



NVIDIA ConnectX-4 Lx Adapter Cards Firmware Release Notes v14.32.1900

Table of contents

Firmware Compatible Products	3
Changes and New Features	22
Unsupported Features and Commands	22
Bug Fixes in this Firmware Version	24
Known Issues	25
PreBoot Drivers (FlexBoot/UEFI)	32
Supported Non-Volatile Configurations	33
Changes and New Feature History	37
Bug Fixes History	42

Release Notes Update History

Revision	Date	Description
14.32.1900	September 02, 2024	Initial release of this Release Notes version, This version introduces .Changes and New Features.

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 [Firmware Download](#).

Document Revision History

A list of the changes made to this document are provided in [Document Revision History](#).

Firmware Compatible Products

The chapter contains the following sections:

These are the release notes for the NVIDIA® ConnectX®-4 Lx adapters firmware. This firmware supports the following protocols:

- Ethernet - 1GbE, 10GbE, 25GbE, 40GbE, 50GbE
- PCI Express 3.0, supporting backwards compatibility for v2.0 and v1.1

Supported Devices

This firmware supports the devices and protocols listed below:

NVIDIA SKU	Legacy OPN	PSID	Device Description
900-9X414-0052-SN1	MCX442 1A- XCHN	MT_00 000005 88	ConnectX-4 Lx EN network interface card for OCP2.0; Type 1 with Host Management; 10GbE dual-port SFP28; PCIe3.0 x8; UEFI Enabled (x86; ARM)
900-9X4B0-0052-0S0	MCX412 1A- XCAS	MT_24 201100 04	ConnectX-4 Lx EN network interface card; 10GbE dual-port SFP28; PCIe3.0 x8; ROHS R6
900-9X429-0014-0T0	MCX413 1A- BCAT	MT_24 301100 27	ConnectX-4 Lx EN network interface card; 40GbE single-port QSFP28; PCIe3.0 x8; ROHS R6
900-9X414-0013-SN2	MCX441 1A- ACHN	MT_00 000005 01	ConnectX-4 Lx EN network interface card for OCP2.0; Type 1 with Host Management; 25GbE single-port SFP28; PCIe3.0 x8; UEFI Enabled
900-9X422-	MCX462 1A-	MT_00 000002	ConnectX-4 Lx EN network interface card for OCP 3.0; with host management; 25GbE Dual-port SFP28; PCIe3.0

NVIDIA SKU	Legacy OPN	PSID	Device Description
0053-SB0	ACAB	38	x8; ROHS R6
900-9X414-0053-SN0	MCX442 1A- ACUN	MT_00 000002 75	ConnectX-4 Lx EN network interface card for OCP; without host management; 25GbE dual-port SFP28; PCIe3.0 x8; UEFI Enabled
900-9X4B0-0052-ST0	MCX412 1A- XCHT	MT_00 000004 14	ConnectX-4 Lx EN network interface card; with host management 10GbE dual-port SFP28; PCIe3.0 x8; ROHS R6
900-9X4B0-0053-0T1	MCX412 1A- ACAT	MT_24 201100 34	ConnectX-4 Lx EN network interface card; 25GbE dual-port SFP28; PCIe3.0 x8; ROHS R6
900-9X473-0015-MN0	MCX443 1M- GCAN	MT_25 101110 32	ConnectX-4 Lx EN network interface card for OCP with Multi-Host; 50GbE single-port QSFP28; PCIe3.0 x8; ROHS R6
900-9X414-0053-SN3	MCX442 1A- ACQN	MT_24 701120 34	ConnectX-4 Lx EN network interface card for OCP with Host Management; 25GbE dual-port SFP28; PCIe3.0 x8; ROHS R6
900-9X4B0-0053-ST0	MCX412 1A- ACUT	MT_00 000002 66	ConnectX-4 Lx EN network interface card; 25GbE dual-port SFP28; PCIe3.0 x8; UEFI Enabled; tall bracket
900-9X473-0015-SN0	MCX443 1A- GCUN	MT_00 000005 06	ConnectX-4 Lx EN network interface card for OCP2.0; Type 1; with Host Management; 50GbE single-port QSFP28; PCIe3.0 x8; UEFI Enabled
900-9X422-0052-SB0	MCX462 1A- XCAB	MT_00 000005 37	ConnectX-4 Lx EN network interface card for OCP 3.0; with host management; 10GbE Dual-port SFP28; PCIe3.0 x8; Thumbscrew bracket
900-9X4B0-0012-0T1	MCX411 1A- XCAT	MT_24 101100 04	ConnectX-4 Lx EN network interface card; 10GbE single-port SFP28+; PCIe3.0 x8; ROHS R6
900-9X4B0-0013-ST0	MCX411 1A- ACUT	MT_00 000002 67	ConnectX-4 Lx EN network interface card; 25GbE single-port SFP28; PCIe3.0 x8; UEFI Enabled; tall bracket

NVIDIA SKU	Legacy OPN	PSID	Device Description
900-9X429-0015-0T0	MCX413 1A- GCAT	MT_24 301100 32	ConnectX-4 Lx EN network interface card; 50GbE single-port QSFP28; PCIe3.0 x8; ROHS R6
900-9X414-0013-SN1	MCX441 1A- ACAN	MT_24 501110 34	ConnectX-4 Lx EN network interface card for OCP; 25GbE single-port SFP28; PCIe3.0 x8; ROHS R6
900-9X414-0013-SN3	MCX441 1A- ACQN	MT_24 501120 34	ConnectX-4 Lx EN network interface card for OCP with Host Management; 25GbE single-port SFP28; PCIe3.0 x8; ROHS R6
900-9X473-0015-SN1	MCX443 1A- GCAN	MT_24 901110 32	ConnectX-4 Lx EN network interface card for OCP; 50GbE single-port QSFP28; PCIe3.0 x8; ROHS R6
900-9X4B0-0013-0Q0	MCX411 1A- ACAT	MT_24 101100 34	ConnectX-4 Lx EN network interface card; 25GbE single-port SFP28; PCIe3.0 x8; ROHS R6
900-9X414-0052-SN0	MCX442 1A- XCQN	MT_24 701100 04	ConnectX-4 Lx EN network interface card for OCP with Host Management; 10GbE dual-port SFP28; PCIe3.0 x8; ROHS R6
900-9X414-0053-SN1	MCX442 1A- ACAN	MT_24 701110 34	ConnectX-4 Lx EN network interface card for OCP; 25GbE dual-port SFP28; PCIe3.0 x8; ROHS R6
900-9X414-0013-SN0	MCX441 1A- ACUN	MT_00 000002 68	ConnectX-4 Lx EN network interface card for OCP; without host management; 25GbE single-port SFP28; PCIe3.0 x8; UEFI Enabled; no bracket
900-9X4B0-0053-ST2	MCX412 1A- ACHT	MT_00 000006 47	ConnectX-4 Lx EN network interface card; with host management; 25GbE dual-port SFP28; PCIe3.0 x8; UEFI enabled;

Driver Software, Tools and Switch Firmware

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
ConnectX-4 Lx Firmware	14.32.1900 / 14.32.1010 / 14.31.1014
MLNX_OFED	5.5-1.0.3.2 / 5.4-3.0.3.0 / 5.4-1.0.3.0 Note: For the list of the supported Operating Systems, please refer to the driver's Release Notes.
MLNX_EN (MLNX_OFED based code)	5.4-1.0.3.0 / 5.3-1.0.0.1 Note: For the list of the supported Operating Systems, please refer to the driver's Release Notes.
WinOF-2	2.80.50000 / 2.70.50000 / 2.60.50000 Note: For the list of the supported Operating Systems, please refer to the driver's Release Notes.
MFT	4.18.0 / 4.17.0 / 4.16.3 Note: For the list of the supported Operating Systems, please refer to the driver's Release Notes.
FlexBoot	3.6.502 Note: Please be aware that not all firmware binaries contain FlexBoot or UEFI, support may vary between cards.
UEFI	14.25.17 Note: Please be aware that not all firmware binaries contain FlexBoot or UEFI, support may vary between cards.
MLNX-OS	3.9.0900 onwards
Onyx	3.9.0900 onwards
NVIDIA Quantum™ Firmware	27.2008.2102 onwards
SwitchX-IB™ Firmware	11.2008.2102 onwards
SwitchX-IB 2 Firmware	15.2008.2102 onwards

Supported Cables and Modules

Please refer to the LinkX® Cables and Transceivers web page (<http://www.mellanox.com/products/interconnect/cables-configurator.php>) for the list of supported cables.

Validated and Supported 1GbE Cables

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

Validated and Supported 10GbE Cables

Speed	Cable OPN	Description
10GE	MFM1T02A-LR	NVIDIA SFP+ optical module for 10GBASE-LR
10GE	MFM1T02A-SR	NVIDIA SFP+ optical module for 10GBASE-SR
10GE	MAM1Q00A-QSA	NVIDIA cable module, ETH 10GbE, 40Gb/s to 10Gb/s, QSFP to SFP+
10GE	MC2309124-005	NVIDIA passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 5m
10GE	MC2309124-007	NVIDIA passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 7m
10GE	MC2309130-001	NVIDIA passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 1m
10GE	MC2309130-002	NVIDIA passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 2m
10GE	MC2309130-003	NVIDIA passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 3m
10GE	MC2309130-00A	NVIDIA passive copper hybrid cable, ETH 10GbE, 10Gb/s, QSFP to SFP+, 0.5m
10GE	MC3309124-004	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 4m

Speed	Cable OPN	Description
10GE	MC3309124-005	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 5m
10GE	MC3309124-006	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 6m
10GE	MC3309124-007	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 7m
10GE	MC3309130-001	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1m
10GE	MC3309130-002	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 2m
10GE	MC3309130-003	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 3m
10GE	MC3309130-00A	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 0.5m
10GE	MC3309130-0A1	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1.5m
10GE	MC3309130-0A2	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 2.5m
10GE	MCP2100-X001B	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1m, Blue Pulltab, Connector Label
10GE	MCP2100-X002B	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 2m, Blue Pulltab, Connector Label
10GE	MCP2100-X003B	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 3m, Blue Pulltab, Connector Label
10GE	MCP2101-X001B	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1m, Green Pulltab, Connector Label
10GE	MCP2104-X001B	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1m, Black Pulltab, Connector Label
10GE	MCP2104-X002B	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 2m, Black Pulltab, Connector Label

Speed	Cable OPN	Description
10GE	MCP2104-X003B	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 3m, Black Pulltab, Connector Label
10GE	MCP2104-X01AB	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 1.5m, Black Pulltab, Connector Label
10GE	MCP2104-X02AB	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 2.5m, Black Pulltab, Connector Label

Validated and Supported 25GbE Cables

Note

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

Speed	Cable OPN	Description
25GbE	MAM1Q00A-QSA28	NVIDIA cable module, ETH 25GbE, 100Gb/s to 25Gb/s, QSFP28 to SFP28
25GbE	MCP2M00-A001	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1m, 30AWG
25GbE	MCP2M00-A001E30N	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1m, Black, 30AWG, CA-N
25GbE	MCP2M00-A002	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2m, 30AWG
25GbE	MCP2M00-A002E30N	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2m, Black, 30AWG, CA-N
25GbE	MCP2M00-A003E26N	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 3m, Black, 26AWG, CA-N

Speed	Cable OPN	Description
25GbE	MCP2M00-A003E30L	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 3m, Black, 30AWG, CA-L
25GbE	MCP2M00-A004E26L	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 4m, Black, 26AWG, CA-L
25GbE	MCP2M00-A005E26L	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 5m, Black, 26AWG, CA-L
25GbE	MCP2M00-A00A	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 0.5m, 30AWG
25GbE	MCP2M00-A00AE30N	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 0.5m, Black, 30AWG, CA-N
25GbE	MCP2M00-A01AE30N	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1.5m, Black, 30AWG, CA-N
25GbE	MCP2M00-A02AE26N	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2.5m, Black, 26AWG, CA-N
25GbE	MCP2M00-A02AE30L	NVIDIA Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2.5m, Black, 30AWG, CA-L
25GbE	MFA2P10-A003	NVIDIA active optical cable 25GbE, SFP28, 3m
25GbE	MFA2P10-A005	NVIDIA active optical cable 25GbE, SFP28, 5m
25GbE	MFA2P10-A007	NVIDIA active optical cable 25GbE, SFP28, 7m
25GbE	MFA2P10-A010	NVIDIA active optical cable 25GbE, SFP28, 10m
25GbE	MFA2P10-A015	NVIDIA active optical cable 25GbE, SFP28, 15m
25GbE	MFA2P10-A020	NVIDIA active optical cable 25GbE, SFP28, 20m
25GbE	MFA2P10-A030	NVIDIA active optical cable 25GbE, SFP28, 30m
25GbE	MFA2P10-A050	NVIDIA active optical cable 25GbE, SFP28, 50m
25GbE	MMA2P00-AS	NVIDIA transceiver, 25GbE, SFP28, LC-LC, 850nm, SR, up to 100m
25GbE	SFP25G-AOC10M-TG	NVIDIA customized active optical cable 25GbE, SFP28, 10m, Aqua
25GbE	SFP25G-AOC30M-TG	NVIDIA customized active optical cable 25GbE, SFP28, 30m, Aqua

Speed	Cable OPN	Description
25GbE	SFP25G-AOC07M-TG	NVIDIA customized active optical cable 25GbE, SFP28, 7m, Aqua
25GbE	SFP25G-AOC05M-TG	NVIDIA customized active optical cable 25GbE, SFP28, 5m, Aqua
25GbE	SFP25G-AOC03M-TG	NVIDIA customized active optical cable 25GbE, SFP28, 3m, Aqua
25GbE	SFP25G-AOC20M-TG	NVIDIA customized active optical cable 25GbE, SFP28, 20m, Aqua
25GbE	MMA2P00-AS_FF	NVIDIA transceiver, 25GbE, SFP28, LC-LC, 850nm, SR, up to 100m
25GbE	MMA2P00-AS-SP	NVIDIA transceiver, 25GbE, SFP28, LC-LC, 850nm, SR, up to 100m, single package
25GbE	MMA2L20-AR	NVIDIA optical transceiver, 25GbE, 25Gb/s, SFP28, LC-LC, 1310nm, LR up to 10km

Validated and Supported 40GbE Cables

Speed	Cable OPN	Description
40GbE	MC2206128-004	NVIDIA passive copper cable, VPI, up to 40Gb/s, QSFP, 4m
40GbE	MC2206128-005	NVIDIA passive copper cable, VPI, up to 40Gb/s, QSFP, 5m
40GbE	MC2206130-001	NVIDIA passive copper cable, VPI, up to 40Gb/s, QSFP, 1m
40GbE	MC2206130-002	NVIDIA passive copper cable, VPI, up to 40Gb/s, QSFP, 2m
40GbE	MC2206130-003	NVIDIA passive copper cable, VPI, up to 40Gb/s, QSFP, 3m

Speed	Cable OPN	Description
40GE	MC2206130-00A	NVIDIA passive copper cable, VPI, up to 40Gb/s, QSFP, 0.5m
40GE	MC2210126-004	NVIDIA passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 4m
40GE	MC2210126-005	NVIDIA passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 5m
40GE	MC2210128-003	NVIDIA passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 3m
40GE	MC2210130-001	NVIDIA passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 1m
40GE	MC2210130-002	NVIDIA passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 2m
40GE	MC2210310-003	NVIDIA active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 3m
40GE	MC2210310-005	NVIDIA active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 5m
40GE	MC2210310-010	NVIDIA active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 10m
40GE	MC2210310-015	NVIDIA active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 15m
40GE	MC2210310-020	NVIDIA active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 20m
40GE	MC2210310-030	NVIDIA active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 30m
40GE	MC2210310-050	NVIDIA active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 50m
40GE	MC2210310-100	NVIDIA active fiber cable, ETH 40GbE, 40Gb/s, QSFP, 100m
40GE	MC2210411-SR4E	NVIDIA optical module, 40Gb/s, QSFP, MPO, 850nm, up to 300m

Speed	Cable OPN	Description
40GE	MC2609125-005	NVIDIA passive copper hybrid cable, ETH 40GbE to 4x10GbE, QSFP to 4xSFP+, 5m
40GE	MC2609130-001	NVIDIA passive copper hybrid cable, ETH 40GbE to 4x10GbE, QSFP to 4xSFP+, 1m
40GE	MC2609130-003	NVIDIA passive copper hybrid cable, ETH 40GbE to 4x10GbE, QSFP to 4xSFP+, 3m
40GE	MCP1700-B001E	NVIDIA passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 1m, Black Pulltab
40GE	MCP1700-B002E	NVIDIA passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 2m, Black Pulltab
40GE	MCP1700-B003E	NVIDIA passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 3m, Black Pulltab
40GE	MCP1700-B01AE	NVIDIA passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 1.5m, Black Pulltab
40GE	MCP1700-B02AE	NVIDIA passive copper cable, ETH 40GbE, 40Gb/s, QSFP, 2.5m, Black Pulltab
40GE	MMA1B00-B150D	NVIDIA transceiver, 40GbE, QSFP+, MPO, 850nm, SR4, up to 150m, DDMI
40GE	MCP7900-X01AA	NVIDIA passive copper hybrid cable, ETH 40GbE to 4x10GbE, QSFP to 4xSFP+, 1.5m, Blue Pulltab, customized label
40GE	MCP7904-X002A	NVIDIA passive copper hybrid cable, ETH 40GbE to 4x10GbE, QSFP to 4xSFP+, 2m, Black Pulltab, customized label
40GE	MCP7904-X003A	NVIDIA passive copper hybrid cable, ETH 40GbE to 4x10GbE, QSFP to 4xSFP+, 3m, Black Pulltab, customized label
40GE	MCP7904-X01AA	NVIDIA passive copper hybrid cable, ETH 40GbE to 4x10GbE, QSFP to 4xSFP+, 1.5m, Black Pulltab, customized label
40GE	MCP7904-X02AA	NVIDIA passive copper hybrid cable, ETH 40GbE to 4x10GbE, QSFP to 4xSFP+, 2.5m, Black Pulltab, customized label
40GE	MC2210511-LR4	NVIDIA Optical Module 40Gb/s FDR 10 QSFP LC-LC 1310nm LR4 up to 10km

Speed	Cable OPN	Description
40GE	MC6709309-005	NVIDIA passive fiber hybrid cable, MPO to 8xLC, 5m
40GE	MC6709309-010	NVIDIA passive fiber hybrid cable, MPO to 8xLC, 10m
40GE	MC6709309-020	NVIDIA passive fiber hybrid cable, MPO to 8xLC, 20m
40GE	MC6709309-030	NVIDIA passive fiber hybrid cable, MPO to 8xLC, 30m

Validated and Supported 56GbE Cables

Note

The 56GbE cables are used to raise 40GbE link speed as the 56GbE speed is not supported.

Speed	Cable OPN	Description
56GE	MC2207126-004	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 4m
56GE	MC2207128-003	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 3m
56GE	MC2207128-0A2	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 2.5m
56GE	MC2207130-001	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 1m

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

Speed	Cable OPN	Description
56GE	MCP1700-F001D	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 1m, Yellow Pulltab
56GE	MCP1700-F002C	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 2m, Red Pulltab
56GE	MCP1700-F002D	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 2m, Yellow Pulltab
56GE	MCP1700-F003C	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 3m, Red Pulltab
56GE	MCP1700-F003D	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 3m, Yellow Pulltab
56GE	MCP170L-F001	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 1m
56GE	MCP170L-F002	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 2m
56GE	MCP170L-F003	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 3m
56GE	MCP170L-F00A	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 0.5m
56GE	MCP170L-F01A	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, LSZH, 1.5m

Supported 3rd Party Cables and Modules

Speed	Cable OPN	Description
10GbE	74752-9096	Dell Active DAC SFP+, Cisco PN SFP-H10GB-CU5M, Molex PN 74752-9096
10GbE	74752-9096 (SFP-H10GB-SU5M)	Cisco-Molex INC Active DAC SFP+ 5m

Speed	Cable OPN	Description
10GbE	74752-9521	CISCO-MOLEX SFP28/SFP+ 10G Passive copper cable
10GbE	74752-9521 (SFP-H10GB-CU5M)	Cisco 10GBASE SFP+ modules
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-1M	Arista Compatible 10G SFP+ Passive Cable 1m
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	CAB-SFP-SFP-5M	Arista Compatible 10G SFP+ Passive Cable 5m
10GbE	FTLX1471D3BCL-ME	10GBASE-LR SFP+ 1310nm 10km DOM Transceiver Module
10GbE	FTLX1471D3BCL-ME	10GBASE-LR SFP+ 1310nm 10km DOM Transceiver Module
10GbE	FTLX8570D3BCL-C2	Cisco FET-10G 10-2566-02 FTLX8570D3BCL-C2 10Gbps Fabric Extender SFP+ 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

Speed	Cable OPN	Description
10GbE	SFP-H10GB-CU1M	Cisco 1-m 10G SFP+ Twinax cable assembly, passive
10GbE	SFP-H10GB-CU3M	Cisco 3-m 10G SFP+ Twinax cable assembly, passive
10GbE	SFP-H10GB-CU5M	Cisco 5-m 10G SFP+ Twinax cable assembly, passive
25GbE	FTLF8536P4BCL	TRANSCEIVER 25GBE SFP SR
25GbE	LTF8507-PC07	HISENSE ACTIVE FIBER CABLE, 25GBE
25GbE	SFP-H25G-CU3M	CISCO 25GBASE-CR1 COPPER CABLE 3-METER NDCCGJ-C403
40GbE	00D5811-N13445C	IBM Cable DAC 40GbE QSFP+ to QSFP+ Passive Copper 7m
40GbE	L45593-D108-D30	PASSIVE COPPER CABLE ETH 40GBE QSFP 3M
40GbE	L45593-D118-D10	QSFP-H40G-CU1M
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	Cisco QSFP-4SFP10G-CU5M
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-CU1M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 1-meter, passive

Speed	Cable OPN	Description
40GbE	QSFP-H40G-CU3M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 3-meter, passive
40GbE	QSFP-H40G-CU5M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 5-meter, passive
100GbE	10137498-2005LF	HPE 100GbE 2m copper cable
100GbE	10137498-2010LF	HPE 100GbE 4m copper cable
100GbE	CBL-00195-02	100GbE QSFP28 to QSFP28 copper cable 3M

Tested Switches

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

Speed	Switch Silicon	OPN # / Name	Description	Vendor
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

Tested 100GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
100GbE	Spectrum-3	MSN4600-XXXX	64-port Non-blocking 100GbE Open Ethernet Switch System	NVIDIA
100GbE	Spectrum-2	MSN3700C-XXXX	32-port Non-blocking 100GbE Open Ethernet Switch System	NVIDIA
100GbE	Spectrum-2	MSN3420-XXXX	48 SFP + 12 QSFP ports Non-blocking 100GbE Open Ethernet Switch System	NVIDIA
100GbE	Spectrum	MSN2410-XXXX	48-port 25GbE + 8-port 100GbE Open Ethernet Switch System	NVIDIA
100GbE	Spectrum	MSN2700-XXXX	32-port Non-blocking 100GbE Open Ethernet Switch System	NVIDIA
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

Speed	Switch Silicon	OPN # / Name	Description	Vendor
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

PRM Revision Compatibility

This firmware version complies with the following Programmer's Reference Manual:

- Mellanox Adapters Programmer's Reference Manual (PRM), Rev 0.53 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.

Changes and New Features

Note

SR-IOV - Virtual Functions (VF) per Port - The maximum Virtual Functions (VF) per port is 127. For further information, see [RoCE Limitations](#).

Changes and New Feature in this Firmware Version

Feature/Change	Description
14.32.1900	
Security Hardening Enhancements	This release contains important reliability improvements and security hardening enhancements. NVIDIA recommends upgrading your devices firmware to this release to improve the devices' firmware security and reliability.

Unsupported Features and Commands

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
- DC in Multi-Host, SR-IOV, and Ethernet (RoCE)
- RoCE LAG for VFs
- Mutlihost Ethernet

Unsupported Commands

- QUERY_MAD_DEMUX
- SET_MAD_DEMUX
- CREATE_RQ - MEMORY_RQ_RMP
- MODIFY_LAG_ASYNC_EVENT

Bug Fixes in this Firmware Version

For a list of old Bug Fixes, please see [Bug Fixes History](#).

This version does not include Bug Fixes.

Known Issues

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	45	32	20	127	100

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	26	127	55

Known Issues

Internal Ref.	Issue
2867868	Description: Resetting the firmware with bond configuration may lead to mkey errors in dmesg.
	Workaround: N/A
	Keywords: Bond configuration, firmware reset, mkey errors
	Discovered in Version: 14.29.1016
2730077	Description: When connecting a ConnectX-4 Lx adapter card against Cisco N9K switch with Finisar module type: FTLF8536P4BCL, the linkup time in AutoNeg NO-FEC mode is up to 2min.
	Workaround: N/A

Internal Ref.	Issue
	<p>Keywords: AN, linkup time, Finisar module</p> <p>Discovered in Version: 14.29.1016</p>
2378593	<p>Description: Sub 1sec firmware update (fast reset flow) is not supported when updating from previous releases to the current one. Doing so may cause network disconnection events.</p> <p>Workaround: Use full reset flow for firmware upgrade/downgrade.</p> <p>Keywords: Sub 1sec firmware update</p> <p>Discovered in Version: 14.29.1016</p>
2396506	<p>Description: On systems with high PCIe latency (2us or above), lower bandwidth may be experienced.</p> <p>Workaround: If such issue is observed:</p> <ol style="list-style-type: none"> 1. Enable ZTT to overcome the high latency. Run: <code>mlxconfig -d <mst device> set ZERO_TOUCH_TUNING_ENABLE=1</code> 2. Reset or power cycle the firmware for change to take effect <p>Keywords: Performance, ZTT</p> <p>Discovered in Version: 14.29.1016</p>
2213356	<p>Description: The following are the Steering Dump limitations:</p> <ul style="list-style-type: none"> • Supported only on ConnectX-5 adapter cards • Requires passing the version (FW/Stelib/MFT) and device type to stelib • Re-format is not supported • Advanced multi-port feature is not supported – LAG/ROCE_AFFILIATION/MPFS_LB/ESW_LB (only traffic vhca <-> wire) • Packet types supported: <ul style="list-style-type: none"> ◦ Layer 2 Eth ◦ Layer 3 IPv4/Ipv6/Grh ◦ Layer 4 TCP/UDP/Bth/GreV0/GreV1 ◦ Tunneling VXLAN/Geneve/GREv0/Mpls • FlexParser protocols are not supported (e.g AliVxlan/VxlanGpe etc..). • Compiles only on x86

Internal Ref.	Issue
	Workaround: N/A
	Keywords: Steering Bump
	Discovered in Version: 14.29.1016
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
1840289	Discovered in Version: 14.27.1016
	Description: Since Packet Pacing enforce max_tc value is "1", features that require multiple TCs will not be active when this mode is available.
	Workaround: N/A
1796628	Keywords: Packet Pacing
	Discovered in Version: 14.26.1040
	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.
1699214	Workaround: N/A
	Keywords: Performance, unicast loopback traffic, multicast loopback traffic
	Discovered in Version: 14.26.1040
1689186	Description: NODNIC VF is partially tested. It is fully tested only in ConnectX-5 adapter cards.
	Workaround: N/A
	Keywords: NODNIC VF
1699214	Discovered in Version: 14.25.1020
	Description: Changing priority to TC map during traffic might cause packet drops.
	Workaround: N/A

Internal Ref.	Issue
	<p>Keywords: QoS</p> <p>Discovered in Version: 14.25.1020</p>
1604699	<p>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.</p> <p>Workaround: N/A</p> <p>Keywords: Counters</p> <p>Discovered in Version: 14.25.1020</p>
-	<p>Description: In Ethernet mode, at 10/40GbE speeds, only NO-FEC in Force mode is supported. Other user configurations are overridden.</p> <p>Workaround: N/A</p> <p>Keywords: Ethernet, 10GbE, 40GbE, RS-FEC</p> <p>Discovered in Version: 14.25.1020</p>
1498399	<p>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).</p> <p>Workaround: N/A</p> <p>Keywords: XRC SRQ/RMP ODP</p> <p>Discovered in Version: 14.25.1020</p>
1426283	<p>Description: An mlxconfig configuration followed by the init 0 command might cause the HCA to remain in high power instead of entering the standby state. Such behavior will require a server reboot after running the mlxconfig tool.</p> <p>Workaround: N/A</p> <p>Keywords: mlxconfig</p> <p>Discovered in Version: 14.24.1000</p>
1546492	<p>Description: Executing the update_lid command while the IB port sniffer utility is active can stop the utility.</p>

Internal Ref.	Issue
	Workaround: N/A
	Keywords: IB Sniffer
	Discovered in Version: 14.24.1000
1537898	Description: Initializing a function while the IB port sniffer utility is active can stop the utility.
	Workaround: N/A
	Keywords: IB Sniffer
1332714/1 345824	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
	Discovered in Version: 14.23.1020
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
1350794	<p data-bbox="342 1461 1469 1568">Description: Encapsulation / Decapsulation support in steering has the following limitations:</p> <ul data-bbox="391 1591 1446 1892" style="list-style-type: none"> <li data-bbox="391 1591 1446 1677">• Encapsulation / Decapsulation can be open on the FDB only if all VFs are non active. <li data-bbox="391 1680 1446 1766">• Encapsulation / Decapsulation supports single mode only: FDB / NIC. Opening tables of both types is not supported. <li data-bbox="391 1768 1446 1854">• RoCE is supported only when no Encapsulation / Decapsulation is opened on the FDB. <li data-bbox="391 1856 1446 1892">• Encapsulation / Decapsulation per device support:

Internal Ref.	Issue																																			
	<table border="1"> <thead> <tr> <th></th> <th></th> <th>NIC</th> <th>FDB</th> <th></th> </tr> </thead> <tbody> <tr> <td>ConnectX-4</td> <td>encap</td> <td>NO</td> <td>YES</td> <td>non-MH</td> </tr> <tr> <td></td> <td>decap</td> <td>NO</td> <td>NO</td> <td></td> </tr> <tr> <td>ConnectX-4 Lx</td> <td>encap</td> <td>NO</td> <td>YES</td> <td>non-MH</td> </tr> <tr> <td></td> <td>decap</td> <td>NO</td> <td>YES</td> <td></td> </tr> <tr> <td>ConnectX-5</td> <td>encap</td> <td>YES</td> <td>YES</td> <td></td> </tr> <tr> <td></td> <td>decap</td> <td>YES</td> <td>YES</td> <td></td> </tr> </tbody> </table>			NIC	FDB		ConnectX-4	encap	NO	YES	non-MH		decap	NO	NO		ConnectX-4 Lx	encap	NO	YES	non-MH		decap	NO	YES		ConnectX-5	encap	YES	YES			decap	YES	YES	
		NIC	FDB																																	
ConnectX-4	encap	NO	YES	non-MH																																
	decap	NO	NO																																	
ConnectX-4 Lx	encap	NO	YES	non-MH																																
	decap	NO	YES																																	
ConnectX-5	encap	YES	YES																																	
	decap	YES	YES																																	
	Workaround: N/A																																			
	Keywords: Steering Encapsulation / Decapsulation																																			
	Discovered in Version: 14.23.1020																																			
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: 14.24.1000																																			
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: 14.23.1020																																			
1355883	Description: Running the QUERY_VPORT_COUNTER command with clear bit results in discard counters being reset.																																			
	Workaround: N/A																																			
	Keywords: Discard counters																																			
	Discovered in Version: 14.22.1002																																			
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.																																			

Internal Ref.	Issue
	<p>Workaround: N/A</p> <p>Keywords: Ethernet multicast loopback packet</p> <p>Discovered in Version: 14.22.1002</p>
1047184	<p>Description: RDMA resq_local_length_error and resp_remote_invalid_request counters do not function properly.</p> <p>Workaround: N/A</p> <p>Keywords: RDMA counters</p> <p>Discovered in Version: 14.21.1000</p>
1168594	<p>Description: RoCE Dual Port Mode (a.k.a Multi-Port vHCA: MPV) is not supported in Multi-Host setups.</p> <p>Workaround: N/A</p> <p>Keywords: Multi-Port vHCA, Multi-Host</p> <p>Discovered in Version: 14.21.1000</p>

PreBoot Drivers (FlexBoot/UEFI)

FlexBoot Changes and New Features

For further information, please refer to the [FlexBoot Release Notes](#).

UEFI Changes and Major New Features

For further information, please refer to the [UEFI Release Notes](#).

Supported Non-Volatile Configurations

Configuration	mlxconfig Parameter Name	Class	TLV ID
NV_MEMIC_CONF	MEMIC_BAR_SIZE	GLOBAL (0)	0x6
	MEMIC_SIZE_LIMIT		
NV_HOST_CHAINING_CONF	HOST_CHAINING_MODE		0x8
	HOST_CHAINING_DESCRIPTOR		
	HOST_CHAINING_TOTAL_BUFFER_SIZE		
NV_FLEX_PARS_CONF	FLEX_PARSER_PROFILE_ENABLE		0xe
	FLEX_IPV4_OVER_VXLAN_PORT		
NV_ROCE_1_5_CONF	ROCE_NEXT_PROTOCOL		0x10
NV_INTERNAL_RESOURCE_CONF	ESWITCH_HAIRPIN_DESCRIPTOR		0x13
	ESWITCH_HAIRPIN_TOT_BUFFER_SIZE		
NV_GLOBAL_PCI_CONF	NON_PREFETCHABLE_PF_BAR		0x80
	NUM_OF_VFS		
	SRIOV_EN		
	PF_LOG_BAR_SIZE		
	VF_LOG_BAR_SIZE		
	NUM_PF_MSIX		
	NUM_VF_MSIX		

Configuration	mlxconfig Parameter Name	Class	TLV ID
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		0x10 a
	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		0x19 0
	DCR_LIFO_SIZE		
NV_VPI_LINK_TYPE	LINK_TYPE	PHYSICAL_PORT (2)	0x12
NV_ROCE_CC	ROCE_CC_PRIO_MASK		0x10 7
	ROCE_CC_ALGORITHM		
NV_ROCE_CC_ECN	CLAMP_TGT_RATE_AFTER_TIME_I NC	0x10 8	
	CLAMP_TGT_RATE		
	RPG_TIME_RESET		
	RPG_BYTE_RESET		
	RPG_THRESHOLD		
	RPG_MAX_RATE		
	RPG_AI_RATE		
	RPG_HAI_RATE		
	RPG_GD		

Configuration	mlxconfig Parameter Name	Class	TLV ID
	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)	0x11 2
	MULTI_PORT_VHCA_EN		
NV_EXTERNAL_PORT_CTRL	PORT_OWNER		0x19 2
	ALLOW_RD_COUNTERS		
	RENEG_ON_CHANGE		
	TRACER_ENABLE		
NV_ROM_BOOT_CONF2	IP_VER		0x19 5
	BOOT_UNDI_NETWORK_WAIT		
NV_ROM_UEFI_CONF	UEFI_HII_EN		0x19 6
NV_ROM_UEFI_DEBUG_LEVEL	BOOT_DBG_LOG		0x20 6
	UEFI_LOGS		
NV_ROM_BOOT_CONF1	BOOT_VLAN		0x22 1
	LEGACY_BOOT_PROTOCOL		
	BOOT_RETRY_CNT		
	BOOT_LACP_DIS		
	BOOT_VLAN_EN		
NV_ROM_IB_BOOT_CONF	BOOT_PKEY	0x22 2	
NV_PCI_CONF	ADVANCED_PCI_SETTINGS		0x80
SAFE_MODE_CONF	SAFE_MODE_THRESHOLD	HOST (7)	0x82
	SAFE_MODE_ENABLE		

Changes and New Feature History

i Note

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
14.32.1010	
Congestion Control Key	Added a Congestion Control Key to all Congestion Control MADs to authenticate that they are originated from a trusted source.
SMP Firewall	Added an SMP firewall to block the option of sending SMPs (MADS sent on QP0 from the Subnet Manager) from unauthorized hosts to prevent fake SMPs from being recognized as the SM.
Vendor Specific MADs: Class 0x9	Vendor Specific MADs Class 0x9 is no longer supported by the firmware. If case the firmware detects such MAD, the firmware will return a "NOT SUPPORTED" error to the user.
Asserts' Severity Level	Added 3 new assert filters (Health buffer, NVlog, FW trace). The assert will be exposed now if its severity level is equal to or above the new filter. The filters are configurable by the ini file. The "Health buffer" filter is also configurable by new access register.

Feature/Change	Description
Rate Limit per VM instead of VM-TC	Enabled Rate Limit per VM instead of VM-TC. This capability is implemented by adding support to a new Scheduling element type: rate limit elements that will connect to the rate_limit and will share its rate limit.
Asymmetrical VFs per PF	Added support for asymmetrical VFs per PF. To enable it: PF_NUM_OF_VF_VALID must be true, and PF_NUM_OF_VF to a non-zero value.
mlxlink Support to read/write Access Registers by LID	Added 2 new MAD access registers to enable mlxlink to read/write access registers by LID (to the whole subnet).
Bug Fixes	See <i>Bug Fixes in this Firmware Version</i> section.

Feature/Change	Description
14.31.1014	
Using NC-SI Commands for Debugging PCI Link Failures	Implemented a new NC-SI command <code>get_debug_info</code> to get mstdump via the NC-SI protocol to debug a device if the PCI link fails for any given reason.
Enable/Disable RDMA via the UEFI HII System Settings	Added support for Enabling/Disabling NIC and RDMA (port/partition) via the UEFI HII system settings. Note: Values set in this option only take effect when is Ethernet mode.
Increased the Maximum Number of MSIX per VF	Increased the maximum number of MSIX per VF to 127. Note that increasing the number of MSIX per VF (<code>NUM_VF_MSIX</code>) affects the configured number of VFs (<code>NUM_OF_VFS</code>). The firmware may reduce the configured number of MSIX per VF and/or the number of VFs with respect to maximum number of MSIX vectors supported by the device (<code>MAX_TOTAL_MSIX</code>).

Feature/Change	Description
Adaptive Routing (AR): multi_path, data_in_order	Added a new bit ("data_in_order") to query the QP and allow a process/library to detect when the AR is enabled.
flex_parser for GENEVE Hardware Offload and ICMP	Added a new flex parser to support GENEVE hardware offload and ICMP.
Non-Page-Supplier-FLR	When the non-page-supplier-FLR function is initiated, the firmware triggers a page event to the page supplier to indicate that all pages should be returned for the FLR function. Pages are returned by the driver to the kernel without issuing the MANAGE_PAGES commands to the firmware.
Bug Fixes	See Bug Fixes section.
14.30.1004	
Performance: Steering	Added support for a new NV config mode "icm_cache_mode_large_scale_steering" that enables less cache misses and improves performance for cases when working with many steering rules. This capability is enabled using the mlxconfig parameter "ICM_CACHE_MODE" .
Bug Fixes	See Bug Fixes section.
14.29.2002	
Bug Fixes	See Bug Fixes section.
14.29.1016	
Ethernet wqe_too_small Mode	Added a new counter per vPort that counts the number of packets that reached the Ethernet RQ but cannot fit into the WQE due to their large size. Additionally, we added the option to control if such packet will cause "CQE with Error" or "CQE MOCK".

Feature/Change	Description
Counters	Added support for the <code>cq_overrun</code> counter. The counter represents the number of times CQs enter an error state due to overflow that occur when the device tries to post a CQE into a full CQ buffer.
Pause Frames from VFs	[Beta] Enabled the capability to allow Virtual Functions to send Pause Frames packets.
Auto-Sensing	Enabled 10/25GbE auto-sensing with 3rd party when using 10/25GbE optical cables.
Steering Dump	Hardware steering dump output used for debugging and troubleshooting. Please see Known Issue 2213356 for its limitations.
14.28.2006	
Cables	Added support for Rate-Select on SFP+ (AFBR-57F5UMZ-EP1) cables to raise 10GbE link speed.
NC-SI	Added a new NC-SI command (<code>get_device_id</code>) to report a unique device identifier.
NC-SI	Added new NC-SI commands (<code>get_lldp_nb</code> , <code>set_lldp_nb</code>) to query the current status of LLDP and to enable/disable it.
Bug Fixes	See Bug Fixes .
14.28.1300	
Bug Fixes	See Bug Fixes .
14.28.1002	
Power Consumption	Added <code>mlxconfig</code> support for power reduction: <ul style="list-style-type: none"> • PCI CAP • AUTO_POWER_SAVE_LINK_DOWN
Resource dump	Added the following segments, as appeared in the PRM, to the Resource Dump: <ul style="list-style-type: none"> • PRM_QUERY_QP • PRM_QUERY_CQ

Feature/Change	Description
	<ul style="list-style-type: none"> • PRM_QUERY_MKEY • QUERY_VNIC_ENV
NC-SI 1.2 New Command	Implemented the following new command from NS-SI 1.2 specification: <ul style="list-style-type: none"> • Get PF Assignment
Parallel VF init/Teardown Performance Optimization	Improved init_hca performance in Parallel Function initialization.
Bug Fixes	See Bug Fixes .

Bug Fixes History

Note

This section includes history of bug fixes of 3 major releases back. For older releases history, please refer to the relevant firmware versions Release Notes in <https://docs.mellanox.com/category/adapterfw>.

Internal Ref.	Issue
2796324	Description: Fixed an issue that resulted in firmware getting stuck and causing unexpected behavior when connecting an optical transceiver that support RXLOS, and the remote side port was down.
	Keywords: cables, RXLOS
	Discovered in Version: 14.31.1014
	Fixed in Release: 14.32.1010
2784304	Description: Fixed an issue that prevented the system from creating more than 128K QPs.
	Keywords: QP
	Discovered in Version: 14.31.1014
	Fixed in Release: 14.32.1010
2843888	Description: Fixed a rare case where the the system got stuck when a peer port went down while using an Optical module.
	Keywords: Cables
	Discovered in Version: 14.31.1014
	Fixed in Release: 14.32.1010

Internal Ref.	Issue
2678394	Description: Limited the external loopback speed to the used module's capabilities.
	Keywords: Cables
	Discovered in Version: 14.31.1014
	Fixed in Release: 14.32.1010
2771990	Description: Improved linkup time when using the fast linkup capability.
	Keywords: Linkup time
	Discovered in Version: 14.31.1014
	Fixed in Release: 14.32.1010
2710956	Description: Added support for the slow_restart and slow_restart_idle parameters to enable Zero Touch RoCE capability.
	Keywords: Zero Touch RoCE, RoCE, slow_restart
	Discovered in Version: 14.31.1014
	Fixed in Release: 14.32.1010
2772447	Description: Fixed an issue that resulted in the Packet Pacing rate being used if the asymmetric VFs was enabled.
	Keywords: Packet Pacing rate, asymmetric VF
	Discovered in Version: 14.31.1014
	Fixed in Release: 14.32.1010
2823281	Description: Fixed an issue that resulted in wrong RNR timeout when trying to set it during the rts2rts_qp transition.
	Keywords: RNR timeout
	Discovered in Version: 14.31.1014
	Fixed in Release: 14.32.1010
2795721	Description: Fixed an issue with RSS on IPSec flows in ConnectX-4 Lx leading to performance degradation. In this scenario, the SPI optimization caused packets from a given host to hash to the same CPU core.

Internal Ref.	Issue
	The fix was to ignore SPI optimization according to l4_type in ConnectX-4 Lx adapter cards.
	Keywords: SPI
	Discovered in Version: 14.31.1014
	Fixed in Release: 14.32.1010
2748449	Description: Altered the GetInventory NC-SI command to not report leading 0xf in firmware version when it starts with 0.
	Keywords: NC-SI, GetInventory, leading 0, FW version
	Discovered in Version: 14.31.1014
	Fixed in Release: 14.32.1010

Internal Ref.	Issue
2450264	Description: Fixed an issue that caused TX PRBS not to change after reconfiguring it. Now all PRBS mode are enabled in test mode.
	Keywords: PRBS
	Discovered in Version: 14.30.1004
	Fixed in Release: 14.31.1014
2641734	Description: Fixed the rate select mechanism in QSFP modules.
	Keywords: Cables
	Discovered in Version: 14.30.1004
	Fixed in Release: 14.31.1014
2600783	Description: Fixed classification issues for "Passive" cables to be more robust.
	Keywords: Cables
	Discovered in Version: 14.30.1004
	Fixed in Release: 14.31.1014
2391109	Description: Fixed an issue that caused a fatal error, and eventually resulted in the HCA hanging when a packet was larger than a strided

Internal Ref.	Issue
	receive WQE that was being scattered.
	Keywords: Strided RQ, MTU
	Discovered in Version: 14.30.1004
	Fixed in Release: 14.31.1014
2621704	<p>Description: Fixed the resource number size (a 64 bit number) to avoid a scenario where it overwrote it with a 32 bit number and erased the high bits when de-allocating the resource number. In this scenario, when two resource numbers had identical low 32 bits, and because the high bits were cleared, it resulted in the same idx. Consequently, when two idxes were identical, then it freed the same idx twice.</p>
	Keywords: Resource number size, free_4k page
	Discovered in Version: 14.30.1004
	Fixed in Release: 14.31.1014
2619161	<p>Description: Initialized the rate table in the static configuration so it will be configured at the link-not-up scenarios.</p>
	Keywords: RoCE, static configuration, rate table
	Discovered in Version: 14.30.1004
	Fixed in Release: 14.31.1014
2444837	<p>Description: Set the cap to 0 for high index functions to avoid too many parallel VF NODNIC functions.</p>
	Keywords: NODNIC, VF, ETH PXE
	Discovered in Version: 14.29.1016
	Fixed in Release: 14.30.1004
2479788	<p>Description: Fixed an issue that caused Tx to hang when a duplicate packet rollback occurred.</p>
	Keywords: Transport, Rxt Checks
	Discovered in Version: 14.29.1016
	Fixed in Release: 14.30.1004

Internal Ref.	Issue
2379573	Description: Low performance might be experienced when upgrading from previous firmware version to 14.29.1000 when using “Fast FW Reset”.
	Keywords: “Fast FW Reset”, performance
	Discovered in Version: 14.29.1016
	Fixed in Release: 14.29.2002
2384583	Description: Fixed an issue that prevented events from being sent when only the DCBX oper version was changed.
	Keywords: Events, DCBX
	Discovered in Version: 14.29.1016
	Fixed in Release: 14.29.2002
2301590	Description: Congestion Control may not work properly if the card supports two ports and each PF for each port is not raised at the same time.
	Keywords: Congestion Control
	Discovered in Version: 14.29.1016
	Fixed in Release: 14.29.2002
2339971	Description: Fixed an issue that prevented MCAM from reporting support for MFBA, MFBE, MFPA registry keys although they were available through the CMDIF interface.
	Keywords: MCAM
	Discovered in Version: 14.29.1016
	Fixed in Release: 14.29.2002
2410395	Description: Fixed an issue that prevented a SFP28 cable from linking up in a 25GbE speed.
	Keywords: Cables
	Discovered in Version: 14.29.1016
	Fixed in Release: 14.29.2002

Internal Ref.	Issue
2093381	<p>Description: Modified the calculation of NUM_VF_MSIX to take into account NVME, Virtio Net/Blk, HotPlug PFs & VFs. Since max_total_msix is the maximum number used for all PFs and VFs (Port, NVME, Virtio Net/Blk, HotPlug), if there are not enough MSIX for all the devices, the number of port VF MSIX may be lowered (less than NUM_VF_MSIX) in order to not exceed the max_total_msix.</p> <p>Note: In case of compatibility issues with an old driver requiring more than 4 MSI-X, you should consider lowering number of PFs/VFs on any of the configurable functions (NVME, Virtio Net/Blk).</p>
	<p>Keywords: MSIX</p>
	<p>Discovered in Version: 14.29.1016</p>
	<p>Fixed in Release: 14.29.2002</p>
2360496	<p>Description: Changed the default value of DCQCN's NP parameter min_time_between_cnps to 4 on all devices to support larger scalability of cluster.</p>
	<p>Keywords: RoCE, Congestion control, DCQCN</p>
	<p>Discovered in Version: 14.28.1002</p>
	<p>Fixed in Release: 14.29.1016</p>
2336284	<p>Description: Fixed an issue that caused packets to drop due to header size issues and/or failing checks. The issue was caused due to a Linux issue that caused VFs to set the wrong header size value in wqe_inline_header_mode input.</p>
	<p>Keywords: ETH, VF, Linux VM, DSCP, wqe_inline_header_mode</p>
	<p>Discovered in Version: 14.28.1002</p>
	<p>Fixed in Release: 14.29.1016</p>
2245422	<p>Description: When MKEY_BY_NAME is enabled by NVCONFIG and a large number of VFs are configured, VM restart (VF/PF FLR) will take longer than when MKEY_BY_NAME is disabled.</p>
	<p>Keywords: SR-IOV</p>
	<p>Discovered in Version: 14.28.1002</p>
	<p>Fixed in Release: 14.29.1016</p>

Internal Ref.	Issue
2252559	Description: On rare cases, a fatal error related to errors from the PCI transport layer might be reported during FLR.
	Keywords: FLR, PCI transport layer, errors
	Discovered in Version: 14.26.1040
	Fixed in Release: 14.29.1016
2245422	Description: When MKEY_BY_NAME is enabled by NVCONFIG and a large number of VFs are configured, VM restart (VF/PF FLR) will take longer than when MKEY_BY_NAME is disabled
	Keywords: SR-IOV
	Discovered in Version: 14.28.1002
	Fixed in Release: 14.29.1016
2321713	Description: Fixed an issue that caused caused the device to go to dead IRISC as one of the firmware semaphores could not be released when a speed change or port state change was triggered.
	Keywords: IRISC, firmware semaphore,
	Discovered in Version: 14.28.1002
	Fixed in Release: 14.29.1016
1979562	Description: Fixed an issue that prevented the DHCP from assigning IPv6 address to the BMC during the initialization phase.
	Keywords: DHCP, IPv6 address, BMC
	Discovered in Version: 14.28.1002
	Fixed in Release: 14.28.2006
2080917	Description: Fixed and issue that resulted in driver startup failure when working in pass-through mode and dual port devices.
	Keywords: Pass-through mode, dual port devices
	Discovered in Version: 14.28.1002
	Fixed in Release: 14.28.2006
1911080	Description: Fixed a rare race condition that caused an erroneous write to the firmware image during certain power-down scenarios. This

Internal Ref.	Issue
	<p>resulted in firmware being recognized as corrupted and prevented the adapter card from being recognized by the system due to missing valid Flash images.</p> <p>Keywords: Flash images, firmware corruption</p> <p>Discovered in Version: 14.24.1000</p> <p>Fixed in Release: 14.28.1300</p>
1906153	<p>Description: Fixed an issue that caused the sent packet to hang while the device entered FLR mode.</p> <p>Keywords: Driver fails to load; FLR stuck; packet hang</p> <p>Discovered in Version: 14.27.1016</p> <p>Fixed in Release: 14.28.1002</p>
2108543	<p>Description: Enabled Bar configuration bitwise by applying the write_en bitmask.</p> <p>Keywords: Bitwise BAR Programming</p> <p>Discovered in Version: 14.27.1016</p> <p>Fixed in Release: 14.28.1002</p>
2126484	<p>Description: Fixed a rare case where the the device hanged while running the sw reset flow under heavy stress and with many open resources.</p> <p>Keywords: sw reset</p> <p>Discovered in Version: 14.27.1016</p> <p>Fixed in Release: 14.28.1002</p>
2119975	<p>Description: Fixed low PXE performance while using the VSC to trigger the send_ring_doorbells.</p> <p>Keywords: NODNIC, DOORBELL</p> <p>Discovered in Version: 14.27.1016</p> <p>Fixed in Release: 14.28.1002</p>
2135671	<p>Description: Fixed an error that prevented the completions (CQ) from being completed due to a race condition in the firmware transport error</p>

Internal Ref.	Issue
	handlers, and the error stressors, where the error stressors would hang the firmware transport error handler flow.
	Keywords: Error stressors; Transport flow; driver timeout
	Discovered in Version: 14.27.1016
	Fixed in Release: 14.28.1002
2119135	Description: Fixed an issue that cause fragmented IP packets to drop.
	Keywords: Fragmented IP packet
	Discovered in Version: 14.27.1016
	Fixed in Release: 14.28.1002
2107103	Description: Fixed an issue that prevented the desched_threshold field from working properly.
	Keywords: DCQCN
	Discovered in Version: 14.27.1016
	Fixed in Release: 14.28.1002
2165169	Description: Added the option to use the unicast MAC from the NC-SI cmd Set MAC Address to establish OS to BMC passthrough.
	Keywords: OS to BMC passthrough
	Discovered in Version: 14.27.1016
	Fixed in Release: 14.28.1002
211746	Description: Fixed a firmware fatal assert that showed an IRISC HANG due to init_hca waiting on the timers flow lock release.
	Keywords: Firmware assert
	Discovered in Version: 14.27.1016
	Fixed in Release: 14.28.1002

© Copyright 2024, NVIDIA. PDF Generated on 09/02/2024