanox Passive Copper Cable VPI 100Gb/s QSFP LSZH 1M                         |
| EDR   | MCP1600-E002    | Mellanox Passive Copper Cable VPI 100Gb/s QSFP LSZH 2M                         |
| EDR   | MCP1600-E003    | Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 3M                         |
| EDR   | MCP1600-E004A26 | Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 4m, Blue, 26AWG |
| EDR   | MCP1600-E005    | Mellanox Passive Copper Cable VPI 100GB/S QSFP LSZH 5M                         |
| EDR   | MCP1600-E00A    | Mellanox Passive Copper Cable VPI 100Gb/s QSFP LSZH 0.5M                       |
| EDR   | MCP1600-E01A    | Mellanox® Passive Copper cable, VPI, up to 100Gb/s, QSFP, LSZH, 1.5m           |
| EDR   | MCP1600-E02A    | Mellanox® Passive Copper cable, VPI, up to 100Gb/s, QSFP, LSZH, 2.5m           |
| EDR   | MCP1OPT-E002    | Mellanox® Passive Copper cable, VPI, up to 100Gb/s, QSFP, LSZH, 2m             |
| EDR   | MFA1A00-E005    | Mellanox Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 5m                      |
| EDR   | MFA1A00-E006    | Mellanox Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 6m                      |
| EDR   | MFA1A00-E010    | Mellanox Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 10m                     |
| EDR   | MFA1A00-E015    | Mellanox Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 15m                     |
| EDR   | MFA1A00-E020    | Mellanox Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 20m                     |



**Table 13 - Validated and Supported EDR / 100Gb/s Cables**

| Speed | Cable OPN #  | Description  |
|-------|--------------|--|
| EDR   | MFA1A00-E030 | Mellanox Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 30m                   |
| EDR   | MFA1A00-E050 | Mellanox Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 50m                   |
| EDR   | MFA1A00-E100 | Mellanox Active Fiber Cable, VPI, up to 100Gb/s, QSFP, 100m                  |
| EDR   | MMA1B00-E100 | Mellanox® Transceiver, IB EDR, up to 100Gb/s, QSFP28, MPO, 850nm, up to 100m |
| EDR   | MMA1L30-CM   | Mellanox® optical module, 100Gb/s, QSFP28, LC-LC, 1310nm, CWDM4, up to 2km   |

## 1.2.12 Validated and Supported HDR / 200Gb/s Cables

**Table 14 - Validated and Supported HDR Cables**

| Speed | OPN # / Name                 | Description  |
|-------|------------------------------|--|
| HDR   | MCP1650-H001E30              | Mellanox® Passive Copper cable, IB HDR, up to 200Gb/s, QSFP28, PVC, 1m, white pultab, 30AWG                        |
| HDR   | MCP1650-H002E26              | Mellanox Passive Copper Cable, IB HDR, up to 200GB/S, QSFP56, LSZH, 2M, black pultab, 26AWG                        |
| HDR   | MCP1650-H003E26 / P06149-B26 | Mellanox® Passive Copper cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, 3m, black pulltab, 26AWG                      |
| HDR   | MCP1650-H00AE30              | Mellanox Passive Copper Cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, 0.5M, black pultab, 30AWG                      |
| HDR   | MCP1650-H01AE30              | Mellanox® Passive Copper cable, IB HDR, up to 200Gb/s, QSFP28, PVC, 1.5m, white pultab, 30AWG                      |
| HDR   | MCP1650-H01AE30              | Mellanox Passive Copper Cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, 1.5M, black pultab, 30AWG                      |
| HDR   | MCP1650-H02AE26              | Mellanox® Passive Copper cable, IB HDR, up to 200Gb/s, QSFP28, PVC, 2.5m, white pultab, 26AWG                      |
| HDR   | MCP7H50-H003R26              | Mellanox® passive copper hybrid cable, IB HDR 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, colored, 3m, 26AWG   |
| HDR   | MCP7H50-H01AR30              | Mellanox® passive copper hybrid cable, IB HDR 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, colored, 1.5m, 30AWG |



HDR links raise with RS-FEC.

## 1.3 Tested Switches

### 1.3.1 Tested 10GbE Switches

**Table 15 - Tested 10GbE Switches**

| Speed    | Switch Silicon | OPN # / Name | Description                 | Vendor   |
|----------|----------------|--------------|-----------------------------|----------|
| 10/40GbE | N/A            | 3064         | 48-port 10Gb/40Gb Switch    | Cisco    |
| 10/40GbE | N/A            | 7050Q        | 16-port 40Gb Switch         | Arista   |
| 10/40GbE | N/A            | 7050S        | 48-port 10Gb/40Gb Switch    | Arista   |
| 10GbE    | N/A            | 5548         | Cisco 10GB ETH switch       | Cisco    |
| 10GbE    | N/A            | G8264        | BNT 10/40GB ETH switch      | BNT      |
| 10GbE    | N/A            | QFX3500      | Juniper 10/40GB ETH switch  | Juniper  |
| 10GbE    | N/A            | S4810P-AC    | 48-port 10Gb/40Gb Switch    | Force10  |
| 10GbE    | SwitchX®       | SX1016X-1BFR | 64-Port 10GbE Switch System | Mellanox |

### 1.3.2 Tested 40GbE Switches

**Table 16 - Tested 40GbE Switches**

| Speed    | Switch Silicon | OPN # / Name | Description                    | Vendor   |
|----------|----------------|--------------|--------------------------------|----------|
| 10/40GbE | N/A            | 3064         | 48-port 10Gb/40Gb Switch       | Cisco    |
| 10/40GbE | N/A            | 7050Q        | 16-port 40Gb Switch            | Arista   |
| 10/40GbE | N/A            | 7050S        | 48-port 10Gb/40Gb Switch       | Arista   |
| 40GbE    | N/A            | 3132Q        | Cisco 40GB ETH switch          | Cisco    |
| 40GbE    | N/A            | 7050QX       | 32-port 40Gb Switch            | Arista   |
| 40GbE    | N/A            | G8316        | BNT 40GB RackSwitch G8316      | BNT      |
| 40GbE    | N/A            | S6000        | 32-port 40Gb Switch            | Dell     |
| 40GbE    | SwitchX®       | SX1036B-1BFR | 36-Port 40/56GbE Switch System | Mellanox |

### 1.3.3 Tested 100GbE Switches

**Table 17 - Tested 100GbE Switches**

| Speed  | Switch Silicon | OPN # / Name | Description  | Vendor |
|--------|----------------|--------------|--|--------|
| 100GbE | N/A            | 7060CX       | 32-port 100Gb Switch   | Arista |
| 100GbE | N/A            | 93180YC-EX   | 48 x 10/25-Gbps fiber ports and 6 x 40/100-Gbps Quad Small Form-Factor Pluggable 28 (QSFP28) ports | Cisco  |
| 100GbE | N/A            | C3232C       | High-Density, 100 Gigabit Ethernet Switch  | Cisco  |
| 100GbE | N/A            | CE8860-4C-EI | 24x10GE (SFP+) or 25GE (SFP28) and 2x100GE switch  | Huawei |

**Table 17 - Tested 100GbE Switches**

| Speed  | Switch Silicon | OPN # / Name      | Description  | Vendor   |
|--------|----------------|-------------------|--|----------|
| 100GbE | Spectrum       | SN2410-CB2F       | 48-port 25GbE + 8-port 100GbE Open Ethernet ToR Switch System              | Mellanox |
| 100GbE | Spectrum       | SN2700-CS2R       | 32-port Non-blocking 100GbE Open Ethernet Spine Switch System              | Mellanox |
| 100GbE | Spectrum       | SN2740-CB2F1      | 32-port Non-blocking 100GbE Open Ethernet Spine Switch System              | Mellanox |
| 100GbE | N/A            | Wedge 100-32X R04 | 32-port 100GbE QSFP28 - Leaf/Spine Switch, power-to-port airflow, DC Power | Edgecore |
| 100GbE | N/A            | QFX5200-32C-32    | 32-port 100GbE Ethernet Switch System                                      | Juniper  |
| 100GbE | N/A            | S6820-56HF        | 48 SFP+ + 8 QSFP Ports 100GbE Switch Ethernet                              | H3C      |

### 1.3.4 Tested QDR Switches

**Table 18 - Tested QDR Switches**

| Speed | Switch Silicon  | OPN # / Name | Description   | Vendor   |
|-------|-----------------|--------------|---|----------|
| QDR   | N/A             | 12300        | 36-Port 40Gb QDR InfiniBand Switch, Management Module, Dual Power | QLogic   |
| QDR   | InfiniScale® IV | IS5025Q-1SFC | 36-port 40Gb/s InfiniBand Switch Systems                          | Mellanox |
| QDR   | InfiniScale® IV | Switch 4036  | Grid Director™ 4036E  | Mellanox |

### 1.3.5 Tested FDR Switches

**Table 19 - Tested FDR Switches**

| Speed | Switch Silicon | OPN # / Name | Description                                  | Vendor   |
|-------|----------------|--------------|--|----------|
| FDR   | SwitchX®       | SX6018F-1SFR | 18-port 56Gb/s InfiniBand/VPI Switch Systems | Mellanox |
| FDR   | SwitchX®       | SX6036F-1BFR | 36-port 56Gb/s InfiniBand/VPI Switch Systems | Mellanox |
| FDR   | SwitchX®       | SX6506       | 108-Port 56Gb/s InfiniBand Director Switch   | Mellanox |
| FDR   | SwitchX®-2     | SX6710-FB2F2 | 36-port 56Gb/s InfiniBand/VPI Switch Systems | Mellanox |

### 1.3.6 Tested EDR / 100Gb/s Switches

**Table 20 - Tested EDR Switches**

| Speed | Switch Silicon | OPN # / Name | Description  | Vendor   |
|-------|----------------|--------------|--|----------|
| EDR   | Switch-IB      | MSB7790-EB2F | 36-port EDR 100Gb/s InfiniBand Switch Systems                    | Mellanox |
| EDR   | Switch-IB 2    | MSB7800-ES2R | 36-port Non-blocking Managed EDR 100Gb/s InfiniBand Smart Switch | Mellanox |

## 1.4 Tools, Switch Firmware and Driver Software

The following are the drivers' software, tools, switch/HCA firmware versions tested that you can upgrade from or downgrade to when using firmware Rev 16.26.1040:

**Table 21 - Tools, Switch Firmware and Driver Software**

|                                | Supported Version   |
|--------------------------------|---|
| MLNX_OFED                      | 4.7-x.0.0.0 / 4.6-1.0.1.1   |
| MLNX_EN (MLNX_OFED based code) | 4.7-x.0.0.0 / 4.6-1.0.1.1   |
| WinOF-2                        | 2.30 / 2.20   |
| MFT                            | 4.13.0 / 4.12.0   |
| MLNX-OS                        | 3.8.2004  |
| Onyx                           | 3.8.2004  |
| ConnectX-5 Firmware            | 16.25.1020 / 16.24.1000   |
| SwitchX-IB™ Firmware           | 11.2000.2046  |
| SwitchX-IB 2 Firmware          | 15.2000.2046  |
| Linux Inbox Drivers            | <ul style="list-style-type: none"> <li>• RH7.6</li> <li>• Ubuntu 16.04.05</li> </ul>                                |
| Windows Inbox Drivers          | <ul style="list-style-type: none"> <li>• Windows 2012</li> <li>• Windows 2012 R2</li> <li>• Windows 2016</li> </ul> |

## 1.5 Supported FlexBoot, UEFI



Please be aware that not all firmware binaries contain FlexBoot or UEFI, support may vary between cards (see [Section 1.1, “Supported Devices”, on page 7.](#))

Firmware Rev 16.26.1040 supports the following FlexBoot:

**Table 22 - Supported FlexBoot, UEFI**

| Expansion ROM | Supported Version |
|---------------|-------------------|
| FlexBoot      | 3.5.803           |
| UEFI          | 14.19.14          |

## 1.6 Revision Compatibility

Firmware Rev 16.26.1040 complies with the following programmer’s reference manual:

- *Mellanox Adapters Programmer’s Reference Manual (PRM), Rev 0.47 or later*, which has Command Interface Revision 0x5. The command interface revision can be retrieved by means of the QUERY\_FW command and is indicated by the field `cmd_interface_rev`.

## 2 Changes and New Features in Rev 16.26.1040

**Table 23 - Changes and New Features in Rev 16.26.1040**

| Feature/Change                              | Description   |
|---|---|
| <b>Rev. 16.26.1040</b>                      |   |
| <b>Address Translation Service (ATS)</b>    | Added Address Translation Service (ATS) support for MKEY and UMEM.  |
| <b>VPD</b>                                  | Added support for exposing the VPD on the VF.   |
| <b>ICMD and Diagnostic Counters</b>         | Enabled the firmware by using the ICMD commands to deal with diagnostic counters similar to cmdif. They can be called via the vsec space. The counters' values are returned only via the tracer. The ICMD Query Caps indicate support and expose the list of the supported counters.  |
| <b>Hairpin Drop Counter</b>                 | Added support for Hairpin Drop Counter.   |
| <b>User Context Object (DEVX)</b>           | This is a containerized sandbox per user, to access PRM command securely by using General Object commands, UMEM and UCTX contexts. The allowed functionalities of this capability depend on the user permissions.<br>The following functionalities are still managed by the Kernel: <ul style="list-style-type: none"> <li>• Resource cleaning</li> <li>• UCTX stamping</li> <li>• Blocking the physical address and IRQ from these UCTX</li> </ul> |
| <b>DEVX Support for Asynchronous Events</b> | Added support for reporting the supported affiliated and unaffiliated asynchronous events to DEVX users through the command interface.  |
| <b>Hairpin and TM RNDV QPs in DEVX</b>      | Added support for Hairpin and TM RNDV QPs to work with DevX.  |
| <b>Software Managed Steering Tables</b>     | Added support for creating software managed steering tables in eSwitch/FDB.   |
| <b>Zero-Touch-RoCE Counters</b>             | Zero-Touch-RoCE counters are now available to the user for debuggability purposes when using the Zero-Touch-RoCE feature.   |
| <b>Security Hardening Enhancements</b>      | This release contains important reliability improvements and security hardening enhancements.<br>Mellanox recommends upgrading your device firmware to this release to improve the device firmware security and reliability.  |
| <b>Bug Fixes</b>                            | See <a href="#">Section 4, “Bug Fixes History”, on page 30</a>  |

### 3 Known Issues

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

For a list of old firmware Know Issues, please see ConnectX-5 Firmware Archived Known Issues file ([http://www.mellanox.com/pdf/firmware/ConnectX5-Firmware\\_Archived\\_Known\\_Issues.pdf](http://www.mellanox.com/pdf/firmware/ConnectX5-Firmware_Archived_Known_Issues.pdf))

**Table 24 - Ethernet Rate Limit per VF in RoCE Mode Limitations**

| Adapter Card | Dual Port Device       |          |                         |          | Single Port Device |          |
|--------------|------------------------|----------|-------------------------|----------|--------------------|----------|
|              | w/o LAG (TOTAL_VFS>32) |          | With LAG (TOTAL_VFS<32) |          | w/o LAG            |          |
|              | w/o QoS                | Full QoS | w/o QoS                 | Full QoS | w/o QoS            | Full QoS |
| ConnectX-5   | 127                    | 127      | 64                      | 64       | 127                | 127      |

**Table 25 - Ethernet Rate Limit per VF in InfiniBand Mode Limitations**

| Adapter Card | Dual Port Device |          | Single Port Device |          |
|--------------|------------------|----------|--------------------|----------|
|              | w/o LAG          |          | w/o LAG            |          |
|              | w/o QoS          | Full QoS | w/o QoS            | Full QoS |
| ConnectX-5   | 127              | 127      | 127                | 127      |

**Table 26 - Known Issues (Sheet 1 of 7)**

| Internal Ref. | Issue  |
|---------------|--|
| 1888306       | <b>Description:</b> Occasionally Bluescreen might occur when using mlxfwreset for Socket Direct devices on Windows.  |
|               | <b>Workaround:</b> N/A   |
|               | <b>Keywords:</b> mlxfwreset, Socket Direct,  |
|               | <b>Discovered in Version:</b> 16.26.1040   |
| 1895191       | <b>Description:</b> SX_RDMA is not supported when Dual Port RoCE feature is enabled. Due to this behavior, packets sent on other port are be processed by the SX_RDMA table. |
|               | <b>Workaround:</b> Use SX_RDMA with Dual Port GVMI instead.  |
|               | <b>Keywords:</b> SX_RDMA, Dual Port RoCE, Dual Port GVMI   |
|               | <b>Discovered in Version:</b> 16.26.1040   |
| 1919403       | <b>Description:</b> Hardware arbitration is currently disabled in OCP3.0 cards. It will be supported on future releases for the same hardware.                               |
|               | <b>Workaround:</b> N/A   |
|               | <b>Keywords:</b> Hardware arbitration, OCP3.0  |
|               | <b>Discovered in Version:</b> 16.26.1040   |
| 1912117       | <b>Description:</b> The sw_reset option is not supported when ATS is enabled.  |
|               | <b>Workaround:</b> N/A   |
|               | <b>Keywords:</b> ATS, sw_reset   |
|               | <b>Discovered in Version:</b> 16.26.1040   |

**Table 26 - Known Issues (Sheet 2 of 7)**

| Internal Ref. | Issue  |
|---------------|--|
| 1836465       | <p><b>Description:</b> When using the hairpin feature, and using VLAN strip or using the “modify esw vport context” command, the packets can have an incorrect VLAN header. Meaning, using VLAN push/pop may not work properly when using vport context VLAN.</p> <p>The features that may be affected by this and not work properly are:</p> <ul style="list-style-type: none"> <li>• Host chaining</li> <li>• Mirroring in FDB</li> <li>• TTL modify in FDB</li> <li>• VGT+</li> </ul> |
|               | <b>Workaround:</b> N/A   |
|               | <b>Keywords:</b> E-switch vport context, VLAN  |
|               | <b>Discovered in Version:</b> 16.26.1040   |
| 1842278       | <p><b>Description:</b> DC LAG can function only in case there is a single PF per port without any active VFs.</p>  |
|               | <b>Workaround:</b> N/A   |
|               | <b>Keywords:</b> DC LAG  |
| 1796628       | <p><b>Description:</b> Due to performance considerations, unicast loopback traffic will go through the NIC SX tables, and multicast loopback traffic will skip the NIC SX tables.</p>  |
|               | <b>Workaround:</b> N/A   |
|               | <b>Keywords:</b> Performance, unicast loopback traffic, multicast loopback traffic   |
| 1797493       | <p><b>Description:</b> Firmware asserts may occur when setting the PF_BAR2_SIZE value higher than the maximum supported size (maximum PF_BAR2_SIZE is 4 for .</p>  |
|               | <b>Workaround:</b> Configure within limits (NIC PF_BAR_SIZE <= 4).   |
|               | <b>Keywords:</b> Multi-GVMI, Sub-Function, SFs, BAR2   |
| 1761271       | <p><b>Description:</b> CWD4 AOM cable is currently not supported.</p>  |
|               | <b>Workaround:</b> N/A   |
|               | <b>Keywords:</b> Modules/Cables  |
| 1752009       | <p><b>Description:</b> When working with Multi-GVMI and SR-IOV, and with a high number of Virtual Functions and sub-functions, the driver start may fail for the VFs/sub-functions.</p>  |
|               | <b>Workaround:</b> Decrease the number of configured VFs or sub-functions using mlxconfig PF_BAR2_SIZE or NUM_OF_VFS.  |
|               | <b>Keywords:</b> Multi-GVMI, SR-IOV  |
| 1762142       | <p><b>Description:</b> PF / ECPF FLR does not clear all its dependent sub-functions. QUERY_ESW_FUNCTIONS and ALLOC/DEALLOC_SF commands might fail / show allocated SFs after PF FLR.</p>   |
|               | <b>Workaround:</b> Perform a graceful shutdown, and not an FLR.  |
|               | <b>Keywords:</b> Multi-GVMI, SF, Sub-Functions, FLR  |
|               | <b>Discovered in Version:</b> 16.25.1020   |



**Table 26 - Known Issues (Sheet 3 of 7)**

| Internal Ref.   | Issue  |
|-----------------|--|
| 1768814/1772474 | <b>Description:</b> Due to hardware limitation, REG_C cannot be passed over loopback when the FDB action is forwarded to multiple destinations.  |
|                 | <b>Workaround:</b> N/A   |
|                 | <b>Keywords:</b> Connection-Tracking   |
|                 | <b>Discovered in Version:</b> 16.25.1020   |
| 1770736         | <b>Description:</b> When a PF or ECPF with many VFs (SR-IOV), and/or SFs (Multi-GVMI) triggers an FLR, PCIe completion timeout might occur.  |
|                 | <b>Workaround:</b> Increase the PCIe completion timeout.   |
|                 | <b>Keywords:</b> Multi-GVMI, SR-IOV, Sub-Function, Virtual Function, PF FLR  |
|                 | <b>Discovered in Version:</b> 16.25.1020   |
| 1716334         | <b>Description:</b> When mlxconfig.PF_BAR2_EN is enabled, configuring more than 255 PCI functions will raise an assert.  |
|                 | <b>Workaround:</b> When working with BAR2, configure SR-IOV to align to the 255 PCI functions limitation.<br>mlxconfig.NUM_OF_VFS controls the number of configured SR-IOV VFs. e.g.: <ul style="list-style-type: none"> <li>Smart NICs: 2 External Host PFs, 2 ARM ECPFs, 125 VFs per PF.</li> <li>Non-smart NICs: 2 External Host PFs, 126 VFs per PF</li> </ul> |
|                 | <b>Keywords:</b> Multi-GVMI, PF_BAR2_EN, Sub-Functions, SR-IOV, VFs  |
|                 | <b>Discovered in Version:</b> 16.25.1020   |
| 1699214         | <b>Description:</b> NODNIC VF is partially tested. It is fully tested only in ConnectX-5 adapter cards.  |
|                 | <b>Workaround:</b> N/A   |
|                 | <b>Keywords:</b> NODNIC VF   |
|                 | <b>Discovered in Version:</b> 16.25.1020   |
| 1699214         | <b>Description:</b> NODNIC VF is partially tested. It is fully tested only in ConnectX-5 adapter cards.  |
|                 | <b>Workaround:</b> N/A   |
|                 | <b>Keywords:</b> NODNIC VF   |
|                 | <b>Discovered in Version:</b> 16.25.1020   |
| 1749691         | <b>Description:</b> On rare occasions, when using Socket-Direct devices, inband burning through the external port might fail.  |
|                 | <b>Workaround:</b> N/A   |
|                 | <b>Keywords:</b> Socket-Direct, inband burning   |
|                 | <b>Discovered in Version:</b> 16.25.1020   |
| 1689186         | <b>Description:</b> Changing priority to TC map during traffic might cause packet drops.   |
|                 | <b>Workaround:</b> N/A   |
|                 | <b>Keywords:</b> QoS   |
|                 | <b>Discovered in Version:</b> 16.25.1020   |
| 1604699         | <b>Description:</b> Ethernet RFC 2819 counter ether_stats_oversize_pkts and Ethernet IEEE 802.3 counter a_frame_too_long_errors share the same resource. Clearing each of them will affect the other.  |
|                 | <b>Workaround:</b> N/A   |
|                 | <b>Keywords:</b> Counters  |
|                 | <b>Discovered in Version:</b> 16.25.1020   |

**Table 26 - Known Issues (Sheet 4 of 7)**

| Internal Ref. | Issue   |
|---------------|---|
| 1558250       | <b>Description:</b> eSwitch owner may receive NIC_VPORT_CONTEXT events from vPorts that are not necessarily armed using the nic_vport_context_arm_change_even_tbit.   |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> Port event, NODNIC   |
|               | <b>Discovered in Version:</b> 16.25.1020  |
| -             | <b>Description:</b> In Ethernet mode, at 10/40GbE speeds, only NO-FEC in Force mode is supported. Other user configurations are overridden.   |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> Ethernet, 10GbE, 40GbE, RS-FEC   |
|               | <b>Discovered in Version:</b> 16.25.1020  |
| 1574876       | <b>Description:</b> DC RoCE LAG is functional only if the router posts VRRP address as the source MAC.  |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> DC RoCE LAG  |
|               | <b>Discovered in Version:</b> 16.25.1020  |
| 1498399       | <b>Description:</b> If the XRC switches between SRQ/RMPs while there is an outstanding ODP on the responder XRC QP, a CQE with an error might be generated (that is not a PFAULT abort).  |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> XRC SRQ/RMP ODP  |
|               | <b>Discovered in Version:</b> 16.25.1020  |
| 1546401       | <b>Description:</b> vport_tc and para_vport_tc are not supported in this version.   |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> SR-IOV vport_tc and para_vport_tc  |
|               | <b>Discovered in Version:</b> 16.24.1000  |
| 1546492       | <b>Description:</b> Executing the update_lid command while the IB port sniffer utility is active can stop the utility.  |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> IB Sniffer   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
| 1537898       | <b>Description:</b> Initializing a function while the IB port sniffer utility is active can stop the utility.   |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> IB Sniffer   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
| 1523577       | <b>Description:</b> When modifying the TTL in the NIC RX, the CQE checksum is not recalculated automatically. The limitation is indicated by the ttl_checksum_correction bit. If the ttl_checksum_correction=0, the capability is not functioning properly. |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> multi_prio_sq, VF  |
|               | <b>Discovered in Version:</b> 16.24.1000  |

**Table 26 - Known Issues (Sheet 5 of 7)**

| Internal Ref.   | Issue  |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|---|--|-------|-----|--------|--------|--|------------|-------|----|-----|--------|-------|----|----|--|---------------|-------|----|-----|--------|-------|----|-----|--|------------|-------|-----|-----|--|-------|-----|-----|--|
| 1414290   | <b>Description:</b> When getting an inline scatter CQE on IB striding RQ, the stride index in the CQE will be zero.  |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Workaround:</b> N/A   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Keywords:</b> Scatter CQE   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Discovered in Version:</b> 16.24.1000   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
| 1475490   | <b>Description:</b> Reboot is not supported on any host during the PLDM firmware burning process.  |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Workaround:</b> N/A   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Keywords:</b> PLDM  |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Discovered in Version:</b> 16.23.1020   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
| 1332714   | <b>Description:</b> The maximum “read” size of MTRC_STDB is limited to 272 Bytes.  |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Workaround:</b> Set the MTRC_STDB.read_size to the maximum value of 0x110=272 Bytes   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Keywords:</b> Access register, MTRC_STDB, tracer to dmesg, fwtrace to dmesg   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Discovered in Version:</b> 16.23.1020   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
| 1408994   | <b>Description:</b> FTE with both forward (FWD) and encapsulation (ENCAP) actions is not supported in the SX NIC Flow Table.   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Workaround:</b> N/A   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Keywords:</b> SX NIC Flow Table   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Discovered in Version:</b> 16.23.1020   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
| 1350794   | <b>Description:</b> Encapsulation / Decapsulation support in steering has the following limitations: <ul style="list-style-type: none"> <li>• Encapsulation / Decapsulation can be open on the FDB only if all VFs are non active.</li> <li>• Encapsulation / Decapsulation supports single mode only: FDB / NIC. Opening tables of both types is not supported.</li> <li>• Encapsulation / Decapsulation per device support: <table border="1" data-bbox="510 1134 981 1344"> <thead> <tr> <th></th> <th></th> <th>NIC</th> <th>FDB</th> <th></th> </tr> </thead> <tbody> <tr> <td rowspan="2">ConnectX-4</td> <td>encap</td> <td>NO</td> <td>YES</td> <td>non MH</td> </tr> <tr> <td>decap</td> <td>NO</td> <td>NO</td> <td></td> </tr> <tr> <td rowspan="2">ConnectX-4 Lx</td> <td>encap</td> <td>NO</td> <td>YES</td> <td>non MH</td> </tr> <tr> <td>decap</td> <td>NO</td> <td>YES</td> <td></td> </tr> <tr> <td rowspan="2">ConnectX-5</td> <td>encap</td> <td>YES</td> <td>YES</td> <td></td> </tr> <tr> <td>decap</td> <td>YES</td> <td>YES</td> <td></td> </tr> </tbody> </table> </li> </ul> |       |     | NIC    | FDB    |  | ConnectX-4 | encap | NO | YES | non MH | decap | NO | NO |  | ConnectX-4 Lx | encap | NO | YES | non MH | decap | NO | YES |  | ConnectX-5 | encap | YES | YES |  | decap | YES | YES |  |
|   |  |       | NIC | FDB    |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | ConnectX-4   | encap | NO  | YES    | non MH |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   |  | decap | NO  | NO     |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
| ConnectX-4 Lx   | encap  | NO    | YES | non MH |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | decap  | NO    | YES |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
| ConnectX-5  | encap  | YES   | YES |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | decap  | YES   | YES |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
| <b>Workaround:</b> N/A                                  |  |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
| <b>Keywords:</b> Steering Encapsulation / Decapsulation |  |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
| <b>Discovered in Version:</b> 16.23.1020                |  |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
| 1027553   | <b>Description:</b> While using e-switch vport sVLAN stripping, the RX steering values on the sVLAN might not be accurate.   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Workaround:</b> N/A   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Keywords:</b> e-sw vport sVLAN stripping, RX steering   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Discovered in Version:</b> 16.24.1000   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
| 1799917   | <b>Description:</b> Untagged CVLAN packets in the Steering Flow Tables do not match the SVLAN tagged packets.  |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Workaround:</b> N/A   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Keywords:</b> Steering Flow Tables, CVLAN/SVLAN packets   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |
|   | <b>Discovered in Version:</b> 16.23.1020   |       |     |        |        |  |            |       |    |     |        |       |    |    |  |               |       |    |     |        |       |    |     |  |            |       |     |     |  |       |     |     |  |

**Table 26 - Known Issues (Sheet 6 of 7)**

| Internal Ref. | Issue   |
|---------------|---|
| 1504073       | <b>Description:</b> When using ConnectX-5 with LRO over PPC systems there might be backpressure to the NIC due to delayed PCI writes operations. In this case bandwidth might drop from line-rate to ~35Gb/s. Packet loss or pause frames might also be observed. |
|               | <b>Workaround:</b> Look for an indication of PCI back pressure (“outbound_pci_stalled_wr” counter in ethtools advancing). Disabling LRO helps reduce the back pressure and its effects.   |
|               | <b>Keywords:</b> Flow Control, LRO  |
|               | <b>Discovered in Version:</b> 16.23.1020  |
| 1277762       | <b>Description:</b> An Ethernet multicast loopback packet is not counted (even if it is not a local loopback packet) when running the <code>nic_receive_steering_discard</code> command.  |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> Ethernet multicast loopback packet   |
|               | <b>Discovered in Version:</b> 16.22.1002  |
| 1190753       | <b>Description:</b> When a dual-port VHCA sends a RoCE packet on its non-native port, and the packet arrives to its affiliated vport FDB, a mismatch might happen on the rules that match the packet source vport.  |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> RoCE, vport FDB  |
|               | <b>Discovered in Version:</b> 16.22.1002  |
| 1306342       | <b>Description:</b> Signature-accessing WQEs sent locally to the NVMeF target QPs that encounter signature errors, will not send a SIGERR CQE.  |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> Signature-accessing WQEs, NVMeF target   |
|               | <b>Discovered in Version:</b> 16.22.1002  |
| 1168594       | <b>Description:</b> RoCE Dual Port Mode (a.k.a Multi-Port vHCA: MPV) is not supported in Multi-Host setups.   |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> Multi-Port vHCA, Multi-Host  |
|               | <b>Discovered in Version:</b> 16.21.1000  |
| 1072337       | <b>Description:</b> If a packet is modified in e-sw flow steering, the SX sniffer Flow Table (of the VF) will see the sniffed packet after the modification.  |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> SX sniffer Flow Table  |
|               | <b>Discovered in Version:</b> 16.21.1000  |
| 1171013       | <b>Description:</b> Signature Handover Operations is not supported when FPP (Function-Per-Port) mode is disabled.   |
|               | <b>Workaround:</b> N/A  |
|               | <b>Keywords:</b> Signature Handover Operations, FPP   |
|               | <b>Discovered in Version:</b> 16.21.1000  |

**Table 26 - Known Issues (Sheet 7 of 7)**

| Internal Ref. | Issue  |
|---------------|--|
| 1059975       | <p><b>Description:</b> NVMeF limitation:</p> <ul style="list-style-type: none"> <li>• Transaction size - up to 128KB per IO (non-inline)</li> <li>• Support up to 16K connections</li> <li>• Support single namespace per drive</li> <li>• Staging buffer size must be at least 16MB in order to allow SRQ size of 64 entries</li> </ul> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> NVMeF</p> <p><b>Discovered in Version:</b> 16.22.1010</p> |

## 4 Bug Fixes History

Table 27 lists the bugs fixed in this release. For a list of old firmware Bug Fixes, please see ConnectX-5 Firmware Archived Bug Fixes file

[http://www.mellanox.com/pdf/firmware/ConnectX5-Firmware\\_Archived\\_Bug\\_Fixes.pdf](http://www.mellanox.com/pdf/firmware/ConnectX5-Firmware_Archived_Bug_Fixes.pdf)

**Table 27 - Bug Fixes History (Sheet 1 of 5)**

| Internal Ref. | Issue  |
|---------------|--|
| 1803791       | <b>Description:</b> On rare occasions, when firmware coalesce Host stuck events occur, a async event might be delayed to be reported, and not be triggered until the next time the PCIe hangs on one of the hosts. |
|               | <b>Keywords:</b> PCIe Error Notification   |
|               | <b>Discovered in Version:</b> 16.25.1020   |
|               | <b>Fixed in Release:</b> 16.26.1040  |
| 1870690       | <b>Description:</b> Fixed an issue that resulted in unexpected queue pairs transitioned to error in lossy tests.   |
|               | <b>Keywords:</b> RoCE Lossy  |
|               | <b>Discovered in Version:</b> 16.25.1020   |
|               | <b>Fixed in Release:</b> 16.26.1040  |
| 1723460       | <b>Description:</b> Limited the maximum amount of dumps created on a PF.   |
|               | <b>Keywords:</b> Dump files  |
|               | <b>Discovered in Version:</b> 16.25.1020   |
|               | <b>Fixed in Release:</b> 16.26.1040  |
| 1824111       | <b>Description:</b> Renamed the GMP Mellanox Vendor Specific External Capability mask enum from <code>IsDiagnosticCountersSupported</code> to <code>IsDiagnosticDataSupported</code> .                             |
|               | <b>Keywords:</b> GMP Mellanox Vendor Specific External Capability mask DiagnosticData  |
|               | <b>Discovered in Version:</b> 16.25.1020   |
|               | <b>Fixed in Release:</b> 16.26.1040  |
| 1784290       | <b>Description:</b> Fixed a stability issue in RoCE retransmissions under stress affecting Zero-Touch-RoCE.  |
|               | <b>Keywords:</b> Zero-Touch-RoCE   |
|               | <b>Discovered in Version:</b> 16.25.1020   |
|               | <b>Fixed in Release:</b> 16.26.1040  |
| 1822787       | <b>Description:</b> Fixed an issue that caused a function to misbehave when a PCIe TLP was set with a poisoned indication.   |
|               | <b>Keywords:</b> PCIe TLP  |
|               | <b>Discovered in Version:</b> 16.25.1020   |
|               | <b>Fixed in Release:</b> 16.26.1040  |
| 1788895       | <b>Description:</b> Fixed an issue that caused large number of packet to drop when running Jumbo frames with TTL rewrite.  |
|               | <b>Keywords:</b> Jumbo frames, TTL   |
|               | <b>Discovered in Version:</b> 16.25.1020   |
|               | <b>Fixed in Release:</b> 16.26.1040  |
| 1806874       | <b>Description:</b> Limited the number of the elements in the QoS tree 2K.<br><b>Note:</b> Creating more than 250 Vport_tc for every TC is not allowed.  |
|               | <b>Keywords:</b> VQoS tree , 255Vfs  |
|               | <b>Discovered in Version:</b> 16.25.1020   |
|               | <b>Fixed in Release:</b> 16.26.1040  |

**Table 27 - Bug Fixes History (Sheet 2 of 5)**

| Internal Ref. | Issue   |
|---------------|---|
| 1769862       | <b>Description:</b> Fixed the query QP flow. Instead of always taking port 1, FW will now reply the proper port, 1 or 2, for the dual port RoCE net device.   |
|               | <b>Keywords:</b> Counters   |
|               | <b>Discovered in Version:</b> 16.24.1020  |
|               | <b>Fixed in Release:</b> 16.26.1040   |
| 1771326       | <b>Description:</b> Fixed an issue that caused the nack counters to constantly be reported as "0".  |
|               | <b>Keywords:</b> Counters   |
|               | <b>Discovered in Version:</b> 16.24.1020  |
|               | <b>Fixed in Release:</b> 16.26.1040   |
| 1775228       | <b>Description:</b> In a rare scenario when the driver is executing the "2err" command and the QP is in SQ drain state, the firmware might post event of broken WQ instead of sending error CQEs on all the WQEs. |
|               | <b>Keywords:</b> QP, WQE, CQE   |
|               | <b>Discovered in Version:</b> 16.25.1020  |
|               | <b>Fixed in Release:</b> 16.26.1040   |
| 1750224       | <b>Description:</b> Fixed an issue that resulted in packets sometimes being dropped while setting the software owner Flow Table as root.  |
|               | <b>Keywords:</b> Flow Table   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1771921       | <b>Description:</b> Fixed an issue that prevented users with non-port owner privilege from using the "read DCBX access registry key" REGID_DCBX_APP/REGID_DCBX_PARAM.   |
|               | <b>Keywords:</b> DCBX   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16..25.1020  |
| 1726413       | <b>Description:</b> Fixed an issue that resulted in inaccurate counters' value (negative value) when querying the vPort counters during traffic.  |
|               | <b>Keywords:</b> Counters   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1747677       | <b>Description:</b> Fixed an issue that prevented the firmware from working with multiple priorities when in Host Chaining mode.  |
|               | <b>Keywords:</b> Host Chaining  |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1682830       | <b>Description:</b> Fixed a rare issue that caused the RX to hang when the server went into the Standby mode (aux power).   |
|               | <b>Keywords:</b> RX   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |

**Table 27 - Bug Fixes History (Sheet 3 of 5)**

| Internal Ref. | Issue   |
|---------------|---|
| -             | <b>Description:</b> Fixed a wrong meter rate for small packets by excluding packet CRC size.  |
|               | <b>Keywords:</b> Meter rate   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1682727       | <b>Description:</b> Fixed an issue that slowed the driver's unload process.<br>The dmesg showed the following message: <i>teardown_hca, resource is still in use, gvmi=XX, type=0x2b</i>  |
|               | <b>Keywords:</b> Driver unload  |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1717993       | <b>Description:</b> Fixed an issue that required pages' calculations on SR-IOV, thus causing the adapter to return 'not enough resources' to the driver's VF.   |
|               | <b>Keywords:</b> SR-IOV   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1657685       | <b>Description:</b> Fixed an issue that cause an unexpected pause counter to be transmitted when the HCA was in tear down mode.   |
|               | <b>Keywords:</b> Pause counter, Ethernet  |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1679530       | <b>Description:</b> Fixed an issue that cause performance degradation in the system when configured priority to TC mapping where some priorities were mapped to TC7. This issue occurred in scenarios such as hairpin QP in system, e.g. host chaining / ttl WA |
|               | <b>Keywords:</b> Performance  |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1615586       | <b>Description:</b> Fixed a rare issue that caused the QP to falsely transition into the error state as a result of handling duplicate read/atomic request followed by memory key invalidation.   |
|               | <b>Keywords:</b> CQE  |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1689159       | <b>Description:</b> Fixed the gateway configurations when sending a MAD IB packet after a ETH packet, to prevent MAD transactions timeout when one port was set as IB, and one port as ETH.   |
|               | <b>Keywords:</b> Dual port with IB and ETH, MAD Timeout   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1697465       | <b>Description:</b> Fixed an issue related to RoCE Lossy Access Register that caused the logic of the condition not to be returned according to the systems' requirements.  |
|               | <b>Keywords:</b> RoCE Lossy Access Register   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |



**Table 27 - Bug Fixes History (Sheet 4 of 5)**

| Internal Ref. | Issue   |
|---------------|---|
| 1706349       | <b>Description:</b> Fixed an ECN compatibility issue with Broadcom RNICs.   |
|               | <b>Keywords:</b> ECN  |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1582423       | <b>Description:</b> In the case of multi-destinations transmission where the last destination is set to encapsulation & wire, the packet for the last destination will not be send. |
|               | <b>Keywords:</b> Remote mirroring   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1678824       | <b>Description:</b> Fixed an issue that prevented the user to enable the port after disabling it in the VF NODNIC.  |
|               | <b>Keywords:</b> VF NODNIC  |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1606289       | <b>Description:</b> Enlarged the number of modify fields to 16 to avoid IPv6 header rewrite failure.  |
|               | <b>Keywords:</b> IPv6 header rewrite  |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1627973       | <b>Description:</b> Fixed an issue that prevented IB QP counters for Acks/Responses from working as a results the NACK/OOS counters showed as zero.                                 |
|               | <b>Keywords:</b> IB QP counters for Acks/Responses  |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1554104       | <b>Description:</b> Set the stateless offloads cap to be permanently '1'.   |
|               | <b>Keywords:</b> Stateless offloads cap   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1541193       | <b>Description:</b> Fixed an issue that cause high connection setup latency on guest VMs.   |
|               | <b>Keywords:</b> Connection latency   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.25.1020   |
| 1554104       | <b>Description:</b> Fixed a rare case that cause the transmission side to hang and an assertion in the RDMA ODP.  |
|               | <b>Keywords:</b> RDMA ODP   |
|               | <b>Discovered in Version:</b> 16.24.1000  |
|               | <b>Fixed in Release:</b> 16.24.1000   |
| 1434510       | <b>Description:</b> Fixed a PCIe bifurcation issue.   |
|               | <b>Keywords:</b> PCIe   |
|               | <b>Discovered in Version:</b> 16.22.1002  |
|               | <b>Fixed in Release:</b> 16.24.1000   |

**Table 27 - Bug Fixes History (Sheet 5 of 5)**

| Internal Ref.       | Issue   |
|---------------------|---|
| 1424906             | <b>Description:</b> Fixed a deadlock in RX related to the “send-invalidate” flow, that resulted in RX getting stuck.  |
|                     | <b>Keywords:</b> RX, deadlock   |
|                     | <b>Discovered in Version:</b> 16.23.1020  |
|                     | <b>Fixed in Release:</b> 16.24.1000   |
| 929504              | <b>Description:</b> Fixed rare errors in RX that resulted in double completion.   |
|                     | <b>Keywords:</b> RX   |
|                     | <b>Discovered in Version:</b> 16.23.1020  |
|                     | <b>Fixed in Release:</b> 16.24.1000   |
| 1284452/<br>1282926 | <b>Description:</b> Fixed an issue that caused the mlxconfig tool to present all possible expansion ROM images, instead of presenting only the existing images.   |
|                     | <b>Keywords:</b> mlxconfig  |
|                     | <b>Discovered in Version:</b> 16.22.1002  |
|                     | <b>Fixed in Release:</b> 16.24.1000   |
| 1475993             | <b>Description:</b> Aligned the default tuning type in PHY TEST MODE to the device protocol.  |
|                     | <b>Keywords:</b> PHY  |
|                     | <b>Discovered in Version:</b> 16.23.1020  |
|                     | <b>Fixed in Release:</b> 16.24.1000   |
| 1403211             | <b>Description:</b> When a device is operating in Safe Mode state, and the user issues the mlxfwreset command, the device might fail to come-up correctly after the reset.<br><b>Note:</b> Do not run mlxfwreset when operating in a Safe Mode state. |
|                     | <b>Keywords:</b> mlxfwreset   |
|                     | <b>Discovered in Version:</b> 16.23.1020  |
|                     | <b>Fixed in Release:</b> 16.24.1000   |
| 1295606             | <b>Description:</b> Fixed an issue related to PCIe "Surprise link down" event reporting capability.   |
|                     | <b>Keywords:</b> PCIe   |
|                     | <b>Discovered in Version:</b> 16.22.1002  |
|                     | <b>Fixed in Release:</b> 16.24.1000   |
| 1434863             | <b>Description:</b> Fixed an issue that resulted in the link partner experiencing false active linkup when plugging in a base-T cable to a closed port.   |
|                     | <b>Keywords:</b> Interfaces   |
|                     | <b>Discovered in Version:</b> 16.22.1002  |
|                     | <b>Fixed in Release:</b> 16.24.1000   |
| 1424873             | <b>Description:</b> Modifying VMQoS rate limiter parameters during traffic might cause transmission failure.  |
|                     | <b>Keywords:</b> VMQoS, rate limiter  |
|                     | <b>Discovered in Version:</b> 16.22.1002  |
|                     | <b>Fixed in Release:</b> 16.24.1000   |

## 5 Firmware Changes and New Feature History

**Table 28 - Firmware Changes and New Feature History (Sheet 1 of 14)**

| Feature/Change   | Description   |
|--|---|
| <b>Rev. 16.25.1020</b>   |   |
| <b>OOB Support AR in IB Networks</b>                               | AddedSL based support for enabling adaptive routing only for traffic running on specific set of SLs. This capability is applicable to application and is enforced by the OpenSM (MADs).   |
| <b>ConnectX-5 OCP 3.0</b>  | Added the ability to read from the server whether or not the Socket-Direct is the requested mode on ConnectX-5 OCP 3.0 adapter cards.   |
| <b>ASAP<sup>2</sup> Offloading VXLAN Decapsulation with HW LRO</b> | Added support for performing hardware Large Receive Offload (HW LRO) on VFs with HW-decapsulated VXLAN.<br>For further information on the VXLAN decapsulation feature, please refer to ASAP <sup>2</sup> User Manual under <a href="http://www.mellanox.com">www.mellanox.com</a> -> Products -> Software -> ASAP2.   |
| <b>Multiple Output Ports - Enhance ASAP<sup>2</sup></b>            | Added the ability to send to multiple destinations encapsulated data and each destination can have its own encapsulation data.  |
| <b>VSC Security</b>  | VSC security includes the mechanisms which will prevent a reasonable host from affecting other hosts from using VSC.  |
| <b>Extended Responder Not Ready (RNR)</b>                          | Enabled Responder Not Ready (RNR) configuration. Now the number of RNRs can be set higher than 6 and lower than indefinite.   |
| <b>Zero Touch RoCE</b>   | Zero touch RoCE enables RoCE to operate on fabrics where no PFC nor ECN are configured. This makes RoCE configuration a breeze while still maintaining its superior high performance.   |
| <b>ODP support for SRQ &amp; XRC</b>                               | Added support for send opcode operations targeting a SRQ/RMP with the receive WQEs using ODP memory. In case the receive WQE receives an ODP, the device will generate ODP notifications (EQE) and PFAULT will abort CQEs.<br><b>Note:</b> It is recommended to prefetch the memory used by the receive WQEs to reduce ODP occurrence as these have significant latencies and will cause a performance degradation. |
| <b>TTL RX</b>  | Enables the device to modify incoming packets' TTL from the uplink to a vport using eswitch rules when <code>sw_fdb_ipv4_ttl_modify_enable</code> is set to true.   |
| <b>QP Counters and Firmware Errors per PID</b>                     | QP counters (RDMA errors) and flow counters (traffic) are now set per PID or UID and available through the driver and the user space tool.  |
| <b>Firmware Burning using DMA Pages</b>                            | This new capability accelerates the firmware burning process by using Direct Memory Access (DMA) pages.   |
| <b>Auto-Sensing when using 25/10GbE Optical Modules</b>            | This new capability accelerates the network to auto-sense the port speed and use it when using a 25/10GbE optical module. Meaning, if the used module is 25GbE but the port is a 10GbE port, the speed used for that network will be 10GbE.   |
| <b>Package ID</b>  | Enabled Package ID configuration using server strap according OCP 3.0.  |
| <b>DPDK UIO</b>  | This capability provides a solution for improving user space drivers development, generic user space IO device services.  |
| <b>Management Query Information Strings (MQIS)</b>                 | Added MQIS support, "Part Number" and "Description" information for secured adapter cards.  |
| <b>SR-IOV Offloading</b>   | Added the ability to probe one representor each time for a specific representor ID when bonding is enabled.   |

**Table 28 - Firmware Changes and New Feature History (Sheet 2 of 14)**

| Feature/Change  | Description  |
|---|--|
| <b>Virtio Supported Hardware Offloads (DPDK)</b>                            | <p>Added the following capabilities related to Virtio Supported Hardware Offloads (DPDK):</p> <ul style="list-style-type: none"> <li>• VXLAN encapsulation/decapsulation</li> <li>• VXLAN GPE encapsulation/decapsulation</li> <li>• Packet Modification: <ul style="list-style-type: none"> <li>• NAT (IPv4, IPv6, TCP, UDP)</li> <li>• Routing (L2)</li> <li>• TTL encapsulation/decapsulation</li> <li>• TCP SEQ and ACK</li> </ul> </li> <li>• ICMP offloading: match on ICMP type, code and identifier fields</li> <li>• HA and LB: LAG</li> <li>• Flow Aging</li> <li>• Flow Metering</li> <li>• Firmware fast update</li> </ul> |
| <b>SR-IOV Supported Hardware Offloads</b>                                   | <p>Added the following capabilities related to SR-IOV Supported Hardware Offloads:</p> <ul style="list-style-type: none"> <li>• VXLAN encapsulation/decapsulation</li> <li>• VXLAN GPE IPv4 &amp; IPv6 encapsulation/decapsulation</li> <li>• Packet Modification: <ul style="list-style-type: none"> <li>• NAT (IPv4, IPv6, TCP, UDP)</li> <li>• Routing (L2)</li> <li>• TTL encapsulation/decapsulation</li> <li>• TCP SEQ and ACK</li> </ul> </li> <li>• VF group rate limit support</li> <li>• VF and PF LAG</li> </ul>  |
| <b>Connection Tracking</b>  | Added metadata registry C “header modify” and “match rule” support to be used by the connection tracking functionality.  |
| <b>NODNIC Connectivity</b>  | Enabled NOIDNIC connectivity for VFs.  |
| <b>VDP</b>  | Enabled the option to read adapter’s VPD data (PN, SN etc) from the VM using lspci.  |
| <b>mlxconfig</b>  | Renamed the BOOT_RETRY_CNT1 parameter to BOOT_RETRY_CNT.   |
| <b>vPort: Multiple Guest Virtual Machine Identifier (GVMI) per Function</b> | <p>Added multiple e-switch vPort PCI function by GVMI, per vPort for the use cases below:</p> <ul style="list-style-type: none"> <li>• VM Live Migration SW fallback path</li> <li>• Container offloads</li> <li>• Scalable IOV</li> </ul>   |
| <b>Reduced Firmware Upgrade Time</b>  | <p>Reduced firmware upgrade time using mlxfwreset tool to ~3 seconds. Using this capability requires enabling PARTIAL_RESET_EN in mlxconfig and using MFT version 4.12.0 and up. The “PARTIAL” refers to not resetting the port modules (which is not mandatory for firmware upgrades).</p> <p><b>Note:</b> Currently this capability only supports firmware upgrade and downgrades to firmware versions newer than XX.25.1020.</p>  |
| <b>Bug Fixes</b>  | See <a href="#">Section 4, “Bug Fixes History”, on page 30</a>   |
| <b>Rev. 16.24.1000</b>  |  |
| <b>Layer 3 Encapsulation</b>  | Added support for an additional layer (Layer 3) of packet processing at the hypervisor level that enables adding and removing protocol headers (e.g., the MAC address is removed during encapsulation, and added during decapsulation) for the encapsulated traffic.   |
| <b>QoS Enhanced Transmission Selection (ETS)</b>                            | Enabled QoS ETS for systems with 64 VFs to better allocate bandwidth in the NIC.   |
| <b>Time to Live (TTL) Modification</b>                                      | Enabled TTL modification for received packets.   |

**Table 28 - Firmware Changes and New Feature History (Sheet 3 of 14)**

| Feature/Change   | Description  |
|--|--|
| <b>VF Mirroring</b>  | Mirrors the traffic from/to one VF to a dedicate admin VF for monitoring and traffic analysis. Note that in this process packets are duplicated and different packet modifications apply to different duplications.  |
| <b>Transmission Histogram Counters</b>                                   | Added support for the transmission histogram counter set as part of the Ethernet extended group counters.  |
| <b>Events Generation by the Hardware upon Counter Incrementation</b>     | Enabled the hardware to generate an event upon counter incrementation, in order to reduce an overhead from the software from reading rarely updated counters such as error counters.   |
| <b>NODNIC Connectivity</b>   | Enables NOIDNIC connectivity to the network through the e-switch and not directly to the physical port.  |
| <b>QP and Mkey Values</b>  | Enabled setting the QP and the Mkey values by the software upon these resources creation.  |
| <b>PCIe Atomic</b>   | Enabled advanced PCIe atomic operations. The HCA will perform PCIe atomic operations as a requestor towards the host memory when receiving compatible atomic messages from the network, and according to the configuration of NV_SW_OFFLOAD_CONFIG pci_atomic_mode field and the PCI AtomicOp Requester Enable bit in the Device Control 2 register. |
| <b>TTL Modification for Rx NIC (Steering)</b>                            | Enabled TTL modification in the Rx NIC steering. When modifying the TTL in the Rx NIC, the CQE checksum will not recalculated automatically.<br><b>Note:</b> TTL modification in the FDB for traffic from the network is currently not supported.  |
| <b>TIR Destination from the FDB</b>                                      | Enabled a single TIR destination from the FDB.   |
| <b>WRED</b>  | Changed the WRED default mode to OFF for Multi-Host adapter cards.   |
| <b>TX Steering Rule on in WQE Ethernet Segment</b>                       | Added support for TX steering rule on flow_table_metadata in WQE Ethernet segment.   |
| <b>L3 Encapsulation/Decapsulation in the Reformat Context Allocation</b> | Added L3 encapsulation/decapsulation support in the reformat context allocation. <ul style="list-style-type: none"> <li>L3 encapsulation removes L2 headers and adds generic L3 tunnel encapsulation.</li> <li>L3 decapsulation removes the generic L3 tunnel decapsulation and L2 header.</li> </ul>  |
| <b>Flow Steering Header Modification</b>                                 | Added support for flow steering header modification (header rewrite) for IPv4 TTL header for loopback traffic (VF-VF/VF-PF).<br><b>Note:</b> TTL modification for traffic from the network is currently not supported.   |
| <b>Teardown: Fast Mode</b>   | <b>[Developers only]</b> Moved the fast teardown HCA cap bit to offset 0x1c.4:1.   |
| <b>Virtual Functions/QoS</b>   | Enabled Virtual Functions to read QPDPM/QPDP/QPTS.   |
| <b>Message vs. Payload based flow control QP Configuration</b>           | Added support for requester QP packet based on E2E credits mode. The new flow control supports HCA-to-switch RDMA traffic packet-based End-2-End.  |
| <b>Multi PCI RDMA IB</b>   | This capability enables the user to expose two PCI/IB devices per network port.  |
| <b>Bug Fixes</b>   | See <a href="#">Section 4, “Bug Fixes History”, on page 30</a>   |
| <b>Steering</b>  | Enabled TTL modification in the RX NIC steering.<br><b>Note:</b> TTL modification in the FDB for traffic from the network is currently not supported.  |
| <b>Rev. 16.23.1020</b>   |  |
| <b>RoCE DC</b>   | Added support for Dynamically Connected Transport (DC) in RoCE in ConnectX-5 adapter cards.  |
| <b>Mini-CQE Compression</b>  | Enabled Rx mini-CQE compressed format for striding RQ.   |

**Table 28 - Firmware Changes and New Feature History (Sheet 4 of 14)**

| Feature/Change                                   | Description  |
|--|--|
| <b>TX Steering Rules</b>                         | Enabled TX steering rule on the application meta-data from the WQE. This functionality implements meta-data Reg A steering and meta-data Reg A and Reg B rewriting.  |
| <b>Tunneling Offload/steering Match Criteria</b> | Added MPLS over UDP and MPLS over GRE protocols for tunneling offload/steering match criteria.   |
| <b>InfiniBand</b>                                | Added support for IPoIB non-default Partition Keys (PKeys). Now the PKey values can be modified in the PKey table without the need of recreating the IPoIB (underlay) QPs.   |
| <b>SR-IOV in Multi-Host/Socket-Direct</b>        | <b>[Beta]</b> Added support for SR-IOV (up to 63 VFs) in Multi-Host/Socket-Direct.   |
| <b>Virtualization</b>                            | Reduced firmware's memory consumption to increase the supported number of VFs per PF to up to 100.   |
| <b>Tools/Driver Version</b>                      | Added support for QUERY_DRIVER_VERSION command. This command allows the PF driver to query its VFs driver version which was set by the SET_DRIVER_VERSION command.   |
| <b>Resiliency</b>                                | Shutting Down RDMA QPs with Excessive Retransmissions is a mechanism used to detect excessive retransmissions for an RC connection, and to close the connection in response to it. If the number of retransmissions due to a Local Ack Timeout, NAK-Sequence Error, or Implied NAK, during a specified period, exceeds the specified threshold, the QP will be handled as if the IB spec defined Retry Count was exceeded.   |
| <b>Diagnostic Counters</b>                       | Added new diagnostic counters to evaluate the number of ICMC hits and misses for particular resources.   |
| <b>Bug Fixes</b>                                 | See <a href="#">Section 4, "Bug Fixes History"</a> , on page 30  |
| <b>Rev. 16.22.1002</b>                           |  |
| <b>Disable SL/diff Flow</b>                      | Added support for disable SL/diff flow to avoid performance degradation for single queue using multiple priorities. This functionality should not be used when DCB (PFC, ETS) is enabled.  |
| <b>Software Reset Flow</b>                       | Software Reset Flow enables the device to recover from fatal errors. The flow includes software detection of a fatal error, automatic creations of an mstdump file for future debug by the software, and resetting of the device. The feature is enabled using an mlxconfig command.<br><br><b>Note:</b> The flow is currently not supported on Multi host devices, Socket Direct devices and devices running management traffic (NCSI, MCTP).   |
| <b>Steering Discard Packet Counters</b>          | Any received packet which is dropped by the device is accounted for. To enable this functionality, the following counters were added to count the discard packets (per vport): <ul style="list-style-type: none"> <li>• <code>nic_receive_steering_discard</code>: Number of packets that completed the NIC Receive Flow Table steering, and were discarded because they did not match any flow in the final Flow Table.</li> <li>• <code>receive_discard_vport_down</code>: Number of packets that were steered to a VPort, and discarded because the VPort was not in a state to receive packets.</li> <li>• <code>transmit_discard_vport_down</code>: Number of packets that were transmitted by a vNIC, and discarded because the VPort was not in a state to transmit packets.</li> </ul> |
| <b>Pause Frame Duration and XOFF Resend Time</b> | Increased the Pause Frame Duration and the XOFF Resend Time to the maximum value defined by the specification.   |
| <b>PCI Relax Ordering</b>                        | mlxconfig configuration can now enable or disable forced PCI relaxed ordering in <code>mkey_context</code> .<br>If this feature is enabled, the software per mkey configuration is ignored.  |

**Table 28 - Firmware Changes and New Feature History (Sheet 5 of 14)**

| Feature/Change                         | Description  |
|--|--|
| <b>Push/Pop VLAN</b>                   | Added support for Push/Pop VLAN, new FLOW TABLE ENTRY actions. These new actions are used by the driver to implement Q-in-Q functionality. For further information, please refer to the PRM section <i>Flow Table</i>  |
| <b>QOS "Rate Limit"</b>                | Added support to limit the transmission rate of individual InfiniBand port Service Levels. This capability is configurable through a new vendor-specific MAD (QosConfigSL). For more details see Mellanox Vendor Specific MAD Specification Rev 1.4 or above.  |
| <b>Packet Pacing</b>                   | Added support for Packet Pacing in ConnectX-5 adapter cards. Packet Pacing (traffic shaping) is a rate-limited flow per Send QPs. A rate-limited flow is allowed to transmit a few packets before its transmission rate is evaluated, and the next packet is scheduled for transmission accordingly. Setting and changing the rate is done by modifying the QP.<br><b>Note:</b> Packet Pacing is not functional in ConnectX-5 Multi host adapter cards.  |
| <b>vport Mirroring</b>                 | Packets are mirrored based on certain mirroring policy. The policy is set using the "set FTE command" that supports forward action in the ACL tables (ingress/egress). The firmware support the following destination list format:<br>1. new destination vport (analyzer)<br>2. another Flow Table<br><br>this way, the driver can forward the SX/RX packet related to the vport once it reaches the ACL table (forward it to the analyzer vport).   |
| <b>Resiliency: Special Error Event</b> | Firmware uses error events to monitor the health of core transport engines, both Rx and Tx, and to detect if a system hang occurred and was not cured by other error mechanisms. Upon such detection, events are sent to the driver to perform any required action (e.g., software reset).   |
| <b>QP's Creation Time</b>              | Accelerated QP's creation time.  |
| <b>SR-IOV LID based Routing Mode</b>   | SR-IOV default routing mode is now LID based. The configuration change is available via mlxconfig tool. Note that in such mode, the VF will get its own LID, hence the GRH is not required.<br><br><b>Note:</b> LID based routing support for vports is supported using SM v4.8.1  |
| <b>Expansion ROM</b>                   | Added PXE and UEFI to additional ConnectX- 5 adapter cards. ConnectX-5 now holds PXE and x86-UEFI  |
| <b>Host Chaining</b>                   | Host Chaining allows the user to connect ("chain") one server to another without going through a switch, thus saving switch ports.<br>Host Chaining algorithm is as follow: <ul style="list-style-type: none"> <li>Received packets from the wire with DMAC equal to the host MAC are forwarded to the local host</li> <li>Received traffic from the physical port with DMAC different than the current MAC are forwarded to the other port: <ul style="list-style-type: none"> <li>Traffic can be transmitted by the other physical port</li> <li>Traffic can reach functions on the port's Physical Function</li> </ul> </li> <li>Device allows hosts to transmit traffic only with its permanent MAC</li> <li>To prevent loops, the received traffic from the wire with SMAC equal to the port permanent MAC is dropped (the packet cannot start a new loop)</li> </ul><br>For Host Chaining limitation, see Known Issue #1178792 in <a href="#">Section 3, "Known Issues"</a> , on page 23 |
| <b>Fast path VLs</b>                   | Enabled fast path VLs which have lower latency (less than 2.55us) than slow path VLs. Fast path mapping can be configured using OpenSM configuration file.   |

**Table 28 - Firmware Changes and New Feature History (Sheet 6 of 14)**

| Feature/Change  | Description   |
|---|---|
| <b>Hairpin</b>  | Hairpin enables ingress traffic on the network port to egress on the same port or the 2nd port of the adapter.<br>Hairpin enables hardware forwarding of packets from the receive queue to the transmit queue, thus fully offloading software gateways to the hardware. The queues can be allocated on different PCI functions, thus enabling packets' forwarding between different NIC ports.  |
| <b>Coherent Accelerator Processor Interface (CAPI v2)</b> | The Coherent Accelerator Process Interface (CAPI) enables the user to attach a coherent accelerator to a Power and OpenPower based platforms. This solution delivers performance that exceeds today's I/O-attached acceleration engines.<br><b>Note:</b> This feature is available only with IBM Power 9 CPUs.  |
| <b>NVME-oF Target Offload over DC transport</b>           | The NVMe-oF target offload provides the IO data path functionality of an NVMe over Fabrics Front-End subsystem transferring the IO operations to NVMe PCIe subsystems.  |
| <b>Bug Fixes</b>  | See <a href="#">Section 4, "Bug Fixes History", on page 30</a>  |
| <b>Rev. 16.21.2010</b>                                    |   |
| <b>Query vPort Environments (Debug Counters)</b>          | Debug counters are a group of counters that handle traffic performance issue related to firmware overhead in transport flow.<br>The following are the additional counters added to this firmware version: <ul style="list-style-type: none"> <li>• current_q_under_processor_handle</li> <li>• total_q_under_processor_handle</li> <li>• qp_priority_update_flow</li> </ul>   |
| <b>Address Ordering</b>                                   | Enables address ordering on ConnectX-5 Multi-Host adapter cards. It allows PCIe READ transaction to bypass the PCIe WRITE transaction, when both transactions are not on the same page.   |
| <b>Bug Fixes</b>  | See <a href="#">Section 4, "Bug Fixes History", on page 30</a>  |
| <b>Rev. 16.21.1000</b>                                    |   |
| <b>Coherent Accelerator Processor Interface (CAPI v2)</b> | <b>[Beta]</b> The Coherent Accelerator Process Interface (CAPI) enables the user to attach a coherent accelerator to a Power and OpenPower based platforms. This solution delivers performance that exceeds today's I/O-attached acceleration engines.<br><b>Note:</b> This feature is available only with IBM Power 9 CPUs.  |
| <b>NVME-oF Target Offload</b>                             | Added support for NVMe over Fabrics (NVMe-oF) offload, an implementation of the new NVMe-oF standard target (server) side in the hardware.  |
| <b>Tag Matching RC/DC transport</b>                       | Added support for Tag Matching Offload with RC/DC transport.<br>In Tag Matching, the software holds a list of matching entries called matching list. Each matching entry contains a tag and a pointer to an application buffer. The matching list is used to steer arriving messages to a specific buffer according to the message tag. The action of traversing the matching list and finding the matching entry is called tag matching.<br>For further information, refer to the PRM section " <i>Tag Matching and Rendezvous Offload</i> " |
| <b>Hairpin</b>  | <b>[Beta]</b> Hairpin enables ingress traffic on the network port to egress on the same port or the 2nd port of the adapter.<br>Hairpin enables hardware forwarding of packets from the receive queue to the transmit queue, thus fully offloading software gateways to the hardware. The queues can be allocated on different PCI functions, thus enabling packets' forwarding between different NIC ports.  |



**Table 28 - Firmware Changes and New Feature History (Sheet 7 of 14)**

| Feature/Change   | Description   |
|--|---|
| <b>Virtual Extensible LAN (VXLAN) Encapsulation Offloads over RDMA in SR-IOV</b> | Added support for VXLAN encapsulation offloads over RDMA in SR-IOV. Virtual Extensible LAN (VXLAN) is a network virtualization technology that improves scalability problems associated with large cloud computing deployments. It tunnels Ethernet frames within Ethernet + IP + UDP frames.   |
| <b>Flow Table Entries Enhancements</b>   | Enabled adapter support for up to 1 million Flow Table Entries.<br>For further information, please refer to the PRM section <i>Flow Table</i>   |
| <b>Hardware Accelerated 802.1ad VLAN (Q-in-Q Tunneling)</b>                      | Q-in-Q tunneling allows the user to create a Layer 2 Ethernet connection between two servers. The user can segregate a different VLAN traffic on a link or bundle different VLANs into a single VLAN.   |
| <b>Memory Mapped to InterConnect (MEMIC)</b>                                     | Added support for locked ICMC data buffer in order to improve latency by saving the PCI 'round trip'.<br>For further information, refer to the PRM.   |
| <b>QoS: Tx Rate Limiter</b>  | Added support for VF rate limiter and bandwidth share in ConnectX-5.  |
| <b>Single PF for InfiniBand Dual Port Device</b>                                 | Added support for InfiniBand native (No SR-IOV) dual port device (Function per port is disabled).<br>In this mode virtualization is not supported and ISSI = 0.   |
| <b>Explicit Congestion Notification (ECN)</b>                                    | Enabled ECN by default.   |
| <b>RoCE Dual Port Mode (a.k.a Multi-Port vHCA: MPV)</b>                          | Enables the usage of a dual port Virtual HCA (vHCA) to share RDMA resources (e.g., MR, CQ, SRQ, PDs) across the two Ethernet (RoCE) NIC network ports and display the NIC as a dual port device.<br>For this feature to function properly, the following requirements must be met: <ul style="list-style-type: none"> <li>• Either the LAG or the Dual Port mode is enabled by the driver</li> <li>• Dual port device: both ports must be set as ETH</li> <li>• In ConnectX-4/ConnectX-4 Lx adapter cards, the maximum allowed number of VFs per PF is 32.</li> <li>• Function per port is enabled</li> </ul> <b>Note:</b> This feature is only supported in single host device |
| <b>DSCP</b>  | Added QPDPM register to support dynamic mapping between DSCP and priority.  |
|  | Added trust level for QoS prioritization according to the DSCP or PCP.  |
|  | Added ingress buffer management for: <ul style="list-style-type: none"> <li>• ingress traffic mapping to a buffer according to priority</li> <li>• buffers sizes and lossless parameters</li> </ul>   |
| <b>Steering Rules Rate Improvement</b>   | Improved steering rules update rate to up to 50K rules per sec.   |
| <b>Windows SR-IOV Enhanced eIPoIB</b>  | Enabled Windows SR-IOV Enhanced eIPoIB (without Secure Connection) for Windows-over-Windows setups.   |
| <b>Driver CR Dump</b>  | crdump operation takes a snapshot of the device's crspace dword-by-dword. It enables the driver to collect debug information upon firmware failure.   |
| <b>Secured Firmware Update</b>   | Secure Firmware Updates provides devices with the ability to verify digital signatures of new firmware binaries, in order to ensure that only officially approved versions are installed on the devices.<br><b>Note:</b> This feature is only available in adapter cards that support this feature.   |
| <b>Cables</b>  | Changed the default FEC mode for cables with attenuation 16 and below from RS to FC.  |
| <b>ECN</b>   | Enabled ECN (CongestionControl) by default for all priorities on Ethernet ports.  |
| <b>Bug Fixes</b>   | See <a href="#">Section 4, "Bug Fixes History", on page 30</a>  |
| <b>Rev. 16.20.1010</b>   |   |

**Table 28 - Firmware Changes and New Feature History (Sheet 8 of 14)**

| Feature/Change   | Description  |
|--|--|
| NVME-oF Target Offload                                     | <b>[Beta]</b> Performance and stability improvements.  |
| Tag Matching   | <b>[Beta]</b> Performance and stability improvements. For further information, refer to the PRM section “Tag Matching and Rendezvous Offload”  |
| Tag Matching DC  | <b>[Alpha]</b> Added support for Tag Matching DC (RNDV/EAGER traffic)  |
| Adaptive Routing (Out-Of-Order)                            | Added support for send/receive Out-of-Order RDMA packets required by the Adaptive Routing. Adaptive Routing (AR) enables the switch to select the output port based on the port's load.  |
| DSCP   | Added trust level for QoS prioritization according to the DSCP or PCP.   |
|  | <b>[Beta]</b> Added ingress buffer management for: <ul style="list-style-type: none"> <li>• ingress traffic mapping to a buffer according to priority</li> <li>• buffers sizes and lossless parameters</li> </ul>  |
| Secured Firmware Updates                                   | <b>[Beta]</b> Secure Firmware Updates provides devices with the ability to verify digital signatures of new firmware binaries, in order to ensure that only officially approved versions are installed on the devices.<br><b>Note:</b> This feature is only available in adapter cards that support this feature.        |
| Multi-Host/Socket Direct Routing to be LID based           | <b>[InfiniBand only]</b> Changed the Multi-Host/Socket Direct routing to be LID based instead of GID based. Thus, GRH/GID index is not required.<br><b>Note:</b> This feature requires SM 4.8.1 and above.   |
| Relaxed Ordering   | <b>[Beta]</b> Added support for relaxed ordering write in memory keys.   |
| RDMA Counters  | Enhanced RDMA counter  |
| TLV for PCI class code                                     | Added 2 new per Host TLVs (see <a href="#">Table 30, “Supported Non-Volatile Configurations,” on page 51</a> )   |
| Header Rewrite   | Modifies the packets header.   |
| Fast Teardown  | Enables fast unloading driver by using Teardown HCA with op_mode=1 (force_close). For further information, refer to the PRM.   |
| IPoIB Virtualization                                       | Added support for enhanced IPoIB (QP.ulp == 2) in virtualized system (SR-IOV / Multi-Host / Socket Direct)   |
| SFP Power Flow Improvement (level 2,1)                     | Added support for SFP power class.   |
| Bug Fixes  | See <a href="#">Section 4, “Bug Fixes History”, on page 30</a>   |
| <b>Rev. 16.19.1200</b>                                     |  |
| General  | This is the first GA version of the ConnectX-5/Ex adapter cards.   |
| Bug Fixes  | See section <a href="#">Section 4, “Bug Fixes History”, on page 30</a>   |
| <b>Rev. 16.18.2000 (Beta)</b>                              |  |
| Virtual Extensible LAN (VXLAN) encapsulation/decapsulation | <b>[Beta Level]</b> Virtual Extensible LAN (VXLAN) is a network virtualization technology that improves scalability problems associated with large cloud computing deployments. It tunnels Ethernet frames within Ethernet + IP + UDP frames. Mellanox implements VXLAN encapsulation and decapsulation in the hardware. |
| NVME-oF Target Offload                                     | <b>[Beta Level]</b> Added support for NVMe over Fabrics (NVME-oF) offload, an implementation of the new NVME-oF standard target (server) side in hardware.   |
| Tag Matching   | <b>[Beta Level]</b> Added support for offloading MPI tag matching to HCA. It matches the source send operations to the destination receivers.  |

**Table 28 - Firmware Changes and New Feature History (Sheet 9 of 14)**

| Feature/Change   | Description   |
|--|---|
| <b>VLAN Switch Tagging (VST)</b>                           | Enables the virtual machine interface to have no VLAN tag over it, thus allowing VLAN tagging to be handled by the Hypervisor.  |
| <b>On Demand Paging (ODP)</b>                              | Added supported for Demand Paging (ODP).  |
| <b>Rev. 16.18.1000 (Beta)</b>                              |   |
| <b>NVM Express over Fabrics (NVMeF)</b>                    | NVMeF is a protocol for communicating block storage IO requests over RDMA. For further information, please refer to the PRM section <i>“NVMe over Fabric Target Application Offload (NVMeF)”</i> .  |
| <b>Tag Matching</b>  | In Tag Matching, the software holds a list of matching entries called matching list. Each matching entry contains a tag and a pointer to an application buffer. The matching list is used to steer arriving messages to a specific buffer according to the message tag. The action of traversing the matching list and finding the matching entry is called tag matching.<br>For further information, refer to the PRM section <i>“Tag Matching and Rendezvous Offload”</i> |
| <b>RX Loss (BaseT link down indication)</b>                | Added logical link indication in SFP to BaseT modules and disabled logical link when peer port is down.   |
| <b>SFP Rate</b>  | Added support for 10GbE in 25GbE SFP optical modules  |
| <b>PDDR</b>  | Enables mlxlink tool to collect data on the PHY link status and provides link down reasons and additional link related information.   |
| <b>KR Tx Response</b>                                      | Enabled TX configuration response and movement during Link Training in Ethernet.  |
| <b>Phy Test mode</b>                                       | Added support at lane rate of 12.89Gb.  |
| <b>Performance</b>   | Improved performance for Send Queues (SQs) transmitting multiple priorities in a single Traffic Class (TC) configuration.   |
| <b>Dropless TCP</b>  | Added the ability to avoid packet drops due to temporary lack of posted Receive buffers (WQEs), for trusted Receive Queues (RQs).   |
| <b>Head of Queue (HoQ) per TC</b>                          | Limits the amount of time a packet may head a Traffic Class (TC) transmission queue, without being transmitted. Stale packets are discarded.<br>Active by default for TCs adhering to link level flow control   |
| <b>User Access Region (UAR) 4KB Granularity Allocation</b> | UAR page size currently is set to 4KB and not according to what the system page size determines.  |
| <b>No Driver NIC (NODNIC) Performance Improvement</b>      | Improved performance of: <ul style="list-style-type: none"> <li>• Doorbell from User Access Region (UAR)</li> <li>• Clear interrupt from User Access Region (UAR)</li> </ul>  |
| <b>Counters</b>  | Added support for additional transport counters.  |
| <b>Scatter to CQE on Sender for DC</b>                     | Enabled scatter-to-CQE for sent packets for DC.   |
| <b>CQ modify</b>   | Enabled moderation period modification in CQ modify command.  |
| <b>Network traffic between UEFI-Shell and OS</b>           | Enabled network traffic between UEFI-Shell and OS.  |
| <b>non-RDMA capable VFs</b>                                | Enabled the PF to force disable RoCE for its VFs.   |
| <b>Loopback Enabled/Disabled</b>                           | Enabled VNIC the control to enable/disable its local loopback traffic.  |
| <b>RDMA RX Flow Table</b>                                  | Added the option to open a receive RDMA Flow Table and to forward RoCE traffic to some destination QP.  |

**Table 28 - Firmware Changes and New Feature History (Sheet 10 of 14)**

| Feature/Change                                 | Description  |
|--|--|
| <b>GENEVE &amp; IP-in-IP Stateless Offload</b> | <p>Added support for IP-in-IP and GENEVE network protocols encapsulated into IP frame (L2 tunneling).</p> <p>Encapsulation is suggested as a means to alter the normal IP routing for datagrams, by delivering them to an intermediate destination that would otherwise not be selected based on the (network part of the) IP Destination Address field in the original IP header.</p> <p><b>Note:</b> For driver support, please see the Release Notes/User Manual of the relevant OS driver.</p>   |
| <b>Resilient RoCE</b>                          | <p>Resilient RoCE is the ability to send RoCE traffic over a lossy network (a network without flow control enabled), without the need to enable flow control on the network.</p> <p>The ability is accomplished by enabling ECN on both the Switch and the Host.</p>   |
| <b>Power MGMT</b>                              | <p>Added support for PCIe Express standard "Slot capability register" message (PCIe base rev 3.1, section 6.9 – "Slot Power Limit Control")</p> <p>When ConnectX-5 Ex based adapter is inserted to a PCIe slot that supports the reporting of the slot power limit control, the ConnectX-5 Ex may disable the 2nd port if PCIe slot message, showing that the power in this slot is insufficient. If not, both ports will stay enabled.</p> <p>In cases where ConnectX-5 Ex based adapter is inserted to a PCI slot that <b>DOES NOT</b> support the "Slot capability register" message, and the adapters' configuration is 2 active optic cables/ transceivers, only one port will be enabled (the first inserted optic).</p> <p>Custom and OEM branded card based on ConnectX-5 Ex may be configured by INI to support/not-support the Power management feature.</p> <p>In hosts which do not support the "Slot capability register" Message and have enough power to support 2 active optical cable, the user will have the option to override the configuration resulted from "Slot capability register" by running the following NVconfig command:</p> <ul style="list-style-type: none"> <li>• <code>echo "MLNX_RAW_TLV_FILE" &gt; /tmp/power_conf_tlv.cfg;</code></li> <li>• <code>echo "0x00000004 0x00000088 0x00000000 0xc0000000" &gt;&gt; /tmp/power_conf_tlv.cfg</code></li> <li>• <code>mlxconfig -d &lt;device&gt; -f /tmp/power_conf_tlv.cfg set_raw</code></li> <li>• <code>mlxfwreset -d &lt;device&gt; reset</code></li> </ul> <p>For details on ConnectX-5 Ex power, please refer to ConnectX-5 Ex Datasheet</p> |
| <b>Virtual Functions (VF) per Port</b>         | <p>The maximum Virtual Functions (VF) per port is 64.</p> <p><b>Note:</b> When increasing the number of VFs, the following limitations must be taken into consideration:</p> <pre>server_total_bar_size &gt;= (num_pfs)*(2log_pf_uar_bar_size + 2log_vf_uar_bar_size*total_vfs) server_total_msix &gt;= (num_pfs)*(num_pf_msix + num_vfs_msix *total_vfs)</pre> <p><b>Note:</b> For the maximum number of VFs supported by your driver, please refer to your drivers' Release Notes or User Manual.</p>  |
| <b>QoS per VFs</b>                             | <p><b>[InfiniBand Only]</b> Added support for multiple VLs in SR-IOV/mutlihost environments.</p> <p><b>Note:</b> The number of VLs can be configured by the NVCONFIG. The default VL number is 4 VLs.</p>  |
| <b>HCA Port Flap Counter</b>                   | <p>Added support for Port Flap Counter.</p>  |
| <b>Fixed Buffer Size (KSM)</b>                 | <p>Limits the buffer size for all entries to improve performance.</p> <p>KSM is used when associating Key Length My Virtual Address (KLMs) with fixed memory size.</p>   |

**Table 28 - Firmware Changes and New Feature History (Sheet 11 of 14)**

| Feature/Change   | Description  |
|--|--|
| <b>NULL Mkey</b>   | This entry (null_mkey) is use to indicate non-present KLM/KSM entries. When accessing is, it causes the device to generate page fault event.   |
| <b>Out-of-Band Online Firmware Update: Firmware Update over PLDM</b> | PLDM firmware burning is based on the DMTF spec DSP0267 (draft 9). The feature enables upgrading firmware and expansion ROM images using the PLDM protocol over MCTP (over PCIe). By doing so, a supporting BMC can query and upgrade the firmware without using OS based tools.   |
| <b>New Group in Ports Performance Counters (PPCNT)</b>               | Added a new physical layer statistics counters group. The new group includes BER counters, FEC error correction, clear time, and additional physical layer counters. For further information, please refer to the <a href="#">Ethernet Adapters Programming Manual (PRM)</a> .   |
| <b>Permanent Link Up Mode</b>  | Enables the user to set a certain link up state for an unlimited period of time. This mode has 3 states: <ul style="list-style-type: none"> <li>• Aux power (standby)</li> <li>• Reboot/boot/driver unloaded - the server is active and no driver is up</li> <li>• Driver is up - at least one driver is up (the time between init HCA and teardown or FLR)</li> </ul> |
| <b>No Driver NIC (NODNIC) Performance Improvement</b>                | Added support for: <ul style="list-style-type: none"> <li>• Doorbell from User Access Region (UAR)</li> <li>• Clear interrupt from User Access Region (UAR)</li> </ul>   |
| <b>Firmware Resiliency: Suppress Pauses</b>                          | Allows the user to configure the adapter card to stop sending pauses after x when the receive port is unavailable (in a hang state).   |
| <b>Performance Back-pressure Counters</b>                            | Added support for new performance counters.  |
| <b>Data Center Bridging Exchange (DCBX)</b>                          | DCBX is used by DCB devices to exchange configuration information with directly connected peers. DCBX uses Link Layer Discovery Protocol (LLDP) to exchange parameters between two link peers. For further information, please refer to the PRM.   |
| <b>Access Register: Default Values Revert</b>                        | Allows network port registers to revert to their default values when the driver is restarted or the host is rebooted.  |
| <b>Link up Modes</b>   | Added additional network link up modes. The new modes decide when to keep the network link up. The new modes are: <ul style="list-style-type: none"> <li>• keep_eth_link_up</li> <li>• keep_ib_link_up</li> <li>• keep_link_up_on_boot</li> <li>• keep_link_up_on_standby</li> </ul>   |
| <b>Explicit Congestion Notification (ECN)</b>                        | Explicit Congestion Notification (ECN) is an extension to the Internet Protocol and to the Transmission Control Protocol. ECN allows end-to-end notification of network congestion without dropping packets.   |
| <b>RoCE Link Aggregation (RoCE LAG)</b>                              | RoCE Link Aggregation provides failover and link aggregation capabilities. In this mode, only one IB port, that represents the two physical ports, is exposed to the application layer. For further information, please refer to the PRM.  |
| <b>OVS Offload</b>   | Mellanox Accelerated Switching And Packet Processing (ASAP <sup>2</sup> ) Direct technology allows to offload OVS by handling OVS data-plain in Mellanox NIC hardware (Mellanox Embedded Switch or eSwitch) while maintaining OVS control-plain unmodified.  |
| <b>FCS no scatter / FCS check</b>                                    | Enables the user to control whether or not to scatter Frame Check Sequence (FCS) or to check FCS functionality.  |

**Table 28 - Firmware Changes and New Feature History (Sheet 12 of 14)**

| Feature/Change  | Description   |
|---|---|
| <b>PRBS Patterns Generation and Tuning</b>                      | A new PHY test mode in which the device can generate different PRBS patterns for SerDes tuning purpose. For further information, please refer to PRM registers: PPAOS, PPTT, PPRT.  |
| <b>Management Controller Transport Protocol (MCTP) over PCI</b> | Added support for MCTP host management over PCI   |
| <b>OCBB / OCSD support after mlxfwreset</b>                     | Added support for OCBB/OCSD memory pointers restoration after mlxfwreset  |
| <b>MCTP media migration</b>                                     | Added support for MCTP media migration between SMBUS and PCI  |
| <b>Cables</b>   | Removed the RX amplitude configuration on some cable types  |
| <b>IPoIB checksum and LSO offload</b>                           | Added IPoIB checksum and LSO offload support  |
| <b>Scatter FCS in RQ</b>  | Enables software to scatter or strip FCS in RQ.   |
| <b>CQE Time Stamping</b>  | Keeps track of the creation of a packet. A time-stamping service supports assertions of proof that a datum existed before a particular time.  |
| <b>Priority Flow Control (PFC)</b>                              | Applies pause functionality to specific classes of traffic on the Ethernet link.  |
| <b>RDMA retransmission counters</b>                             | Custom port counters provide the user a clear indication about RDMA send/receive statistics and errors.   |
| <b>Link Layer Discovery Protocol (LLDP)</b>                     | The Link Layer Discovery Protocol (LLDP) is a vendor-neutral Link Layer protocol in the Internet Protocol Suite used by network devices for advertising their identity, capabilities, and neighbors on a IEEE 802 LAN. The protocol is formally defined in IEEE 802.1AB.  |
| <b>Flow Steering Counters</b>                                   | Provides a clear indication of Flow Steering statistics and errors.   |
| <b>WQE Inline Header</b>  | The minimal amount of packet headers inlined in the WQE's Eth Segment.  |
| <b>table-miss Flow</b>  | A flow table may include a table-miss flow entry, which renders all Match Fields wildcards. If a packet does not match a flow entry in a flow table, this is a table miss. The behavior on a table miss depends on the table configuration. A table-miss flow entry in the flow table may specify how to process unmatched packets. |
| <b>Strided WQE User Space</b>                                   | Striding RQ is a receive queue comprised by work queue elements (i.e. WQEs), where multiple packets of LRO segments (i.e. message) are written to the same WQE.   |
| <b>SR-IOV (EN eSwitch &amp; RoCE)</b>                           | Single Root IO Virtualization (SR-IOV) is a technology that allows a physical PCIe device to present itself multiple times through the PCIe bus.  |
| <b>Vector Calculation/ Erasure Coding Offload</b>               | Uses the HCA for offloading erasure coding calculations.  |
| <b>Link params modification via access registers</b>            | The change includes the following:<br><ol style="list-style-type: none"> <li>1. Changed port configuration which required link re-training (such as speed)</li> <li>2. PAOS down</li> <li>3. PAOS up</li> </ol> This change, will cause the link to toggle and new configurations to take effect.                                   |
| <b>Checksum Calculation on Image/Device</b>                     | Flint utility allows performing an MD5 checksum on the non-persistent sections of the firmware image.<br>For further information, please refer to <a href="#">MFT User Manual</a> .   |
| <b>Port Link</b>  | Reduced the port link-up time when negotiating according to Clause 73 (DME)   |

**Table 28 - Firmware Changes and New Feature History (Sheet 13 of 14)**

| Feature/Change                    | Description   |
|-----------------------------------|---|
| <b>Ethernet Network</b>           | <ul style="list-style-type: none"> <li>• Large Receive Offload (LRO)</li> <li>• Large Send Offload (LSO)</li> <li>• Receive Side Scaling (RSS)</li> <li>• Global Pause</li> <li>• RoCEv1.0/RoCEv2.0</li> <li>• Flow Steering</li> <li>• Sniffer Ethernet</li> <li>• Multi packet WQE</li> <li>• Minimal Bandwidth Guarantee (ETS)</li> <li>• Explicit Congestion Notification (ECN)</li> <li>• Priority Flow Control (PFC)</li> </ul>   |
| <b>PCI</b>                        | <ul style="list-style-type: none"> <li>• PCIe Function Level Reset (FLR)</li> <li>• Power Management L2/L3 flow support</li> </ul>  |
| <b>PRM</b>                        | <ul style="list-style-type: none"> <li>• Strided SRQ</li> <li>• Self Loopback support</li> <li>• Transport Domain support</li> <li>• CQ2EQ remapping</li> <li>• Added support for the following commands:               <ul style="list-style-type: none"> <li>• MODIFY/QUERY_ESW_VPORT_CONTEXT</li> <li>• QUERY/MODIFY_CONG_STATUS</li> <li>• QUERY/MODIFY_CONG_PARAMS</li> <li>• QUERY_CONG_STATISTICS</li> <li>• ADD/DELETE_VXLAN_UDP_DPORT</li> </ul> </li> </ul>   |
| <b>Virtualization</b>             | <ul style="list-style-type: none"> <li>• VXLAN/NVGRE Stateless offload</li> <li>In this release, NVGRE is supported through Windows ONLY</li> <li>• SR-IOV EN</li> </ul>  |
| <b>Performance</b>                | <ul style="list-style-type: none"> <li>• CQE zipping</li> </ul>   |
| <b>Misc</b>                       | <ul style="list-style-type: none"> <li>• Wake-on-Lane/Standby</li> <li>• FlexBoot/UEFI support</li> </ul>   |
| <b>Non-Volatile Configuration</b> | <ul style="list-style-type: none"> <li>• Non-Volatile Configuration (NVConfig). For the complete list, please refer to <a href="#">Section 8, on page 51</a>.</li> </ul>  |
| <b>Port management</b>            | <ul style="list-style-type: none"> <li>• Enabled port management. Now one port can be set as Ethernet and one as InfiniBand.</li> </ul>   |
| <b>InfiniBand Network</b>         | <ul style="list-style-type: none"> <li>• Dynamically Connected (DC) transport</li> <li><b>Note:</b> There is no interoperability between ConnectX-5 and ConnectX-4 adapter cards when using DC.</li> <li>• Unreliable Datagram Connection transport</li> <li>• Atomic Operation</li> <li>• CORE-Direct®               <ul style="list-style-type: none"> <li>• Provides Collective Off-loading in HCA</li> <li>• Frees CPU to perform computation in parallel with collective operations</li> </ul> </li> <li>• T10 DIF pipeline Data Integrity Signature off-loading (at beta level)</li> <li>• User Memory Registration (UMR)</li> <li>• Automatic Path Migration</li> <li>• On Demand Paging (ODP) - Memory can now be used without pinning memory beforehand.</li> <li>• Congestion Control</li> <li>• Shrink Address Vectors for RC and UD</li> <li>• Programmable Port/Node GUID</li> </ul> |

**Table 28 - Firmware Changes and New Feature History (Sheet 14 of 14)**

| Feature/Change         | Description  |
|------------------------|--|
| <b>General</b>         | <ul style="list-style-type: none"> <li>• Thermal monitoring and protection</li> <li>• Port LEDs indications</li> <li>• NVConfig Tool</li> <li>• Suspend to RAM (S3) support</li> <li>• Diagnostic counters vendor-specific MAD support, as defined by VS-MAD spec version 1.2</li> <li>• Physical Port Counter - Beta level</li> <li>• Q Counter - Beta level</li> <li>• Firmware burning (using mstflint) when the driver is down</li> <li>• CPLD field upgrade</li> <li>• V Port commands</li> </ul> |
| <b>Host management</b> | <ul style="list-style-type: none"> <li>• NC-SI over RMIII support</li> </ul>   |
| <b>MAD</b>             | <ul style="list-style-type: none"> <li>• Config space address in MAD management class 0x09</li> </ul>  |



## **6 PreBoot Drivers (FlexBoot/UEFI)**

### **6.1 FlexBoot Changes and New Features**

For further information, please refer to FlexBoot Release Notes

<https://docs.mellanox.com> --> Software --> Firmware Management --> PreBoot Drivers (FlexBoot/UEFI)

### **6.2 UEFI Changes and Major New Features**

For further information, please refer to UEFI Release Notes

<https://docs.mellanox.com> --> Software > Firmware Management --> PreBoot Drivers (FlexBoot/UEFI)

## 7 Unsupported Features and Commands

### 7.1 Unsupported Features

The following advanced feature are unsupported in the current firmware version:

**Table 29 - Unsupported Features**

| Feature  | ConnectX-5 |
|--|------------|
| The following service types: <ul style="list-style-type: none"> <li>• SyncUMR</li> <li>• Mellanox transport</li> <li>• RAW IPv6</li> </ul> | No         |
| INT-A not supported for EQs only MSI-X   | No         |
| PCI VPD write flow (RO flow supported)   | No         |
| Streaming Receive Queue (STRQ) and collapsed CQ  | No         |
| Subnet Manager (SM) on VFs   | No         |
| RoCE LAG in Multi-Host/Socket-Direct   | No         |

### 7.2 Unsupported Commands

- QUERY\_MAD\_DEMUX
- SET\_MAD\_DEMUX
- CREATE\_RQ - MEMORY\_RQ\_RMP
- MODIFY\_LAG\_ASYNC\_EVENT

## 8 Supported Non-Volatile Configurations

**Table 30 - Supported Non-Volatile Configurations**

| Configuration             | mlxconfig Parameter Name        | Class             | TLV ID |
|---------------------------|---------------------------------|-------------------|--------|
| NV_MEMIC_CONF             | MEMIC_BAR_SIZE                  | GLOBAL (0)        | 0x6    |
|                           | MEMIC_SIZE_LIMIT                |                   |        |
| NV_HOST_CHAINING_CONF     | HOST_CHAINING_MODE              |                   | 0x8    |
|                           | HOST_CHAINING_DESCRIPTOR        |                   |        |
|                           | HOST_CHAINING_TOTAL_BUFFER_SIZE |                   |        |
| NV_FLEX_PARS_CONF         | FLEX_PARSER_PROFILE_ENABLE      |                   | 0xe    |
|                           | FLEX_IPV4_OVER_VXLAN_PORT       |                   |        |
| NV_ROCE_1_5_CONF          | ROCE_NEXT_PROTOCOL              |                   | 0x10   |
| NV_INTERNAL_RESOURCE_CONF | ESWITCH_HAIRPIN_DESCRIPTOR      |                   | 0x13   |
|                           | ESWITCH_HAIRPIN_TOT_BUFFER_SIZE |                   |        |
| NV_GLOBAL_PCI_CONF        | NON_PREFETCHABLE_PF_BAR         |                   | 0x80   |
|                           | NUM_OF_VFS                      |                   |        |
|                           | SRIOV_EN                        |                   |        |
|                           | PF_LOG_BAR_SIZE                 |                   |        |
|                           | VF_LOG_BAR_SIZE                 |                   |        |
|                           | NUM_PF_MSIX                     |                   |        |
|                           | NUM_VF_MSIX                     |                   |        |
| NV_TPT_CONF               | INT_LOG_MAX_PAYLOAD_SIZE        |                   | 0x82   |
| NV_POWER_CONF             | SW_RECOVERY_ON_ERRORS           | 0x88              |        |
|                           | RESET_WITH_HOST_ON_ERRORS       |                   |        |
|                           | ADVANCED_POWER_SETTINGS         |                   |        |
| NV_SW_OFFLOAD_CONFIG      | CQE_COMPRESSION                 | 0x10a             |        |
|                           | IP_OVER_VXLAN_EN                |                   |        |
|                           | PCI_ATOMIC_MODE                 |                   |        |
|                           | LRO_LOG_TIMEOUT0                |                   |        |
|                           | LRO_LOG_TIMEOUT1                |                   |        |
|                           | LRO_LOG_TIMEOUT2                |                   |        |
|                           | LRO_LOG_TIMEOUT3                |                   |        |
| NV_IB_DC_CONF             | LOG_DCR_HASH_TABLE_SIZE         | 0x190             |        |
|                           | DCR_LIFO_SIZE                   |                   |        |
| NV_VPI_LINK_TYPE          | LINK_TYPE                       | PHYSICAL_PORT (2) | 0x12   |

**Table 30 - Supported Non-Volatile Configurations**

| Configuration   | mlxconfig Parameter Name      | Class             | TLV ID |
|-----------------|-------------------------------|-------------------|--------|
| NV_ROCE_CC      | ROCE_CC_PRIO_MASK             | PHYSICAL_PORT (2) | 0x107  |
|                 | ROCE_CC_ALGORITHM             |                   |        |
| NV_ROCE_CC_ECN  | CLAMP_TGT_RATE_AFTER_TIME_INC | PHYSICAL_PORT (2) | 0x108  |
|                 | CLAMP_TGT_RATE                |                   |        |
|                 | RPG_TIME_RESET                |                   |        |
|                 | RPG_BYTE_RESET                |                   |        |
|                 | RPG_THRESHOLD                 |                   |        |
|                 | RPG_MAX_RATE                  |                   |        |
|                 | RPG_AI_RATE                   |                   |        |
|                 | RPG_HAI_RATE                  |                   |        |
|                 | RPG_GD                        |                   |        |
|                 | RPG_MIN_DEC_FAC               |                   |        |
|                 | RPG_MIN_RATE                  |                   |        |
|                 | RATE_TO_SET_ON_FIRST_CNP      |                   |        |
|                 | DCE_TCP_G                     |                   |        |
|                 | DCE_TCP_RTT                   |                   |        |
|                 | RATE_REDUCE_MONITOR_PERIOD    |                   |        |
|                 | INITIAL_ALPHA_VALUE           |                   |        |
|                 | MIN_TIME_BETWEEN_CNPS         |                   |        |
|                 | CNP_802P_PRIO                 |                   |        |
| CNP_DSCP        |                               |                   |        |
| NV_LLDP_NB_CONF | LLDP_NB_DCBX                  | PHYSICAL_PORT (2) | 0x10a  |
|                 | LLDP_NB_RX_MODE               |                   |        |
|                 | LLDP_NB_TX_MODE               |                   |        |
| NV_LLDP_NB_DCBX | DCBX_IEEE                     | PHYSICAL_PORT (2) | 0x18e  |
|                 | DCBX_CEE                      |                   |        |
|                 | DCBX_WILLING                  |                   |        |
| NV_KEEP_LINK_UP | KEEP_ETH_LINK_UP              | PHYSICAL_PORT (2) | 0x190  |
|                 | KEEP_IB_LINK_UP               |                   |        |
|                 | KEEP_LINK_UP_ON_BOOT          |                   |        |
|                 | KEEP_LINK_UP_ON_STANDBY       |                   |        |
| NV_QOS_CONF     | NUM_OF_VL                     | PHYSICAL_PORT (2) | 0x192  |
|                 | NUM_OF_TC                     |                   |        |
|                 | NUM_OF_PFC                    |                   |        |
| NV_MPFS_CONF    | DUP_MAC_ACTION                | PHYSICAL_PORT (2) | 0x196  |
|                 | SRIOV_IB_ROUTING_MODE         |                   |        |
|                 | IB_ROUTING_MODE               |                   |        |

**Table 30 - Supported Non-Volatile Configurations**

| Configuration           | mlxconfig Parameter Name | Class             | TLV ID |
|-------------------------|--------------------------|-------------------|--------|
| NV_HCA_CONF             | PCI_WR_ORDERING          | HOST-FUNCTION (3) | 0x112  |
|                         | MULTI_PORT_VHCA_EN       |                   |        |
| NV_EXTERNAL_PORT_CTRL   | PORT_OWNER               |                   | 0x192  |
|                         | ALLOW_RD_COUNTERS        |                   |        |
|                         | RENEG_ON_CHANGE          |                   |        |
|                         | TRACER_ENABLE            |                   |        |
| NV_ROM_BOOT_CONF2       | IP_VER                   |                   | 0x195  |
|                         | BOOT_UNDI_NETWORK_WAIT   |                   |        |
| NV_ROM_UEFI_CONF        | UEFI_HII_EN              |                   | 0x196  |
| NV_ROM_UEFI_DEBUG_LEVEL | BOOT_DBG_LOG             |                   | 0x206  |
|                         | UEFI_LOGS                |                   |        |
| NV_ROM_BOOT_CONF1       | BOOT_VLAN                |                   | 0x221  |
|                         | LEGACY_BOOT_PROTOCOL     |                   |        |
|                         | BOOT_RETRY_CNT           |                   |        |
|                         | BOOT_LACP_DIS            |                   |        |
|                         | BOOT_VLAN_EN             |                   |        |
| NV_ROM_IB_BOOT_CONF     | BOOT_PKEY                | 0x222             |        |
| NV_PCI_CONF             | ADVANCED_PCI_SETTINGS    | HOST (7)          | 0x80   |
| SAFE_MODE_CONF          | SAFE_MODE_THRESHOLD      |                   | 0x82   |
|                         | SAFE_MODE_ENABLE         |                   |        |