ConnectX-5 Adapter Cards Firmware Release Notes v16.27.2008

Table of Contents

Release Notes Update History

Revision	Date	Description
16.27.2008	April 06, 2020	Initial release of this Release Notes version, This version introduces Changes and New Features and Bug Fixes.

Overview

Firmware which is added at the time of manufacturing, is used to run user programs on the device and can be thought of as the software that allows hardware to run. Embedded firmware is used to control the functions of various hardware devices and systems, much like a computer's operating system (OS) controls the function of software applications. Firmware may be written into read-only memory (ROM), erasable programmable read-only memory (EPROM) or flash memory.

Firmware Download

Please visit www.mellanox.com → Support & Education → Firmware Download

Document Revision History

A list of the changes made to this document are provided in Document Revision History.

1 Firmware Compatible Products

These are the release notes for the ConnectX®-4 adapters firmware Rev

Unable to render include or excerpt-include. Could not retrieve page.

. This firmware supports the following protocols:

- InfiniBand - SDR, QDR, FDR10, FDR, EDR
- Ethernet - 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 56GbE[1], 100GbE
- PCI Express 3.0, supporting backwards compatibility for v2.0 and v1.1

[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.

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)	Present (Enabled)	Exists
MCX566A-CDAB	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 (Enabled)	Present (Enabled)	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

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 (Enabled)	Present (Enabled)	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

Device Part Number	PSID	Device Name	FlexBoot	UEFI x86	UEFI ARM	Enable/ disable exprom Feature
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
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-ACAT	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-ACAN	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-ECAN	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

Device Part Number	PSID	Device Name	FlexBoot	UEFI x86	UEFI ARM	Enable/ disable exprom Feature
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
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
MNV303212A-ADLT	MT_0000000158	Innova-2 Flex Open for Application Acceleration, dual-port SFP28, 25GbE, KU15P, 8GB, No Crypto, PCI4.0 x8, HHHL, active heat sink, tall bracket	Present (Enabled)	Present (Disabled)	Not Present	Exists
MNV303611A-EDLT	MT_0000000250	Innova-2 Flex Open VPI, dual-port QSFP28, EDR / 100GbE, KU15P, No memory, No Crypto, PCI4.0 x8, HHHL, passive heat sink, tall bracket	Present (Enabled)	Present (Disabled)	Not Present	Exists
MCX566M-GDAI	MT_0000000262	ConnectX®-5 Ex EN network interface card for OCP 3.0 with Multi-Host, with host management, 50GbE Dual-port QSFP28, PCIe 4.0/3.0 x16, Internal Lock bracket	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX515A-CCUT	MT_0000000519	ConnectX®-5 EN network interface card, 100GbE single-port QSFP28, PCIe3.0 x16, UEFI Enabled (ARM, x86), tall bracket	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists

Device Part Number	PSID	Device Name	FlexBoot	UEFI x86	UEFI ARM	Enable/ disable exprom Feature
MCX556A-ECUT	MT_0000000504	ConnectX®-5 VPI adapter card, EDR IB (100Gb/s) and 100GbE, dual-port QSFP28, PCIe3.0 x16, UEFI enabled, tall bracket	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX565M-CDAI	MT_0000000347	ConnectX®-5 Ex EN network interface card for OCP 3.0, with Multi-Host and host management, 100GbE Single-port QSFP28, PCIe4.0 x16, Internal Lock bracket	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists

1.1 Supported Mellanox 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.

1.1.1 Validated and Supported QDR Cables

Speed	Cable OPN	Description
QDR	MC2206125-007	Mellanox® passive copper cable, IB QDR, 40Gb/s, QSFP,7m

1.1.2 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
FDR10	MC2206310-003	Mellanox® active fiber cable, IB QDR/FDR10, 40Gb/s, QSFP, 3m
FDR10	MC2206310-005	Mellanox® active fiber cable, IB QDR/FDR10, 40Gb/s, QSFP, 5m

Speed	Cable OPN	Description
FDR10	MC2206310-010	Mellanox® active fiber cable, IB QDR/FDR10, 40Gb/s, QSFP, 10m
FDR10	MC2206310-015	Mellanox® active fiber cable, IB QDR/FDR10, 40Gb/s, QSFP, 15m
FDR10	MC2206310-020	Mellanox® active fiber cable, IB QDR/FDR10, 40Gb/s, QSFP, 20m
FDR10	MC2206310-030	Mellanox® active fiber cable, IB QDR/FDR10, 40Gb/s, QSFP, 30m
FDR10	MC2206310-050	Mellanox® active fiber cable, IB QDR/FDR10, 40Gb/s, QSFP, 50m
FDR10	MC2206310-100	Mellanox® active fiber cable, IB QDR/FDR10, 40Gb/s, QSFP, 100m
FDR10	MC2210411-SR4E	Mellanox® optical module, 40Gb/s, QSFP, MPO, 850nm, up to 300m

1.1.3 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	MC2207128-0A2	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 2.5m
FDR	MC2207130-001	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 1m
FDR	MC2207130-002	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 2m
FDR	MC2207130-00A	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 0.5m
FDR	MC2207130-0A1	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 1.5m
FDR	MC220731V-003	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 3m
FDR	MC220731V-005	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 5m
FDR	MC220731V-007	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 7m
FDR	MC220731V-010	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 10m
FDR	MC220731V-012	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 12m
FDR	MC220731V-015	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 15m

Speed	Cable OPN	Description
FDR	MC220731V-020	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 20m
FDR	MC220731V-025	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 25m
FDR	MC220731V-030	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 30m
FDR	MC220731V-040	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 40m
FDR	MC220731V-050	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 50m
FDR	MC220731V-075	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 75m
FDR	MC220731V-100	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 100m
FDR	MCP1700-F001C	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 1m, Red Pulltab
FDR	MCP1700-F001D	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 1m, Yellow Pulltab
FDR	MCP1700-F002C	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 2m, Red Pulltab
FDR	MCP1700-F002D	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 2m, Yellow Pulltab
FDR	MCP1700-F003C	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 3m, Red Pulltab
FDR	MCP1700-F003D	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 3m, Yellow Pulltab
FDR	MCP170L-F001	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 1m
FDR	MCP170L-F002	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 2m
FDR	MCP170L-F003	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 3m
FDR	MCP170L-F00A	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 0.5m
FDR	MCP170L-F01A	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 1.5m
FDR	MMA1B00-F030D	Mellanox® transceiver, FDR, QSFP+, MPO, 850nm, SR4, up to 30m, DDMI

1.1.4 Validated and Supported EDR / 100Gb/s Cables

Speed	Cable OPN	Description
EDR	MCP1600-E001	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 1m 30AWG
EDR	MCP1600-E001E30	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 1m, Black, 30AWG
EDR	MCP1600-E002	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 2m 28AWG
EDR	MCP1600-E002E30	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 2m, Black, 30AWG
EDR	MCP1600-E003	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 3m 26AWG
EDR	MCP1600-E003E26	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 3m, Black, 26AWG
EDR	MCP1600-E004E26	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 4m, Black, 26AWG
EDR	MCP1600-E005E26	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 5m, Black, 26AWG
EDR	MCP1600-E00A	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 0.5m 30AWG
EDR	MCP1600-E00AE30	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 0.5m, Black, 30AWG
EDR	MCP1600-E00BE30	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 0.75m, Black, 30AWG
EDR	MCP1600-E01A	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 1.5m 30AWG
EDR	MCP1600-E01AE30	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 1.5m, Black, 30AWG
EDR	MCP1600-E01BE30	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 1.25m, Black, 30AWG
EDR	MCP1600-E02A	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 2.5m 26AWG
EDR	MCP1600-E02AE26	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 2.5m, Black, 26AWG
EDR	MFA1A00-E001	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 1m
EDR	MFA1A00-E003	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 3m
EDR	MFA1A00-E005	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 5m
EDR	MFA1A00-E010	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 10m

Speed	Cable OPN	Description
EDR	MFA1A00-E015	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 15m
EDR	MFA1A00-E020	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 20m
EDR	MFA1A00-E030	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 30m
EDR	MFA1A00-E050	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 50m
EDR	MFA1A00-E100	Mellanox® active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 100m
EDR	MMA1B00-E100	Mellanox® transceiver, IB EDR, up to 100Gb/s, QSFP28, MPO, 850nm, SR4, up to 100m

⚠ EDR links raise with RS-FEC.

1.1.5 Validated and Supported HDR / 200Gb/s Cables

Speed	Cable OPN	Description
HDR	MCP7H50-H001R30	Mellanox® passive copper hybrid cable, IB HDR 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, colored, 1m, 30AWG
HDR	MCP7H50-H002R26	Mellanox® passive copper hybrid cable, IB HDR 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, colored, 2m, 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.

1.1.6 Validated and Supported 1GbE Cables

Speed	Cable OPN	Description
1GbE	MC3208011-SX	Mellanox® Optical module, ETH 1GbE, 1Gb/s, SFP, LC-LC, SX 850nm, up to 500m
1GbE	MC3208411-T	Mellanox® module, ETH 1GbE, 1Gb/s, SFP, Base-T, up to 100m

1.1.7 Validated and Supported 10GbE Cables

Speed	Cable OPN	Description
10GE	MFM1T02A-LR	Mellanox® SFP+ optical module for 10GBASE-LR
10GE	MFM1T02A-SR	Mellanox® SFP+ optical module for 10GBASE-SR
10GE	MAM1Q00A-QSA	Mellanox® cable module, ETH 10GbE, 40Gb/s to 10Gb/s, QSFP to SFP+
10GE	MC2309124-005	Mellanox® passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 5m
10GE	MC2309124-007	Mellanox® passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 7m
10GE	MC2309130-001	Mellanox® passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 1m
10GE	MC2309130-002	Mellanox® passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 2m
10GE	MC2309130-003	Mellanox® passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 3m
10GE	MC2309130-00A	Mellanox® passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 0.5m
10GE	MC3309124-004	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 4m
10GE	MC3309124-005	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 5m
10GE	MC3309124-006	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 6m
10GE	MC3309124-007	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 7m
10GE	MC3309130-001	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1m
10GE	MC3309130-002	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 2m
10GE	MC3309130-003	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 3m
10GE	MC3309130-00A	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 0.5m
10GE	MC3309130-0A1	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1.5m
10GE	MC3309130-0A2	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 2.5m
10GE	MCP2100-X001B	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1m, Blue Pulltab, Connector Label
10GE	MCP2100-X002B	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 2m, Blue Pulltab, Connector Label
10GE	MCP2100-X003B	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 3m, Blue Pulltab, Connector Label

Speed	Cable OPN	Description
10GE	MCP2101-X001B	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1m, Green Pulltab, Connector Label
10GE	MCP2104-X001B	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1m, Black Pulltab, Connector Label
10GE	MCP2104-X002B	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 2m, Black Pulltab, Connector Label
10GE	MCP2104-X003B	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 3m, Black Pulltab, Connector Label
10GE	MCP2104-X01AB	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1.5m, Black Pulltab, Connector Label
10GE	MCP2104-X02AB	Mellanox® passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 2.5m, Black Pulltab, Connector Label

1.1.8 Validated and Supported 25GbE Cables

The 25GbE cables can be supported only when connected to the MAM1Q00A-QSA28 module.

Speed	Cable OPN	Description
25GE	MAM1Q00A-QSA28	Mellanox® cable module, ETH 25GbE, 100Gb/s to 25Gb/s, QSFP28 to SFP28
25GE	MCP2M00-A001	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1m, 30AWG
25GE	MCP2M00-A001E30N	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1m, Black, 30AWG, CA-N
25GE	MCP2M00-A002	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2m, 30AWG
25GE	MCP2M00-A002E30N	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2m, Black, 30AWG, CA-N
25GE	MCP2M00-A003E26N	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 3m, Black, 26AWG, CA-N
25GE	MCP2M00-A003E30L	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 3m, Black, 30AWG, CA-L
25GE	MCP2M00-A004E26L	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 4m, Black, 26AWG, CA-L
25GE	MCP2M00-A005E26L	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 5m, Black, 26AWG, CA-L
25GE	MCP2M00-A00A	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 0.5m, 30AWG
25GE	MCP2M00-A00AE30N	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 0.5m, Black, 30AWG, CA-N
25GE	MCP2M00-A01AE30N	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1.5m, Black, 30AWG, CA-N

Speed	Cable OPN	Description
25GE	MCP2M00-A02AE26N	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2.5m, Black, 26AWG, CA-N
25GE	MCP2M00-A02AE30L	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2.5m, Black, 30AWG, CA-L
25GE	MFA2P10-A003	Mellanox® active optical cable 25GbE, SFP28, 3m
25GE	MFA2P10-A005	Mellanox® active optical cable 25GbE, SFP28, 5m
25GE	MFA2P10-A007	Mellanox® active optical cable 25GbE, SFP28, 7m
25GE	MFA2P10-A010	Mellanox® active optical cable 25GbE, SFP28, 10m
25GE	MFA2P10-A015	Mellanox® active optical cable 25GbE, SFP28, 15m
25GE	MFA2P10-A020	Mellanox® active optical cable 25GbE, SFP28, 20m
25GE	MFA2P10-A030	Mellanox® active optical cable 25GbE, SFP28, 30m
25GE	MFA2P10-A050	Mellanox® active optical cable 25GbE, SFP28, 50m
25GE	MMA2P00-AS	Mellanox® transceiver, 25GbE, SFP28, LC-LC, 850nm, SR, up to 100m
25GE	MMA2L20-AR	Mellanox® optical transceiver, 25GbE, 25Gb/s, SFP28, LC-LC, 1310nm, LR up to 10km

1.1.9 Validated and Supported 40GbE Cables

Speed	Cable OPN	Description
40GE	MC2206128-004	Mellanox® passive copper cable, VPI, up to 40Gb/s, QSFP, 4m
40GE	MC2206128-005	Mellanox® passive copper cable, VPI, up to 40Gb/s, QSFP, 5m
40GE	MC2206130-001	Mellanox® passive copper cable, VPI, up to 40Gb/s, QSFP, 1m
40GE	MC2206130-002	Mellanox® passive copper cable, VPI, up to 40Gb/s, QSFP, 2m
40GE	MC2206130-003	Mellanox® passive copper cable, VPI, up to 40Gb/s, QSFP, 3m
40GE	MC2206130-00A	Mellanox® passive copper cable, VPI, up to 40Gb/s, QSFP, 0.5m
40GE	MC2210126-004	Mellanox® passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 4m
40GE	MC2210126-005	Mellanox® passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 5m
40GE	MC2210128-003	Mellanox® passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 3m
40GE	MC2210130-001	Mellanox® passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 1m
40GE	MC2210130-002	Mellanox® passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 2m

Speed	Cable OPN	Description
40GE	MC2210310-003	Mellanox® active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 3m
40GE	MC2210310-005	Mellanox® active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 5m
40GE	MC2210310-010	Mellanox® active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 10m
40GE	MC2210310-015	Mellanox® active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 15m
40GE	MC2210310-020	Mellanox® active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 20m
40GE	MC2210310-030	Mellanox® active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 30m
40GE	MC2210310-050	Mellanox® active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 50m
40GE	MC2210310-100	Mellanox® active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 100m
40GE	MC2210411-SR4E	Mellanox® optical module, 40Gb/s, QSFP, MPO, 850nm, up to 300m
40GE	MC2609125-005	Mellanox® passive copper hybrid cable, ETH 40GbE to 4x10GbE, QSFP to 4xSFP+, 5m
40GE	MC2609130-001	Mellanox® passive copper hybrid cable, ETH 40GbE to 4x10GbE, QSFP to 4xSFP+, 1m
40GE	MC2609130-003	Mellanox® passive copper hybrid cable, ETH 40GbE to 4x10GbE, QSFP to 4xSFP+, 3m
40GE	MCP1700-B001E	Mellanox® passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 1m, Black Pulltab
40GE	MCP1700-B002E	Mellanox® passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 2m, Black Pulltab
40GE	MCP1700-B003E	Mellanox® passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 3m, Black Pulltab
40GE	MCP1700-B01AE	Mellanox® passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 1.5m, Black Pulltab
40GE	MCP1700-B02AE	Mellanox® passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 2.5m, Black Pulltab
40GE	MMA1B00-B150D	Mellanox® transceiver, 40GbE, QSFP+, MPO, 850nm, SR4, up to 150m, DDMI

1.1.10 Validated and Supported 56GbE Cables

Speed	Cable OPN	Description
56GE	MC2207126-004	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 4m
56GE	MC2207128-003	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 3m

Speed	Cable OPN	Description
56GE	MC2207128-0A2	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 2.5m
56GE	MC2207130-001	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 1m
56GE	MC2207130-002	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 2m
56GE	MC2207130-00A	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 0.5m
56GE	MC2207130-0A1	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 1.5m
56GE	MC220731V-003	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 3m
56GE	MC220731V-005	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 5m
56GE	MC220731V-010	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 10m
56GE	MC220731V-015	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 15m
56GE	MC220731V-020	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 20m
56GE	MC220731V-025	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 25m
56GE	MC220731V-030	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 30m
56GE	MC220731V-040	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 40m
56GE	MC220731V-050	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 50m
56GE	MC220731V-075	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 75m
56GE	MC220731V-100	Mellanox® active fiber cable, VPI, up to 56Gb/s, QSFP, 100m
56GE	MCP1700-F001C	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 1m, Red Pulltab
56GE	MCP1700-F001D	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 1m, Yellow Pulltab
56GE	MCP1700-F002C	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 2m, Red Pulltab
56GE	MCP1700-F002D	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 2m, Yellow Pulltab
56GE	MCP1700-F003C	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 3m, Red Pulltab
56GE	MCP1700-F003D	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 3m, Yellow Pulltab

Speed	Cable OPN	Description
56GE	MCP170L-F001	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 1m
56GE	MCP170L-F002	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 2m
56GE	MCP170L-F003	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 3m
56GE	MCP170L-F00A	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 0.5m
56GE	MCP170L-F01A	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 1.5m

1.1.11 Validated and Supported 100GbE Cables

Speed	Cable OPN	Description
100GE	MCP1600-C001	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP, PVC, 1m 30AWG
100GE	MCP1600-C001E30N	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 1m, Black, 30AWG, CA-N
100GE	MCP1600-C002	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP, PVC, 2m 30AWG
100GE	MCP1600-C002E30N	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 2m, Black, 30AWG, CA-N
100GE	MCP1600-C003	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP, PVC, 3m 28AWG
100GE	MCP1600-C003E26N	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 3m, Black, 26AWG, CA-N
100GE	MCP1600-C003E30L	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 3m, Black, 30AWG, CA-L
100GE	MCP1600-C005E26L	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 5m, Black, 26AWG, CA-L
100GE	MCP1600-C00A	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP, PVC, 0.5m 30AWG
100GE	MCP1600-C00AE30N	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 0.5m, Black, 30AWG, CA-N
100GE	MCP1600-C00BE30N	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 0.75m, Black, 30AWG, CA-N
100GE	MCP1600-C01A	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP, PVC, 1.5m 30AWG
100GE	MCP1600-C01AE30N	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 1.5m, Black, 30AWG, CA-N
100GE	MCP1600-C02A	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP, PVC, 2.5m 30AWG
100GE	MCP1600-C02AE26N	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 2.5m, Black, 26AWG, CA-N

Speed	Cable OPN	Description
100GE	MCP1600-C02AE30L	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28,2.5m, Black, 30AWG, CA-L
100GE	MCP1600-C03A	Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP, PVC, 3.5m 26AWG
100GE	MCP1600-E001	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 1m 30AWG
100GE	MCP1600-E002	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 2m 28AWG
100GE	MCP1600-E003	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 3m 26AWG
100GE	MCP1600-E01A	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 1.5m 30AWG
100GE	MCP1600-E02A	Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 2.5m 26AWG
100GE	MCP7F00-A001R	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, colored pulltabs, 1m, 30AWG
100GE	MCP7F00-A001R30N	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 1m, Colored, 30AWG, CA-N
100GE	MCP7F00-A002R	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, colored pulltabs, 2m, 30AWG
100GE	MCP7F00-A002R30N	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 2m, Colored, 30AWG, CA-N
100GE	MCP7F00-A003R26N	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 3m, Colored, 26AWG, CA-N
100GE	MCP7F00-A003R30L	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 3m, Colored, 30AWG, CA-L
100GE	MCP7F00-A005R26L	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 5m, Colored, 26AWG, CA-L
100GE	MCP7F00-A01AR	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, colored pulltabs,1.5m, 30AWG
100GE	MCP7F00-A01AR30N	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 1.5m, Colored, 30AWG, CA-N
100GE	MCP7F00-A02AR26N	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 2.5m, Colored, 26AWG, CA-N
100GE	MCP7F00-A02AR30L	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 2.5m, Colored, 30AWG, CA-L

Speed	Cable OPN	Description
100GE	MCP7F00-A02ARLZ	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 2.5m, LSZH, Colored, 28AWG
100GE	MCP7F00-A03AR26L	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 3.5m, Colored, 26AWG, CA-L
100GE	MCP7H00-G001	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 1m, 30AWG
100GE	MCP7H00-G001R	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, colored pulltabs, 1m, 30AWG
100GE	MCP7H00-G001R30N	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 1m, Colored, 30AWG, CA-N
100GE	MCP7H00-G002R	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, colored pulltabs, 2m, 30AWG
100GE	MCP7H00-G002R30N	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 2m, Colored, 30AWG, CA-N
100GE	MCP7H00-G003R	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, colored pulltabs, 3m, 28AWG
100GE	MCP7H00-G003R26N	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 3m, Colored, 26AWG, CA-N
100GE	MCP7H00-G003R30L	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 3m, Colored, 30AWG, CA-L
100GE	MCP7H00-G004R26L	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 4m, Colored, 26AWG, CA-L
100GE	MCP7H00-G01AR	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, colored pulltabs, 1.5m, 30AWG
100GE	MCP7H00-G01AR30N	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 1.5m, Colored, 30AWG, CA-N
100GE	MCP7H00-G02AR	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, colored pulltabs, 2.5m, 30AWG
100GE	MCP7H00-G02AR26N	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 2.5m, Colored, 26AWG, CA-N
100GE	MCP7H00-G02AR30L	Mellanox® passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 2.5m, Colored, 30AWG, CA-L
100GE	MFA1A00-C003	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 3m

Speed	Cable OPN	Description
100GE	MFA1A00-C005	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 5m
100GE	MFA1A00-C010	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 10m
100GE	MFA1A00-C015	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 15m
100GE	MFA1A00-C020	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 20m
100GE	MFA1A00-C030	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 30m
100GE	MFA1A00-C050	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 50m
100GE	MFA1A00-C100	Mellanox® active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 100m
100GE	MFA7A20-C003	Mellanox® active fiber hybrid solution, ETH 100GbE to 2x50GbE, QSFP28 to 2xQSFP28, 3m
100GE	MFA7A20-C005	Mellanox® active fiber hybrid solution, ETH 100GbE to 2x50GbE, QSFP28 to 2xQSFP28, 5m
100GE	MFA7A20-C010	Mellanox® active fiber hybrid solution, ETH 100GbE to 2x50GbE, QSFP28 to 2xQSFP28, 10m
100GE	MFA7A20-C020	Mellanox® active fiber hybrid solution, ETH 100GbE to 2x50GbE, QSFP28 to 2xQSFP28, 20m
100GE	MFA7A50-C003	Mellanox® active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 3m
100GE	MFA7A50-C005	Mellanox® active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 5m
100GE	MFA7A50-C010	Mellanox® active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 10m
100GE	MFA7A50-C015	Mellanox® active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 15m
100GE	MFA7A50-C020	Mellanox® active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 20m
100GE	MFA7A50-C030	Mellanox® active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 30m
100GE	MMA1B00-C100D	Mellanox® transceiver, 100GbE, QSFP28, MPO, 850nm, SR4, up to 100m, DDMI
100GbE	MMA1L10-CR	Mellanox® optical transceiver, 100GbE, QSFP28, LC-LC, 1310nm, LR4 up to 10km **Note:** Only revision A2 and above.

1.1.12 Validated and Supported 200GbE Cables

Speed	Cable OPN	Description
200GE	MCP1650-V001E30	Mellanox® Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 1m, black pulltab, 30AWG
200GE	MCP1650-V002E26	Mellanox® Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 2m, black pulltab, 26AWG
200GE	MCP1650-V003E26	Mellanox® Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 3m, black pulltab, 26AWG
200GE	MCP1650-V00AE30	Mellanox® Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 0.5m, black pulltab, 30AWG
200GE	MCP1650-V01AE30	Mellanox® Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 1.5m, black pulltab, 30AWG
200GE	MCP1650-V02AE26	Mellanox® Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 2.5m, black pulltab, 26AWG

1.2 Supported 3rd Party Cables and Modules

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	FTLX8571D3BCL-ME	10gb SFP 850nm Optic Transceiver
10GbE	L45593-D178-B50	QSFP-4SFP10G-CU5M
10GbE	SFP-10G-SR	Cisco 10GBASE-SR SFP+ transceiver module for MMF, 850-nm wavelength, LC duplex connector
10GbE	SFP-H10GB-ACU10M	Cisco 10GBASE-CR1 Active Copper Cable 10-meter
10GbE	SFP-H10GB-ACU7M	Cisco 10GBASE-CR1 Active Copper Cable 7-meter
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-CU4M	Cisco 10GBASE-CR1 Copper Cable 4-meter
10GbE	SFP-H10GB-CU5M	Cisco 5-m 10G SFP+ Twinax cable assembly, passive
25GbE	SFP-25G-AOC5M	Cisco 25GBASE-AOC Active Optical Cable 5-meter
25GbE	SFP-25G-AOC7M	Cisco 25GBASE-AOC Active Optical Cable 7-meter

Speed	Cable OPN	Description
25GbE	SFP-H25G-CU1M	25GBASE-CR1 Copper Cable 1-meter
25GbE	SFP-H25G-CU2.5M	Cisco 25GBASE-CR1 Copper Cable 2.5-meter
25GbE	SFP-H25G-CU2M	25GBASE-CR1 Copper Cable 2-meter
25GbE	SFP-H25G-CU3M	Cisco 25GBASE-CR1 Copper Cable 3-meter
25GbE	SFP-H25G-CU4M	Cisco 25GBASE-CR1 Copper Cable 4-meter
40GbE	2231254-2	PASSIVE COPPER CABLE ETH 40GBE QSFP 3M
40GbE	QSFP-40G-SR4	Cisco 40GBASE-SR4, 4 lanes, 850 nm MMF
40GbE	QSFP-40G-SR-BD	Cisco 40GBASE-SR-BiDi, duplex MMF
40GbE	QSFP-4SFP10G-CU5M	PASSIVE COPPER SPLITTER CABLE ETH 40GBE TO 4X10GBE 5M
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
40GbE	QSFP-H40G-CU5M	PASSIVE COPPER CABLE ETH 40GBE QSFP 5M
56GbE	FTL414QB2N-E5	Finisar FTL414QB2N-E5 56Gb 850nm 100m QSFP+ Transceiver Module ARK
100GbE	CAB-Q-Q-100GbE-3M	Passive 3 meter , QSFP+ to QSFP+ QSFP100 TWINAX 103.125Gbps-CR4
100GbE	FTLF8519P3BTL-N1	1000BASE-SX and 2G Fibre Channel (2GFC) 500m Industrial Temperature SFP Optical Transceiver
100GbE	QSFP-100G-AOC5M	Cisco 100GBASE QSFP Active Optical Cables 5-meter
100GbE	QSFP-100G-AOC7M	Cisco 100GBASE QSFP Active Optical Cables 7-meter
100GbE	QSFP-100G-CU3M	Cisco 100GBASE-CR4 Passive Copper Cable 3-meter
100GbE	QSFP-100G-CU5M	Cisco 100GBASE-CR4 Passive Copper Cable 5-meter
100GbE	QSFP-100G-SR4-S	Cisco 100GBASE SR4 QSFP Transceiver, MPO, 100m over OM4 MMF
100GbE	QSFP-40/100-SRBD	Cisco 100G and 40GBASE SR-BiDi QSFP Transceiver, LC, 100m OM4 MMF
100GbE	SO-QSFP28-LR4	QSFP28, 100GBase, 1310nm, SM, DDM, 10km, LC
100GbE	TR-FC13L-N00	100G QSFP28 Optical Transceivers, QSFP28 LR4 (10km)

1.3 Tested Switches

1.3.1 Tested EDR / 100Gb/s Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
EDR	Switch-IB	MSB7790-XXX	36-port Unmanaged EDR 100Gb/s InfiniBand Switch Systems	NVIDIA
EDR	Switch-IB	MSB7700-XXX	36-port Managed EDR 100Gb/s InfiniBand Switch Systems	NVIDIA
EDR	Switch-IB 2	MSB7800-XXX	36-port Managed EDR 100Gb/s InfiniBand Switch Systems	NVIDIA

1.3.2 Tested 10/40GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
10GbE	N/A	5548UP	32x 10GbE SFP+ Switch System	Cisco
10/40GbE	N/A	7050Q	16 x 40GbE QSFP+ Switch System	Arista
10/40GbE	N/A	7050S	48x 10GbE SFP+ and 4 x 40GbE QSFP+ Switch System	Arista
10/40GbE	N/A	G8264	48x 10GbE SFP+ and 4 x 40GbE QSFP+ Switch System	Lenovo
10/40GbE	N/A	QFX3500	48x 10GbE SFP+ and 4 x 40GbE QSFP+ Switch System	Juniper
10/40GbE	N/A	S4810P-AC	48x 10GbE SFP+ and 4 x 40GbE QSFP+ Switch System	Force10
10/40GbE	N/A	3064	48x 10GbE SFP+ and 4 x 40GbE QSFP+ Switch System	Cisco
10/40GbE	N/A	8164F	48x 10GbE SFP+ and 2 x 40GbE QSFP+ Switch System	Dell
10/40GbE	N/A	S5000	48x 10GbE SFP+ and 4 x 40GbE QSFP+ Switch System	Dell
10/40GbE	N/A	3132Q	4x 10GbE SFP+ and 32 x 40GbE QSFP+ Switch System	Cisco
40GbE	N/A	7050QX	32x 40GbE QSFP+ Switch System	Arista
40GbE	N/A	G8316	16x 40GbE QSFP+ Switch System	Lenovo
40GbE	N/A	S6000	32x 40GbE QSFP+ Switch System	Dell

1.3.3 Tested 100GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
100GbE	Spectrum	MSN2410-XXXX	48-port 25GbE + 8-port 100GbE Open Ethernet Switch System	Mellanox
100GbE	Spectrum	MSN2700-XXXX	32-port Non-blocking 100GbE Open Ethernet Switch System	Mellanox
100GbE	Spectrum	MSN2740-XXXX	32-port Non-blocking 100GbE Open Ethernet Switch System	Mellanox
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
100GbE	N/A	CE6860-1-48S8CQ-EI	Huawei 100GbE Ethernet switch	Huawei
100GbE	N/A	7060CX-32S	32-port 100GbE Ethernet Switch System	Arista
100GbE	N/A	3232C	32-port 100GbE Ethernet Switch System	Cisco
100GbE	N/A	N9K-C9236C	36-port 100GbE Ethernet Switch System	Cisco
100GbE	N/A	93180YC-EX	48-port 25GbE + 6-port 100GbE Ethernet Switch System	Cisco
100GbE	N/A	T7032-IX7	32-port 100GbE Ethernet Switch System	Quanta

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 this firmware version:

	Supported Version
MLNX_OFED	5.0-2.1.8.0/ 5.0-1.0.0.0 / 4.7-3.2.9.0
MLNX_EN (MLNX_OFED based code)	5.0-2.1.8.0 / 5.0-1.0.0.0 / 4.7-3.2.9.0
WinOF-2	2.40.50000 / 2.30 / 2.20
MFT	4.14.0-105 / 4.13.3 / 4.13.0
MLNX-OS	3.8.2204
Onyx	3.8.2306
ConnectX-5 Firmware	16.27.2008 / 16.27.1016 / 16.26.1040
SwitchX-IB™ Firmware	11.2000.2626
SwitchX-IB 2 Firmware	15.2000.2626
Linux Inbox Drivers	• RH7.6 • Ubuntu 16.04.05
Windows Inbox Drivers	• Windows 2012 • Windows 2012 R2 • Windows 2016

1.5 Supported FlexBoot, UEFI

> ⚠ Please be aware that not all firmware binaries contain FlexBoot or UEFI, support may vary between cards. For further information see Supported Devices.

This firmware version is compiled with the following expansion ROMs and versions:

Expansion ROM	Supported Version
FlexBoot	3.5.901
UEFI	14.20.22

1.6 PRM Revision Compatibility

This firmware version 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

2.1 Important Notes

- Security Hardening Enhancements: This release contains important reliability improvements and security hardening enhancements. Mellanox recommends upgrading your devices firmware to this release to improve the devices' firmware security and reliability.
- When upgrading or changing configuration on multi-host adapter cards, for the changes to take effect, PCIe restart must be simultaneously send from both hosts (servers).

  To do so, perform the following:

  1. Shut down the server with the auxiliary card.
  2. Shut down the server with the primary card.
  3. Bring back the server with the primary card.
  4. Bring back the server with the auxiliary card.
- SR-IOV - Virtual Functions (VF) per Port - The maximum Virtual Functions (VF) per port is 127. For further information, see RoCE Limitations.

2.2 Changes and New Feature in this Firmware Version

Feature/Change	Description
16.27.2008	
Parallel VF init/Teardown Performance Optimization	Improved init_hca performance in Parallel Function initialization.
Bug Fixes	See Bug Fixes.

2.3 Unsupported Features and Commands

2.3.1 Unsupported Features

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

- The following service types:
  - SyncUMR
  - Mellanox transport

    - RAW IPv6
- INT-A not supported for EQs only MSI-X
- PCI VPD write flow (RO flow supported)
- Streaming Receive Queue (STRQ) and collapsed CQ
- Subnet Manager (SM) on VFs
- RoCE LAG in Multi-Host/Socket-Direct

2.3.2 Unsupported Commands

- QUERY_MAD_DEMUX
- SET_MAD_DEMUX
- CREATE_RQ - MEMORY_RQ_RMP
- MODIFY_LAG_ASYNC_EVENT

3 Bug Fixes in this Firmware Version

Bug Fixes History lists the bugs fixed in this release. For a list of old Bug Fixes, please see Bug Fixes History.

Internal Ref.	Issue
1848091	**Description:** Although the effective BER (after FEC) is expected to meet our design targets (e.g. 10e-14 or lower), occasionally it may be higher.
	Keywords: Cables
	Discovered in Version: 20.25.6000
	Fixed in Release: 16.27.2008
1856717	**Description:** High BER may occur when connecting cables of type 0.5/1m DAC to an HDR speed.
	Keywords: Cables
	Discovered in Version: 20.25.7020
	Fixed in Release: 16.27.2008
2083691	**Description:** Fixed an issue that prevented the load of the correct PCIe Tx parameters when the speed was changed after the PCIe link was disabled.
	Keywords: PCIe
	Discovered in Version: 16.27.1016
	Fixed in Release: 16.27.2008
2107103	**Description:** Fixed an issue that prevented the desched_threshold field from working properly.
	Keywords: DCQCN
	Discovered in Version: 16.27.1016
	Fixed in Release: 16.27.2008
2101810	**Description:** Fixed an issue that caused the "roce_adp_retrans" counter to present the values of the "local_ack_timeout_err" counter.
	Keywords: RoCE, lossy, q_counter
	Discovered in Version: 16.27.1016
	Fixed in Release: 16.27.2008
2063264	**Description:** If Relaxed Ordering is disabled by running the "setpci" command, it will not be functional even after re-enabling it by running the "setpci" command again.
	Keywords: PCI Relaxed Ordering
	Discovered in Version: 16.27.1016
	Fixed in Release: 16.27.2008

4 Known Issues

> This section includes history of known issues of 3 major releases back. For older releases history, please refer to the relevant firmware versions.

Ethernet Rate Limit per VF in RoCE Mode Limitations

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
127	127	64	64	127	127

Ethernet Rate Limit per VF in InfiniBand Mode Limitations

Dual Port Device		Single Port Device	
w/o LAG		w/o LAG	
w/o QoS	Full QoS	w/o QoS	Full QoS
127	127	127	127

Known Issues

Internal Ref.	Issue
2064538	**Description:** When working with an NVME offload QP that is created with a unaligned page size (page_offset != 0), the QP moves to an error state on the first posted WQE.
	Workaround: Create an NVME offload QP with page an aligned size (page_offset = 0).
	Keywords: NVMF offload, unaligned page size
	Discovered in Version: 16.27.2008
2093458	**Description:** Flow Metering capability is not functional in firmware v16.27.1016.
	Workaround: To use Flow Metering, use older firmware versions.
	Keywords: Flow Metering
	Discovered in Version: 16.27.1016
2071210	**Description:** mlxconfig query for the BOOT_INTERRUPT_DIS TLV shows a wrong value in the "current value" field.
	Workaround: Use "next boot" indication to see the right value.
	Keywords: mlxconfig
	Discovered in Version: 16.27.1016

Internal Ref.	Issue
2058677	**Description:** In Socket Direct supported cards, after performing mlxfwreset, the expansion ROM register might be writable on all hosts for less than 1 second.
	Workaround: N/A
	Keywords: Expansion ROM, Socket Direct
	Discovered in Version: 16.27.1016
2057653	**Description:** quota_exceeded_command and invalid_command counters do not function properly. In this firmware version, the quota_exceeded_command counter's value always remains 0, whereas the invalid_command counter increases only for some Ethernet commands failure events.
	Workaround: N/A
	Keywords: quota_exceeded_command, invalid_command, vnic_env counters
	Discovered in Version: 16.27.1016
1930619	**Description:** PF_BAR2 and ATS cannot be enabled together, i.e. when PF_BAR2 is enabled, ATS cannot be enabled too.
	Workaround: N/A
	Keywords: ATS, SF, BAR2, Multi GVMI
	Discovered in Version: 16.26.1040
-	**Description:** In rare cases, following a server powerup, a fatal error (device's health compromised) message might appear with ext_synd 0x8d1d. The error will be accompanied by a failure to use mlxconfig and in some cases flash burning tools.
	Workaround: N/A
	Keywords: mlxconfig, flash tool, ext_synd 0x8d1d
	Discovered in Version: 16.26.1040
1888306	**Description:** Occasionally, a BlueScreen might occur when using mlxfwreset for Socket Direct devices on Windows.
	Workaround: N/A
	Keywords: mlxfwreset, Socket Direct
	Discovered in Version: 16.26.1040
1919403	**Description:** Hardware arbitration is currently disabled in OCP3.0 cards. It will be supported on future releases for the same hardware.
	Workaround: N/A
	Keywords: Hardware arbitration, OCP3.0
	Discovered in Version: 16.26.1040
1836465	**Description:** 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. The features that may be affected by this and not work properly are: • Host chaining • Mirroring in FDB • TTL modify in FDB • VGT+
	Workaround: N/A

Internal Ref.	Issue
	Keywords: E-switch vport context, VLAN
	Discovered in Version: 16.26.1040
1842278	**Description:** DC LAG can function only in case there is a single PF per port without any active VFs.
	Workaround: N/A
	Keywords: DC LAG
	Discovered in Version: 16.26.1040
1796628	**Description:** Due to performance considerations, unicast loopback traffic will go through the NIC SX tables, and multicast loopback traffic will skip the NIC SX tables.
	Workaround: N/A
	Keywords: Performance, unicast loopback traffic, multicast loopback traffic
	Discovered in Version: 16.26.1040
1797493	**Description:** Firmware asserts may occur when setting the PF_BAR2_SIZE value higher than the maximum supported size (maximum PF_BAR2_SIZE is 4 for .
	Workaround: Configure within limits (NIC PF_BAR_SIZE <= 4).
	Keywords: Multi-GVMI, Sub-Function, SFs, BAR2
	Discovered in Version: 16.26.1040
1761271	**Description:** CWDM4 AOM cable is currently not supported.
	Workaround: N/A
	Keywords: Modules/Cables
	Discovered in Version: 16.26.1040
1762142	**Description:** 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.
	Workaround: Perform a graceful shutdown, and not an FLR.
	Keywords: Multi-GVMI, SF, Sub-Functions, FLR
	Discovered in Version: 16.25.1020
1768814/1772474	**Description:** Due to hardware limitation, REG_C cannot be passed over loopback when the FDB action is forwarded to multiple destinations.
	Workaround: N/A
	Keywords: Connection-Tracking
	Discovered in Version: 16.25.1020
1770736	**Description:** When a PF or ECPF with many VFs (SR-IOV), and/or SFs (Multi-GVMI) triggers an FLR, PCIe completion timeout might occur.
	Workaround: Increase the PCIe completion timeout.
	Keywords: Multi-GVMI, SR-IOV, Sub-Function, Virtual Function, PF FLR
	Discovered in Version: 16.25.1020
1716334	**Description:** When mlxconfig.PF_BAR2_EN is enabled, configuring more than 255 PCI functions will raise an assert.

Internal Ref.	Issue
	Workaround: When working with BAR2, configure SR-IOV to align to the 255 PCI functions limitation. mlxconfig.NUM_OF_VFS controls the number of configured SR-IOV VFs. e.g.: • Smart NICs: 2 External Host PFs, 2 ARM ECPFs, 125 VFs per PF. • Non-smart NICs: 2 External Host PFs, 126 VFs per PF
	Keywords: Multi-GVMI, PF_BAR2_EN, Sub-Functions, SR-IOV, VFs
	Discovered in Version: 16.25.1020
1699214	**Description:** NODNIC VF is partially tested. It is fully tested only in ConnectX-5 adapter cards.
	Workaround: N/A
	Keywords: NODNIC VF
	Discovered in Version: 16.25.1020
1749691	**Description:** On rare occasions, when using Socket-Direct devices, inband burning through the external port might fail.
	Workaround: N/A
	Keywords: Socket-Direct, inband burning
	Discovered in Version: 16.25.1020
1689186	**Description:** Changing priority to TC map during traffic might cause packet drops.
	Workaround: N/A
	Keywords: QoS
	Discovered in Version: 16.25.1020
1604699	**Description:** 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.
	Workaround: N/A
	Keywords: Counters
	Discovered in Version: 16.25.1020
1558250	**Description:** eSwitch owner may receive NIC_VPORT_CONTEXT events from vPorts that are not necessarily armed using the nic vport context arm_change_even tbit.
	Workaround: N/A
	Keywords: Port event, NODNIC
-	**Description:** In Ethernet mode, at 10/40GbE speeds, only NO-FEC in Force mode is supported. Other user configurations are overridden.
	Workaround: N/A
	Keywords: Ethernet, 10GbE, 40GbE, RS-FEC
	Discovered in Version: 16.25.1020
1574876	**Description:** DC RoCE LAG is functional only if the router posts VRRP address as the source MAC.
	Workaround: N/A
	Keywords: DC RoCE LAG
	Discovered in Version: 16.25.1020

Internal Ref.	Issue
1498399	**Description:** 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).
	Workaround: N/A
	Keywords: XRC SRQ/RMP ODP
	Discovered in Version: 16.25.1020
1546401	**Description:** vport_tc and para_vport_tc are not supported in this version.
	Workaround: N/A
	Keywords: SR-IOV vport_tc and para_vport_tc
	Discovered in Version: 16.24.1000
1546492	**Description:** Executing the update_lid command while the IB port sniffer utility is active can stop the utility.
	Workaround: N/A
	Keywords: IB Sniffer
	Discovered in Version: 16.24.1000
1537898	**Description:** Initializing a function while the IB port sniffer utility is active can stop the utility.
	Workaround: N/A
	Keywords: IB Sniffer
	Discovered in Version: 16.24.1000
1523577	**Description:** 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.
	Workaround: N/A
	Keywords: multi_prio_sq, VF
	Discovered in Version: 16.24.1000
1414290	**Description:** When getting an inline scatter CQE on IB striding RQ, the stride index in the CQE will be zero.
	Workaround: N/A
	Keywords: Scatter CQE
	Discovered in Version: 16.24.1000
1475490	**Description:** Reboot is not supported on any host during the PLDM firmware burning process.
	Workaround: N/A
	Keywords: PLDM
	Discovered in Version: 16.23.1020
1332714/1345824	**Description:** The maximum “read” size of MTRC_STDB is limited to 272 Bytes.
	Workaround: Set the MTRC_STDB.read_size to the maximum value of 0x110=272 Bytes
	Keywords: Access register, MTRC_STDB, tracer to dmesg, fwtrace to dmesg

Internal Ref.	Issue
	Discovered in Version: 16.23.1020
1408994	**Description:** FTE with both forward (FWD) and encapsulation (ENCAP) actions is not supported in the SX NIC Flow Table.
	Workaround: N/A
	Keywords: SX NIC Flow Table
	Discovered in Version: 16.23.1020
1350794	**Description:** Encapsulation / Decapsulation support in steering has the following limitations: • Encapsulation / Decapsulation can be open on the FDB only if all VFs are non active. • Encapsulation / Decapsulation supports single mode only: FDB / NIC. Opening tables of both types is not supported. • Encapsulation / Decapsulation per device support: 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
	Workaround: N/A
	Keywords: Steering Encapsulation / Decapsulation
	Discovered in Version: 16.23.1020
1027553	**Description:** While using e-switch vport sVLAN stripping, the RX steering values on the sVLAN might not be accurate.
	Workaround: N/A
	Keywords: e-sw vport sVLAN stripping, RX steering
	Discovered in Version: 16.24.1000
1799917	**Description:** Untagged CVLAN packets in the Steering Flow Tables do not match the SVLAN tagged packets.
	Workaround: N/A
	Keywords: Steering Flow Tables, CVLAN/SVLAN packets
	Discovered in Version: 16.23.1020
1504073	**Description:** 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.
	Workaround: 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.
	Keywords: Flow Control, LRO
	Discovered in Version: 16.23.1020

Internal Ref.	Issue

1178792	**Description**: Host Chaining Limitations:
	• Single MAC address per port is supported
	• Both ports should be configured to Ethernet when host chaining is enabled
	• The following capabilities cannot function when host chaining is enabled:
	  • SR-IOV
	  • DSCP
	  • NODNIC
	  • Load balancing
	  • LAG
	  • Dual Port RoCE (multi port vHCA)

	Workaround: N/A
	Keywords: Host Chaining
	Discovered in Version: 16.22.1002
1277762	**Description:** An Ethernet multicast loopback packet is not counted (even if it is not a local loopback packet) when running the nic_receive_steering_discard command.
	Workaround: N/A
	Keywords: Ethernet multicast loopback packet
	Discovered in Version: 16.22.1002
1190753	**Description:** 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.
	Workaround: N/A
	Keywords: RoCE, vport FDB
	Discovered in Version: 16.22.1002
1306342	**Description:** Signature-accessing WQEs sent locally to the NVMeF target QPs that encounter signature errors, will not send a SIGERR CQE.
	Workaround: N/A
	Keywords: Signature-accessing WQEs, NVMeF target
	Discovered in Version: 16.22.1002
1168594	**Description:** RoCE Dual Port Mode (a.k.a Multi-Port vHCA: MPV) is not supported in Multi-Host setups.
	Workaround: N/A
	Keywords: Multi-Port vHCA, Multi-Host
	Discovered in Version: 16.21.1000
1072337	**Description:** 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.
	Workaround: N/A
	Keywords: SX sniffer Flow Table
	Discovered in Version: 16.21.1000
1171013	**Description:** Signature Handover Operations is not supported when FPP (Function-Per-Port) mode is disabled.
	Workaround: N/A

Internal Ref.	Issue
	Keywords: Signature Handover Operations, FPP
	Discovered in Version: 16.21.1000
1059975	**Description:** NVMeF limitation: • Transaction size - up to 128KB per IO (non-inline) • Support up to 16K connections • Support single namespace per drive • Staging buffer size must be at least 16MB in order to allow SRQ size of 64 entries
	Workaround: N/A
	Keywords: NVMeF
	Discovered in Version: 16.22.1010

5 PreBoot Drivers (FlexBoot/UEFI)

5.1 FlexBoot Changes and New Features

For further information, please refer to the FlexBoot Release Notes.

5.2 UEFI Changes and Major New Features

For further information, please refer to the UEFI Release Notes.

6 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_DESCRIPTORS		
	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_DESCRIPTORS		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
NV_ROCE_CC	ROCE_CC_PRIO_MASK		0x107
	ROCE_CC_ALGORITHM		

Configuration	mlxconfig Parameter Name	Class	TLV ID
NV_ROCE_CC_ECN	CLAMP_TGT_RATE_AFTER_TIME_INC		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		0x10a
	LLDP_NB_RX_MODE		
	LLDP_NB_TX_MODE		
NV_LLDP_NB_DCBX	DCBX_IEEE		0x18e
	DCBX_CEE		
	DCBX_WILLING		
NV_KEEP_LINK_UP	KEEP_ETH_LINK_UP		0x190
	KEEP_IB_LINK_UP		
	KEEP_LINK_UP_ON_BOOT		
	KEEP_LINK_UP_ON_STANDBY		
NV_QOS_CONF	NUM_OF_VL		0x192
	NUM_OF_TC		
	NUM_OF_PFC		
NV_MPFS_CONF	DUP_MAC_ACTION		0x196
	SRIOV_IB_ROUTING_MODE		

Configuration	mlxconfig Parameter Name	Class	TLV ID
	IB_ROUTING_MODE		
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		

7 Changes and New Feature History

⚠ This section includes history of changes and new feature of 3 major releases back. For older releases history, please refer to the relevant firmware versions.

Feature/Change	Description
16.27.1016	
RoCE Selective Repeat	RoCE Selective Repeat introduces a new QP retransmission mode in RoCE: recovery from packet drop by resending the dropped packet and not only all the PSN window (Go-Back-N protocol), This new capability comes with the following limitations: • Selective repeat cannot be used with AR • Does not work with signature (T10-DIF) • Does not work with Tag Matching enabled
RedFish (RDE)	Allows BMC to query and control NIC over RedFish API (https://www.dmtf.org/standards/redfish). Currently, the NIC supports reading data and setting basic Ethernet and InfiniBand parameters.
Power Consumption	Removed power consumption limitation. This new capability removes the necessity to disable the port due to insufficient power. When this feature is enabled, an event will be sent on every transition between the following power states: • Power limit is not advertised • Sufficient power detected • Insufficient power detected
ECMP with RoCE Traffic	Enables matching of source_vhca_port in the FDB flow for ECMP hardware offload on a single FDB.
Live Firmware Patch (LFWP)	Firmware can be patched with critical bugs fixes live with minimal serviceability impact. The patching can be down only within the same major branch.
Hardware Offloaded Rules (Resource Dump)	Added support for dumping hardware steering entries (raw data) using the resource dump API.
VF Command Failure Resource Dump Event	Enables the device to generate events that will notify the user about any resource dumps updates. The resource dump owner (PF/ECPF) will be notified on the VF command failures by events that will contain the needed information to call the resource_dump command. The device will dump the resource in the standard resource dump sequence.
GTP-U Tunneled Packets	Added support for RSS based on inner headers of GTP-U tunneled packets.
Link Down Counter	The `eth_link_down_counter` now counts logical link downs as well.
Relaxed Ordering Read	**[Beta]** Added support for relaxed ordering read when using the create_mkey flag. **Note:** This capability is not supported when using Multi-Host cards and when in PCIe-Switch mode.
Bug Fixes	See Bug Fixes.

Rev. 16.26.4012	
Globally Disable RoCE through MST	Enables the user to globally disable RoCE on init by writing to the access register NCFG_REG.
Zero-Touch RoCE (ZTR) Slow Start	Enabled Zero-Touch RoCE (ZTR) slow start capability for responder flows.
Resource Dump	Extracts and prints data segments generated by the firmware.
Lossless Hairpin QP	Hairpin QP buffer is now available in Flow Control. Host Chaining now supports full fairness between several devices.
Bug Fixes	See Bug Fixes History.
Rev. 16.26.2002	
Bug Fixes	This version included a bug fix for OPN MCX562A-ACA. For further information see Bug Fixes History.
Rev. 16.26.1040	
Address Translation Service (ATS)	Added Address Translation Service (ATS) support for MKEY and UMEM.
VPD	Added support for exposing the VPD on the VF.
ICMD and Diagnostic Counters	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.
Hairpin Drop Counter	Added support for Hairpin Drop Counter.
User Context Object (DEVX)	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. The following functionalities are still managed by the Kernel: • Resource cleaning • UCTX stamping • Blocking the physical address and IRQ from these UCTX
DEVX Support for Asynchronous Events	Added support for reporting the supported affiliated and unaffiliated asynchronous events to DEVX users through the command interface.
Hairpin and TM RNDV QPs in DEVX	Added support for Hairpin and TM RNDV QPs to work with DevX.
Software Managed Steering Tables	Added support for creating software managed steering tables in eSwitch/FDB.
Zero-Touch-RoCE Counters	Zero-Touch-RoCE counters are now available to the user for debuggability purposes when using the Zero-Touch-RoCE feature.
Security Hardening Enhancements	This release contains important reliability improvements and security hardening enhancements. Mellanox recommends upgrading your device firmware to this release to improve the device firmware security and reliability.
Bug Fixes	See Bug Fixes History.

Rev. 16.25.1020	
OOB Support AR in IB Networks	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).
ConnectX-5 OCP 3.0	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.
ASAP2 Offloading VXLAN Decapsulation with HW LRO	Added support for performing hardware Large Receive Offload (HW LRO) on VFs with HW-decapsulated VXLAN. For further information on the VXLAN decapsulation feature, please refer to ASAP2 User Manual under www.mellanox.com -> Products -> Software -> ASAP2.
Multiple Output Ports - Enhance ASAP2	Added the ability to send to multiple destinations encapsulated data and each destination can have its own encapsulation data.
VSC Security	VSC security includes the mechanisms which will prevent a reasonable host from affecting other hosts from using VSC.
Extended Responder Not Ready (RNR)	Enabled Responder Not Ready (RNR) configuration. Now the number of RNRs can be set higher than 6 and lower than indefinite.
Zero Touch RoCE	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.
ODP support for SRQ & XRC	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. Note: 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.
TTL RX	Enables the device to modify incoming packets' TTL from the uplink to a vport using eswitch rules when sw_fdb_ipv4_ttl_modify_enable is set to true.
QP Counters and Firmware Errors per PID	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.
Firmware Burning using DMA Pages	This new capability accelerates the firmware burning process by using Direct Memory Access (DMA) pages.
Auto-Sensing when using 25/10GbE Optical Modules	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.
Package ID	Enabled Package ID configuration using server strap according OCP 3.0.
DPDK UIO	This capability provides a solution for improving user space drivers development, generic user space IO device services.
Management Query Information Strings (MQIS)	Added MQIS support, “Part Number” and “Description" information for secured adapter cards.

SR-IOV Offloading	Added the ability to probe one representor each time for a specific representor ID when bonding is enabled.

Virtio Supported Hardware Offloads (DPDK)	Added the following capabilities related to Virtio Supported Hardware Offloads (DPDK):
	• VXLAN encapsulation/decapsulation
	• VXLAN GPE encapsulation/decapsulation
	• Packet Modification:
	  • NAT (IPv4, IPv6, TCP, UDP)
	  • Routing (L2)
	  • TTL encapsulation/decapsulation
	  • TCP SEQ and ACK
	• ICMP offloading: match on ICMP type, code and identifier fields
	• HA and LB: LAG
	• Flow Aging
	• Flow Metering
	• Firmware fast update

SR-IOV Supported Hardware Offloads	Added the following capabilities related to SR-IOV Supported Hardware Offloads:
	• VXLAN encapsulation/decapsulation
	• VXLAN GPE IPv4 & IPv6 encapsulation/decapsulation
	• Packet Modification:
	  • NAT (IPv4, IPv6, TCP, UDP)
	  • Routing (L2)
	  • TTL encapsulation/decapsulation
	  • TCP SEQ and ACK
	• VF group rate limit support
	• VF and PF LAG

Connection Tracking	Added metadata registry C “header modify” and “match rule” support to be used by the connection tracking functionality.
NODNIC Connectivity	Enabled NOIDNIC connectivity for VFs.
VDP	Enabled the option to read adapter’s VPD data (PN, SN etc) from the VM using lspci.
mlxconfig	Renamed the BOOT_RETRY_CNT1 parameter to BOOT_RETRY_CNT.
vPort: Multiple Guest Virtual Machine Identifier (GVMI) per Function	Added multiple e-switch vPort PCI function by GVMI, per vPort for the use cases below: • VM Live Migration SW fallback path • Container offloads • Scalable IOV
Reduced Firmware Upgrade Time	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). Note: Currently this capability only supports firmware upgrade and downgrades to firmware versions newer than XX.25.1020.
Bug Fixes	See Bug Fixes History.
Rev. 16.24.1000	

Layer 3 Encapsulation	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.
QoS Enhanced Transmission Selection (ETS)	Enabled QoS ETS for systems with 64 VFs to better allocate bandwidth in the NIC.
Time to Live (TTL) Modification	Enabled TTL modification for received packets.
VF Mirroring	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.
Transmission Histogram Counters	Added support for the transmission histogram counter set as part of the Ethernet extended group counters.
Events Generation by the Hardware upon Counter Incrementation	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.
NODNIC Connectivity	Enables NOIDNIC connectivity to the network through the e-switch and not directly to the physical port.
QP and Mkey Values	Enabled setting the QP and the Mkey values by the software upon these resources creation.
PCIe Atomic	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.
TTL Modification for Rx NIC (Steering)	Enabled TTL modification in the Rx NIC steering. When modifying the TTL in the Rx NIC, the CQE checksum will not recalculated automatically. **Note:** TTL modification in the FDB for traffic from the network is currently not supported.
TIR Destination from the FDB	Enabled a single TIR destination from the FDB.
WRED	Changed the WRED default mode to OFF for Multi-Host adapter cards.
TX Steering Rule on in WQE Ethernet Segment	Added support for TX steering rule on flow_table_metadata in WQE Ethernet segment.
L3 Encapsulation/Decapsulation in the Reformat Context Allocation	Added L3 encapsulation/decapsulation support in the reformat context allocation. • L3 encapsulation removes L2 headers and adds generic L3 tunnel encapsulation. • L3 decapsulation removes the generic L3 tunnel decapsulation and L2 header.
Flow Steering Header Modification	Added support for flow steering header modification (header rewrite) for IPv4 TTL header for loopback traffic (VF-VF/VF-PF). **Note:** TTL modification for traffic from the network is currently not supported.

Teardown: Fast Mode	**[Developers only]** Moved the fast teardown HCA cap bit to offset 0x1c.4:1.
Virtual Functions/QoS	Enabled Virtual Functions to read QPDPM/QPDP/QPTS.
Message vs. Payload based flow control QP Configuration	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.
Multi PCI RDMA IB	This capability enables the user to expose two PCI/IB devices per network port.
Steering	Enabled TTL modification in the RX NIC steering. **Note:** TTL modification in the FDB for traffic from the network is currently not supported.
Bug Fixes	See Bug Fixes History.

8 Bug Fixes History

⚠ This section includes history of bug fixes of 3 major releases back. For older releases history, please refer to the relevant firmware versions.

Internal Ref.	Issue
2078455	**Description:** Fixed an issue that caused the device to perform DMA into memory already returned to the OS during the transition between the pre-boot driver to the OS driver.
	Keywords: Memory corruption
	Discovered in Version: 16.26.4012
	Fixed in Release: 16.27.1016
2019824	**Description:** Fixed an out-of-range issue when reading or writing the q_counters.
	Keywords: q_counter, out-of-range
	Discovered in Version: 16.26.4012
	Fixed in Release: 16.27.1016
2004477	**Description:** Fixed an issue that in Multi-Host adapter cards that resulted in wrong handling of PERST on auxiliary links.
	Keywords: Multi-Host, PERST
	Discovered in Version: 16.26.4012
	Fixed in Release: 16.27.1016
1949324	**Description:** Fixed an issue that caused the ZTR counters query to always return 0.
	Keywords: ZTR counters
	Discovered in Version: 16.26.4012
	Fixed in Release: 16.27.1016
2064453	**Description:** Fixed an issue that prevented the adapter card from going into the bypass mode when the BMC disabled the hardware arbitration.
	Keywords: BMC, hardware arbitration, bypass mode
	Discovered in Version: 16.26.4012
	Fixed in Release: 16.27.1016
2047778	**Description:** Fixed an issue that required using a a non-standard tool to perform the first LFWP patching when used the GA firmware version.
	Keywords: LFWP
	Discovered in Version: 16.26.4012
	Fixed in Release: 16.27.1016
2066506	**Description:** Updated the OOOSLMask behaviour. Firmware now will apply OOOSLMask even if it is set to zero (OOOSLMask = 0).
	Keywords: OOOSLMask, port info MAD
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.27.1016

Internal Ref.	Issue
2045815	**Description:** Fixed redundant page consumption in the init_hca.
	Keywords: init_hca, pages
	Discovered in Version: 16.26.4012
	Fixed in Release: 16.27.1016
1912117	**Description:** The sw_reset option is not supported when ATS is enabled.
	Keywords: ATS, sw_reset
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.27.1016
1966962	**Description:** Fixed a possible performance degradation in small packet sizes that occured when ESWITCH_IPV4_TTL_MODIFY_ENABLE configuration was enabled in mlxconfig.
	Keywords: TTL_MODIFY, small packets degradation, performance
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.27.1016
1938614	**Description:** Due to the string DB not being updated after Live-Patch, the tracer cannot function after Live-Patch.
	Keywords: Live-Patch, LFWP, mlxfwreset, strings
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.27.1016
1993707	**Description:** Fixed a rare issue that caused other active functions to receive a malformed CQE during driver (PF or VF) unload or FLR flows.
	Keywords: Malformed CQE
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.27.1016
1899133	**Description:** Fixed an issue that prevented PCI link from being established when the firmware was corrupted.
	Keywords: PCI link
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.27.1016
1950134	**Description:** Fixed an issue that triggered a FW assert and resulted in a wrong deallocation of a resource when Packet Pacing was enabled, and a QP was being destroyed.
	Keywords: Packet Pacing, FW assert
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.27.1016
1973826	**Description:** Fixed an issue that caused the firmware to hang when an FLR occurred at the same time as the teardown. As a result, the teardown flow took a lock, and never released it because it was being aborted by an FLR.
	Keywords: FLR, teardown
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.27.1016

Internal Ref.	Issue
1751262	**Description:** Fixed an issue that caused the link to go down for 5 seconds when the other port link was down in dual port NICs. This occured due to a small noise passing through the PLL and the transceiver which are more sensitive to noise.
	Keywords: Link down
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.26.4012
1953029	**Description:** Fixed a rare issue that caused the device to hang when encountering RDMA read operations with very short retry time. In this case, the driver failed to load due to wait_fw_init (mlx5_load_one failed with error code -16).
	Keywords: RDMA, driver load failure
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1917200	**Description:** Increased the bandwidth for RoCE scenarios using a specific QPN.
	Keywords: RoCE Lossy & ECN
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1932875	**Description:** Enabled lower granularity of rate for Congestion Control.
	Keywords: Congestion Control, RoCE Lossy
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1945417	**Description:** Set the `min_time_between_cnps` default value as 2.
	Keywords: QP bandwidth, RoCE Lossy
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1970428	**Description:** Fixed an issue that prevented the meter counter from being cleared after the flow meter being updated. In this case, the Information rate took up to 10 minutes to take effect.
	Keywords: Flow Meter
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1970488	**Description:** Fixed an issue that caused a unequal cost multipath when device set the `lag_tx_port_affinity` of the QP.
	Keywords: ECMP, LAG, device affinity, `lag_tx_port_affinity`
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1900051	**Description:** Fixed a performance degradation issue that occurred when the flow group match criteria includes both IPV6 and IP_ECN.
	Keywords: IPV6, IP_ECN, firmware steering, performance
	Discovered in Version: 16.26.1040

Internal Ref.	Issue
	Fixed in Release: 16.26.4012
1922301	**Description:** Fixed a low performance issue that occurred occasionally when ran firmware steering with high insertion rate during “adding and removing” low table entries.
	Keywords: Firmware steering, insertion rate
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1752009	**Description:** 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.
	Keywords: Multi-GVMI, SR-IOV
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.26.4012
1895191	**Description:** 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.
	Keywords: SX_RDMA, Dual Port RoCE, Dual Port GVMI
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1919211	**Description:** Fixed a performance issue that occurred when ran numerous VMs with NIC Steering rules.
	Keywords: SR-IOV, NIC Steering, VMs
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1929850	**Description:** Creating an NVMoF offloaded target while running the LFWP flow may cause the device to become unstable.
	Keywords: Live Firmware Patch, LFWP, NVME
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1932516	**Description:** sw_reset action fails in case it is initiated during live-patch flow.
	Keywords: Live Firmware Patch, LFWP, sw_reset, mlxfwreset
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1955978	**Description:** Updated the firmware behavior to report PLDM version 1.1.0 regardless of whether Redfish is enabled or not.
	Keywords: PLDM, Redfish
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.2002
1900730	**Description:** Enabled the option to prevent clock and capture CPLD GPIOs glitch upon firmware reset.
	Keywords: Thermal fault issue, mlxfwreset, firmware reset
	Discovered in Version: 16.26.1040

Internal Ref.	Issue
	Fixed in Release: 16.26.4012
1946509	**Description:** Fixed an issue that slowed the firmware flows when executing many destroy XRQ commands on an XRQ that supported DC transport service.
	Keywords: DC XRQ
	Discovered in Version: 16.26.1040
	Fixed in Release: 16.26.4012
1914893/ 1849498	**Description:** Fixed an issue that caused performance degradation when working in dual-port devices under bidirectional traffic stress.
	Keywords: Performance
	Discovered in Version: 16.25.6000
	Fixed in Release: 16.26.4012
1803791	**Description:** 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.
	Keywords: PCIe Error Notification
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.26.1040
1870690	**Description:** Fixed an issue that resulted in unexpected queue pairs transitioned to error in lossy tests.
	Keywords: RoCE Lossy
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.26.1040
1723460	**Description:** Limited the maximum amount of dumps created on a PF.
	Keywords: Dump files
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.26.1040
1824111	**Description:** Renamed the GMP Mellanox Vendor Specific External Capability mask enum from IsDiagnosticCountersSupported to IsDiagnosticDataSupported.
	Keywords: GMP Mellanox Vendor Specific External Capability mask DiagnosticData
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.26.1040
1784290	**Description:** Fixed a stability issue in RoCE retransmissions under stress affecting Zero-Touch-RoCE.
	Keywords: Zero-Touch-RoCE
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.26.1040
1822787	**Description:** Fixed an issue that caused a function to misbehave when a PCIe TLP was set with a poisoned indication.
	Keywords: PCIe TLP

Internal Ref.	Issue
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.26.1040
1788895	**Description:** Fixed an issue that caused large number of packet to drop when running Jumbo frames with TTL rewrite.
	Keywords: Jumbo frames, TTL
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.26.1040
1806874	**Description:** Limited the number of the elements in the QoS tree 2K. **Note:** Creating more than 250 Vport_tc for every TC is not allowed.
	Keywords: VQoS tree , 255Vfs
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.26.1040
1769862	**Description:** 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.
	Keywords: Counters
	Discovered in Version: 16.24.1020
	Fixed in Release: 16.26.1040
1771326	**Description:** Fixed an issue that caused the nack counters to constantly be reported as “0”.
	Keywords: Counters
	Discovered in Version: 16.24.1020
	Fixed in Release: 16.26.1040
1775228	**Description:** 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.
	Keywords: QP, WQE, CQE
	Discovered in Version: 16.25.1020
	Fixed in Release: 16.26.1040
1750224	**Description:** Fixed an issue that resulted in packets sometimes being dropped while setting the software owner Flow Table as root.
	Keywords: Flow Table
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1771921	**Description:** 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.
	Keywords: DCBX
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1726413	**Description:** Fixed an issue that resulted in inaccurate counters’ value (negative value) when querying the vPort counters during traffic.

Internal Ref.	Issue
	Keywords: Counters
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1747677	**Description:** Fixed an issue that prevented the firmware from working with multiple priorities when in Host Chaining mode.
	Keywords: Host Chaining
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1682830	**Description:** Fixed a rare issue that caused the RX to hang when the server went into the Standby mode (aux power).
	Keywords: RX
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
-	**Description:** Fixed a wrong meter rate for small packets by excluding packet CRC size.
	Keywords: Meter rate
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1682727	**Description:** Fixed an issue that slowed the driver's unload process. The dmesg showed the following message: *teardown_hca, resource is still in use, gvmi=XX, type=0x2b*
	Keywords: Driver unload
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1717993	**Description:** Fixed an issue that required pages' calculations on SR-IOV, thus causing the adapter to return 'not enough resources' to the driver's VF.
	Keywords: SR-IOV
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1657685	**Description:** Fixed an issue that cause an unexpected pause counter to be transmitted when the HCA was in tear down mode.
	Keywords: Pause counter, Ethernet
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1679530	**Description:** 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
	Keywords: Performance
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020

Internal Ref.	Issue
1615586	**Description:** 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.
	Keywords: CQE
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1689159	**Description:** 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.
	Keywords: Dual port with IB and ETH, MAD Timeout
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1697465	**Description:** 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.
	Keywords: RoCE Lossy Access Register
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1706349	**Description:** Fixed an ECN compatibility issue with Broadcom RNICs.
	Keywords: ECN
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1582423	**Description:** 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.
	Keywords: Remote mirroring
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1678824	**Description:** Fixed an issue that prevented the user to enable the port after disabling it in the VF NODNIC.
	Keywords: VF NODNIC
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1606289	**Description:** Enlarged the number of modify fields to 16 to avoid IPv6 header rewrite failure.
	Keywords: IPv6 header rewrite
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1627973	**Description:** Fixed an issue that prevented IB QP counters for Acks/Responses from working as a results the NACK/OOS counters shown as zero.
	Keywords: IB QP counters for Acks/Responses
	Discovered in Version: 16.24.1000

Internal Ref.	Issue
	Fixed in Release: 16.25.1020
1554104	**Description:** Set the stateless offloads cap to be permanently '1'.
	Keywords: Stateless offloads cap
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1541193	**Description:** Fixed an issue that cause high connection setup latency on guest VMs.
	Keywords: Connection latency
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.25.1020
1554104	**Description:** Fixed a rare case that cause the transmission side to hang and an assertion in the RDMA ODP.
	Keywords: RDMA ODP
	Discovered in Version: 16.24.1000
	Fixed in Release: 16.24.1000
1434510	**Description:** Fixed a PCIe bifurcation issue.
	Keywords: PCIe
	Discovered in Version: 16.22.1002
	Fixed in Release: 16.24.1000
1424906	**Description:** Fixed a deadlock in RX related to the “send-invalidate” flow, that resulted in RX getting stuck.
	Keywords: RX, deadlock
	Discovered in Version: 16.23.1020
	Fixed in Release: 16.24.1000
929504	**Description:** Fixed rare errors in RX that resulted in double completion.
	Keywords: RX
	Discovered in Version: 16.23.1020
	Fixed in Release: 16.24.1000
1284452/ 1282926	**Description:** Fixed an issue that caused the mlxconfig tool to present all possible expansion ROM images, instead of presenting only the existing images.
	Keywords: mlxconfig
	Discovered in Version: 16.22.1002
	Fixed in Release: 16.24.1000
1475993	**Description:** Aligned the default tuning type in PHY TEST MODE to the device protocol.
	Keywords: PHY
	Discovered in Version: 16.23.1020
	Fixed in Release: 16.24.1000

Internal Ref.	Issue
1403211	**Description:** 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. **Note:** Do not run mlxfwreset when operating in a Safe Mode state.
	Keywords: mlxfwreset
	Discovered in Version: 16.23.1020
	Fixed in Release: 16.24.1000
1295606	**Description:** Fixed an issue related to PCIe "Surprise link down" event reporting capability.
	Keywords: PCIe
	Discovered in Version: 16.22.1002
	Fixed in Release: 16.24.1000
1434863	**Description:** Fixed an issue that resulted in the link partner experiencing false active linkup when plugging in a base-T cable to a closed port.
	Keywords: Interfaces
	Discovered in Version: 16.22.1002
	Fixed in Release: 16.24.1000
1424873	**Description:** Modifying VMQoS rate limiter parameters during traffic might cause transmission failure.
	Keywords: VMQoS, rate limiter
	Discovered in Version: 16.22.1002
	Fixed in Release: 16.24.1000

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. Neither NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: “NVIDIA”) make any representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice. Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation and/

or Mellanox Technologies Ltd. in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.