NVIDIA ConnectX-6 Adapter Cards Firmware Release Notes v20.31.1014

Table of Contents

1 Release Notes Update History

Revision	Date	Description
20.31.1014	July 02, 2021	Initial release of this Release Notes version, This version introduces Changes and New Features and Bug Fixes.

2 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.

2.1 Firmware Download

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

2.2 Document Revision History

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

3 Firmware Compatible Products

The chapter contains the following sections:

These are the release notes for the NVIDIA® ConnectX®-6 adapters firmware Rev 20.31.1014. This firmware supports the following protocols:

- InfiniBand - SDR, QDR, FDR, EDR, HDR, HDR100
- Ethernet - 1GbE, 10GbE, 25GbE, 40GbE, 50GbE[1], 100GbE[1], 200GbE[2]
- PCI Express 4.0, supporting backwards compatibility for v3.0, v2.0 and v1.1

[1]. Speed that supports both NRZ and PAM4 modes in Force mode and Auto-Negotiation mode.

[2]. Speed that supports PAM4 mode only.

⚠ The minimal required NVIDIA Quantum firmware version is 27.2000.1260

⚠ Please make sure to use a PCIe slot that can supply the required power to the ConnectX-6 adapter card as stated in section Specifications in the adapter card's User Manual.

3.1 Supported Devices

This firmware supports the devices and protocols listed below:

Device Part Number	PSID	Device Name	FlexBoot	UEFI x86	UEFI ARM	Enable/ disable exprom Feature
MCX653106A-EFAT	MT_0000000219	ConnectX-6 VPI adapter card; 100Gb/s (HDR100; EDR IB and 100GbE); dual-port QSFP56; PCIe3.0/4.0 2x8 in a row	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX653105A-ECAT	MT_0000000222	ConnectX-6 VPI adapter card; 100Gb/s (HDR100; EDR IB and 100GbE); single-port QSFP56; PCIe3.0 x16; tall bracket; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX653105A-HDAT / MCX653105A-HDAL	MT_0000000223	ConnectX-6 VPI adapter card; HDR IB (200Gb/s) and 200GbE; single- port QSFP56; PCIe4.0 x16; tall bracket; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX653106A-ECAT	MT_0000000224	ConnectX-6 VPI adapter card; H100Gb/s (HDR100; EDR IB and 100GbE); dual-port QSFP56; PCIe3.0 x16; tall bracket; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX653106A-HDAT / MCX653106A-HDAL	MT_0000000225	ConnectX-6 VPI adapter card; HDR IB (200Gb/s) and 200GbE; dual-port QSFP56; PCIe4.0 x16; tall bracket; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists

Device Part Number	PSID	Device Name	FlexBoot	UEFI x86	UEFI ARM	Enable/ disable exprom Feature
MCX654105A-HCAT	MT_0000000226	ConnectX-6 VPI adapter card; HDR IB (200Gb/s) and 200GbE; single- port QSFP56; Socket Direct 2x PCIe3.0 x16; tall bracket; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX654106A-ECAT	MT_0000000227	ConnectX-6 VPI adapter card; 100Gb/s (HDR100; EDR InfiniBand and 100GbE); dual-port QSFP56; Socket Direct 2x PCIe3.0 x16; tall bracket; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX654106A-HCAT	MT_0000000228	ConnectX-6 VPI adapter card; HDR IB (200Gb/s) and 200GbE; dual-port QSFP56; Socket Direct 2x PCIe3.0 x16; tall bracket; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX653105A-EFAT	MT_0000000237	ConnectX-6 VPI adapter card; 100Gb/s (HDR100; EDR IB and 100GbE); single-port QSFP56; PCIe3.0/4.0 Socket Direct 2x8 in a row; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX651105A-EDAT	MT_0000000473	ConnectX®-6 VPI adapter card, 100Gb/s (HDR100, EDR IB and 100GbE, single-port QSFP56, PCIe4.0 x8, tall bracket	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX653435A-HDAI / MCX653435A-HDAE	MT_0000000296	ConnectX®-6 VPI adapter card, 200Gb/s (HDR IB and 200GbE) for OCP 3.0, with host management, Single-port QSFP56, PCIe4.0 x16, Internal Lock	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX653436A-HDAI / MCX653436A-HDAB	MT_0000000297	ConnectX®-6 VPI adapter card, 200Gb/s (HDR IB and 200GbE) for OCP 3.0, with host management, Dual-port QSFP56, PCIe4.0 x16, Internal Lock	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX653435A-EDAI	MT_0000000295	ConnectX®-6 VPI adapter card, 100Gb/s (HDR100, EDR IB and 100GbE) for OCP 3.0, with host management, Single-port QSFP56, PCIe 3.0/4.0 x16, Internal Lock	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX653435M-HDAI	MT_0000000601	ConnectX-6 VPI adapter card; 200Gb/s (HDR IB and 200GbE) for OCP 3.0; with host management; Single-port QSFP56; Multi Host or Socket Direct; PCIe4.0 x16; Internal Lock	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX613106A-VDAT	MT_0000000236	ConnectX-6 EN adapter card; 200GbE; dual-port QSFP56; PCIe4.0 x16; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists

Device Part Number	PSID	Device Name	FlexBoot	UEFI x86	UEFI ARM	Enable/ disable exprom Feature
MCX613436A-VDAI	MT_0000000294	ConnectX-6 EN adapter card; 200GbE for OCP 3.0; with host management; Dual-port QSFP56; PCIe 4.0 x16; Internal Lock	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX614106A-VCAT	MT_0000000221	ConnectX-6 EN adapter card; 200GbE; dual-port QSFP56; Socket Direct 2x PCIe3.0 x16; tall bracket; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX614106A-CCAT	MT_0000000220	ConnectX-6 EN adapter card; 100GbE; dual-port QSFP56; Socket Direct 2x PCIe3.0 x16; tall bracket; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX614105A-VCAT	MT_0000000284	ConnectX-6 EN adapter card kit; 200GbE; single-port QSFP56; Socket Direct 2x PCIe3.0 x16	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX613105A-VDAT	MT_0000000234	ConnectX-6 EN adapter card; 200GbE; single-port QSFP56; PCIe4.0 x16; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists
MCX613106A-CCAT	MT_0000000235	ConnectX-6 EN adapter card; 100GbE; dual-port QSFP56; PCIe3.0 x16; ROHS R6	Present (Enabled)	Present (Enabled)	Present (Enabled)	Exists

3.2 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
ConnectX-6 Firmware	20.30.1004 / 20.29.2002 /20.29.1016
MLNX_OFED	5.4-1.0.3.0 / 5.3-1.0.0.1 / 5.2-2.2.0.0 **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 / 5.2-2.2.0.0 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.
WinOF-2	2.70.50000 / 2.60.50000 / 2.50.50000 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.
MFT	4.17.0 / 4.16.3 / 4.16.0 **Note:** For the list of the supported Operating Systems, please refer to the driver's Release Notes.

	Supported Version
FlexBoot	3.6.403 **Note:** Please be aware that not all firmware binaries contain FlexBoot or UEFI, support may vary between cards. For further information see Supported Devices.
UEFI	14.24.13 **Note:** Please be aware that not all firmware binaries contain FlexBoot or UEFI, support may vary between cards. For further information see Supported Devices.
MLNX-OS	3.9.0900 onwards
NVIDIA Quantum™ Firmware	27.2008.2102 onwards
SwitchX-IB 2 Firmware	15.2008.2102 onwards
SwitchX-IB Firmware	11.2008.2102 onwards

3.3 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.

3.3.1 Switch and HCAs InfiniBand Cable Connectivity Matrix

NVIDIA Quantum™ based switches and NVIDIA® ConnectX®-6 HCAs support HDR (PAM4, 50Gb/s per lane) and EDR (NRZ, 25Gb/s per lane) technologies. As the ConnectX adapter cards are identified by their maximum supported throughput (e.g. ConnectX-6 VPI 100Gb/s card can support either 2-lanes of 50Gb/s or 4-lanes of 25Gb/s), the exact connectivity will be determined by the cable that is being used.

As a reference:

Speed Mode	Speed Supported	Number of Lanes Used
HDR	200Gb/s InfiniBand	4 lanes of 50Gb/s
HDR100	100Gb/s InfiniBand	2 lanes of 50Gb/s
EDR	100Gb/s InfiniBand	4 lanes of 25Gb/s
FDR	56Gb/s	4 lanes of 14Gb/s

The following tables present the connectivity matrix, between NVIDIA Quantum based switches, ConnectX-6 HCA, and the cables.

3.3.2 Switch-to-Switch Connectivity

Switch	Switch	Cable					
		H cable DAC	H cable AOC	HDR DAC	HDR AOC	EDR DAC/ AOC	FDR DAC/ AOC
NVIDIA Quantum™	NVIDIA Quantum	No such cable	HDR100	HDR	HDR	EDR	N/A
NVIDIA® Switch-IB®/ Switch-IB 2	NVIDIA Quantum	N/A	N/A	EDR	N/A	EDR	N/A
NVIDIA® SWITCHX®-2	NVIDIA Quantum	N/A	N/A	N/A	N/A	N/A	FDR
SWITCHX-2	Switch-IB/ Switch-IB 2	N/A	N/A	N/A	N/A	N/A	FDR

3.3.3 NIC-to-Switch Connectivity Matrix

Adapter	Switch		Cable					
			Y cable DAC/ AOC	HDR DAC	HDR AOC	EDR DAC	EDR AOC	FDR DAC/ AOC
ConnectX-6 200Gb/s	NVIDIA Quantum	HDR Switch	HDR100	HDR	HDR	EDR	EDR	N/A
ConnectX-6 100Gb/s	NVIDIA Quantum		HDR100	EDR	EDR	EDR	EDR	N/A
ConnectX-4/ ConnectX-5	NVIDIA Quantum		N/A	EDR	N/A	EDR	EDR	FDR
ConnectX-3/ ConnectX-3 Pro	NVIDIA Quantum		N/A	N/A	N/A	N/A	FDR[a]	FDR[a]
ConnectX-6	Switch-IB/ Switch-IB 2	EDR Switch	N/A	EDR	N/A	EDR	EDR	N/A
ConnectX-4/ ConnectX-5	Switch-IB/ Switch-IB 2		N/A	EDR	N/A	EDR	EDR	N/A
ConnectX-3/ ConnectX-3 Pro	Switch-IB/ Switch-IB 2		N/A	N/A	N/A	FDR	N/A	FDR
ConnectX-6	SWITCHX-2	FDR Switch	N/A	N/A	N/A	N/A	N/A	FDR
ConnectX-4/ ConnectX-5	SWITCHX-2		N/A	N/A	N/A	N/A	N/A	FDR
ConnectX-3/ ConnectX-3 Pro	SWITCHX-2		N/A	N/A	N/A	N/A	N/A	FDR

a. Connectivity between NVIDIA Quantum and ConnectX-3 and ConnectX-3 Pro is not supported when using ports #27-34.

3.3.3.1 VPI Protocol Support

ConnectX-6 VPI supports having one port as InfiniBand and the second port as Ethernet according to the following matrix of combinations.

> FDR is not supported in VPI mode.

This section provides details on the following tests:

To set the right configuration, run:

```
mlxconfig -d <mst device> s LINK_TYPE_P1=1/2 LINK_TYPE_P2=1/2
```

where:

- `LINK_TYPE_P1` - sets the configuring protocol for port 1
- `LINK_TYPE_P2` - sets the configuring protocol for port 2
- `(1/2)` - values used for the different protocols:
  - 1 - for InfiniBand
  - 2 - for Ethernet

Legend:

	Configuration Combination Support
V	Supported
X	Not supported

Port #1	InfiniBand
Port #2	Ethernet

	Port #2 - Ethernet							
	200GbE/50GbE		100GbE/25GbE		40GbE/10GbE		1GbE	
Port #1 - InfiniBand	#1	#2	#1	#2	#1	#2	#1	#2
HDR / HDR100	V	V	V	V	V	X	V	V
EDR	V	V	V	V	V	X	V	V
FDR*	X	V	X	V	X	X	X	V
QDR/SDR	V	V	V	V	V	X	V	V

* FDR is not supported in VPI mode.

Port #2	InfiniBand
Port #1	Ethernet

	Port #2 - InfiniBand							
	HDR/HDR100		EDR		FDR*		QDR	
Port #1 - Ethernet	#1	#2	#1	#2	#1	#2	#1	#2
200GbE/50GbE	V	V	V	X	V	X	V	V
100GbE/25GbE	V	V	V	X	V	X	V	V
40GbE/10GbE	V	V	V	X	V	X	V	V
1GbE	V	V	V	X	V	X	V	V

* FDR is not supported in VPI mode.

3.3.4 Validated and Supported FDR Cables

Speed	Cable OPN	Description
FDR	MC2207128-003	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 3m
FDR	MC2207130-002	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 2m
FDR	MC220731V-005	NVIDIA active fiber cable, VPI, up to 56Gb/s, QSFP, 5m
FDR	MC220731V-030	NVIDIA active fiber cable, VPI, up to 56Gb/s, QSFP, 30m

3.3.5 Validated and Supported EDR / 100Gb/s Cables

Speed	Cable OPN	Description
EDR	MCP1600-E001	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 1m 30AWG
EDR	MCP1600-E001E30	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 1m, Black, 30AWG
EDR	MCP1600-E002	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 2m 28AWG
EDR	MCP1600-E002E30	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 2m, Black, 30AWG
EDR	MCP1600-E003	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 3m 26AWG
EDR	MCP1600-E003E26	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 3m, Black, 26AWG
EDR	MCP1600-E004E26	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 4m, Black, 26AWG
EDR	MCP1600-E005E26	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 5m, Black, 26AWG

Speed	Cable OPN	Description
EDR	MCP1600-E00A	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 0.5m 30AWG
EDR	MCP1600-E00AE30	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 0.5m, Black, 30AWG
EDR	MCP1600-E00BE30	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 0.75m, Black, 30AWG
EDR	MCP1600-E01A	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 1.5m 30AWG
EDR	MCP1600-E01AE30	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 1.5m, Black, 30AWG
EDR	MCP1600-E01BE30	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 1.25m, Black, 30AWG
EDR	MCP1600-E02A	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 2.5m 26AWG
EDR	MCP1600-E02AE26	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 2.5m, Black, 26AWG
EDR	MFA1A00-E001	NVIDIA active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 1m
EDR	MFA1A00-E003	NVIDIA active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 3m
EDR	MFA1A00-E005	NVIDIA active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 5m
EDR	MFA1A00-E010	NVIDIA active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 10m
EDR	MFA1A00-E015	NVIDIA active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 15m
EDR	MFA1A00-E020	NVIDIA active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 20m
EDR	MFA1A00-E030	NVIDIA active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 30m
EDR	MFA1A00-E050	NVIDIA active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 50m
EDR	MFA1A00-E100	NVIDIA active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 100m
EDR	MMA1B00-E100	NVIDIA transceiver, IB EDR, up to 100Gb/s, QSFP28, MPO, 850nm, SR4, up to 100m
EDR	MFA1A00-E003-TG	NVIDIA customized active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 3m
EDR	MFA1A00-E005-TG	NVIDIA customized active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 5m
EDR	MFA1A00-E010-TG	NVIDIA customized active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 10m
EDR	MFA1A00-E015-TG	NVIDIA customized active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 15m
EDR	MFA1A00-E020-TG	NVIDIA customized active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 20m
EDR	MFA1A00-E030-TG	NVIDIA customized active fiber cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 30m
EDR	MMA1L30-CM	NVIDIA optical module, 100Gb/s, QSFP28, LC-LC, 1310nm, CWDM4, up to 2km
EDR	MMS1C10-CM	NVIDIA active optical module, 100Gb/s, QSFP, MPO, 1310nm, PSM4, up to 500m

⚠ EDR links raise with RS-FEC.

3.3.6 Validated and Supported HDR / 200Gb/s Cables

Speed	Cable OPN #	Description
HDR	MCP1650-H001E30	NVIDIA Passive Copper cable, IB HDR, up to 200Gb/s, QSFP28, PVC, 1m, white pultab, 30AWG
HDR	MCP1650-H002E26	NVIDIA Passive Copper cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, 2M, black pultab, 26AWG
HDR	MCP1650-H00AE30	NVIDIA Passive Copper cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, 0.5M, black pultab, 30AWG
HDR	MCP7H50-H001R30	NVIDIA Passive Copper Hybrid cable, IB HDR 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, colored pulltabs, 1m, 30AWG
HDR	MCP7H50-H01AR30	NVIDIA Passive Copper Hybrid cable, IB HDR 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, colored, 1.5m, 30AWG
HDR	MCP7H50-H002R26	NVIDIA Passive Copper Hybrid cable, IB HDR 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, colored, 2m, 26AWG
HDR	MFS1S00-H003E	NVIDIA Active Fiber cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, black pulltab, 3m
HDR	MFS1S00-H005E	NVIDIA Active Fiber cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, black pulltab, 5m
HDR	MFS1S00-H010E	NVIDIA Active Fiber cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, black pulltab, 10m
HDR	MFS1S00-H100E	NVIDIA Active Fiber cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, black pulltab, 100m
HDR	MFS1S50-H0xxE	NVIDIA Active Fiber Splitter cable, IB HDR, 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, up to 30m
HDR	MFS1S90-H003E	NVIDIA Active Fiber Splitter cable, IB HDR, 2x200Gb/s to 2x200Gb/s, 2xQSFP56 to 2xQSFP56, LSZH, 3m
HDR	MCA7J50-H003R*	NVIDIA Active Copper Hybrid cable, IB HDR 200Gb/s to 2xHDR100 100Gb/s, QSFP56 to 2xQSFP56, 3m, colored
HDR	MCA7J50-H004R*	NVIDIA Active Copper Hybrid cable, IB HDR 200Gb/s to 2xHDR100 100Gb/s, QSFP56 to 2xQSFP56, 4m, colored
HDR	MCA1J00-H003E*	NVIDIA Active Copper cable, IB HDR, up to 200Gb/s, QSFP56, 3m, yellow pulltab
HDR	MCA1J00-H004E*	NVIDIA Active Copper cable, IB HDR, up to 200Gb/s, QSFP56, 4m, yellow pulltab
HDR	MMA1T00-HS	NVIDIA transceiver, HDR, QSFP56, MPO, 850nm, SR4, up to 100m
HDR	MFS1S00-H130E	NVIDIA active fiber cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, black pulltab, 130m
HDR	MFS1S00-H003-LL	NVIDIA active fiber cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, black pulltab, low latency, 3m
HDR	MFS1S00-H005-LL	NVIDIA active fiber cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, black pulltab, low latency, 5m
HDR	MFS1S00-H010-LL	NVIDIA active fiber cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, black pulltab, low latency, 10m

Speed	Cable OPN #	Description
HDR	MFS1S00-H015-LL	NVIDIA active fiber cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, black pulltab, low latency, 15m
HDR	MFS1S00-H020-LL	NVIDIA active fiber cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, black pulltab, low latency, 20m
HDR	MFS1S00-H030-LL	NVIDIA active fiber cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, black pulltab, low latency, 30m

⚠ HDR links raise with RS_FEC.

⚠ *These cables were approved for switch-to-switch connectivity. For switch-to-host connectivity there may be some issues. See Known Issue 2073222/1959529 (see Known Issues)

3.3.7 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

3.3.8 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

Speed	Cable OPN	Description
10GE	MC3309124-004	NVIDIA passive copper cable, ETH 10GbE, 10Gb/s, SFP+, 4m
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
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

3.3.9 Validated and Supported 25GbE Cables

⚠ 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
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
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

Speed	Cable OPN	Description
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

3.3.10 Validated and Supported 40GbE Cables

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

Speed	Cable OPN	Description
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
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
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

3.3.11 Validated and Supported 56GbE Cables

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
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
56GE	MCP1700-F001D	NVIDIA passive copper cable, VPI, up to 56Gb/s, QSFP, 1m, Yellow Pulltab

Speed	Cable OPN	Description
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

3.3.12 Validated and Supported 100GbE Cables

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

Speed	Cable OPN	Description
100GbE	MCP1600-C01AE30N	NVIDIA Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 1.5m, Black, 30AWG, CA-N
100GbE	MCP1600-C02A	NVIDIA Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP, PVC, 2.5m 30AWG
100GbE	MCP1600-C02AE26N	NVIDIA Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 2.5m, Black, 26AWG, CA-N
100GbE	MCP1600-C02AE30L	NVIDIA Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28,2.5m, Black, 30AWG, CA-L
100GbE	MCP1600-C03A	NVIDIA Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP, PVC, 3.5m 26AWG
100GbE	MCP1600-E001	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 1m 30AWG
100GbE	MCP1600-E002	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 2m 28AWG
100GbE	MCP1600-E003	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 3m 26AWG
100GbE	MCP1600-E01A	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 1.5m 30AWG
100GbE	MCP1600-E02A	NVIDIA Passive Copper cable, IB EDR, up to 100Gb/s, QSFP, LSZH, 2.5m 26AWG
100GbE	MCP7F00-A001R	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, colored pulltabs, 1m, 30AWG
100GbE	MCP7F00-A001R30N	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 1m, Colored, 30AWG, CA-N
100GbE	MCP7F00-A002R	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, colored pulltabs, 2m, 30AWG
100GbE	MCP7F00-A002R30N	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 2m, Colored, 30AWG, CA-N
100GbE	MCP7F00-A003R26N	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 3m, Colored, 26AWG, CA-N
100GbE	MCP7F00-A003R30L	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 3m, Colored, 30AWG, CA-L
100GbE	MCP7F00-A005R26L	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 5m, Colored, 26AWG, CA-L
100GbE	MCP7F00-A01AR	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, colored pulltabs,1.5m, 30AWG
100GbE	MCP7F00-A01AR30N	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 1.5m, Colored, 30AWG, CA-N
100GbE	MCP7F00-A02AR26N	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 2.5m, Colored, 26AWG, CA-N
100GbE	MCP7F00-A02AR30L	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 2.5m, Colored, 30AWG, CA-L
100GbE	MCP7F00-A02ARLZ	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 2.5m, LSZH, Colored, 28AWG

Speed	Cable OPN	Description
100GbE	MCP7F00-A03AR26L	NVIDIA passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 3.5m, Colored, 26AWG, CA-L
100GbE	MCP7H00-G001	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 1m, 30AWG
100GbE	MCP7H00-G001R	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, colored pulltabs, 1m, 30AWG
100GbE	MCP7H00-G001R30N	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 1m, Colored, 30AWG, CA-N
100GbE	MCP7H00-G002R	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, colored pulltabs, 2m, 30AWG
100GbE	MCP7H00-G002R30N	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 2m, Colored, 30AWG, CA-N
100GbE	MCP7H00-G003R	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, colored pulltabs, 3m, 28AWG
100GbE	MCP7H00-G003R26N	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 3m, Colored, 26AWG, CA-N
100GbE	MCP7H00-G003R30L	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 3m, Colored, 30AWG, CA-L
100GbE	MCP7H00-G004R26L	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 4m, Colored, 26AWG, CA-L
100GbE	MCP7H00-G01AR	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, colored pulltabs, 1.5m, 30AWG
100GbE	MCP7H00-G01AR30N	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 1.5m, Colored, 30AWG, CA-N
100GbE	MCP7H00-G02AR	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, colored pulltabs, 2.5m, 30AWG
100GbE	MCP7H00-G02AR26N	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 2.5m, Colored, 26AWG, CA-N
100GbE	MCP7H00-G02AR30L	NVIDIA passive copper hybrid cable, ETH 100Gb/s to 2x50Gb/s, QSFP28 to 2xQSFP28, 2.5m, Colored, 30AWG, CA-L
100GbE	MFA1A00-C003	NVIDIA active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 3m
100GbE	MFA1A00-C005	NVIDIA active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 5m
100GbE	MFA1A00-C010	NVIDIA active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 10m
100GbE	MFA1A00-C015	NVIDIA active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 15m
100GbE	MFA1A00-C020	NVIDIA active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 20m
100GbE	MFA1A00-C030	NVIDIA active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 30m
100GbE	MFA1A00-C050	NVIDIA active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 50m
100GbE	MFA1A00-C100	NVIDIA active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 100m
100GbE	MFA7A20-C003	NVIDIA active fiber hybrid solution, ETH 100GbE to 2x50GbE, QSFP28 to 2xQSFP28, 3m
100GbE	MFA7A20-C005	NVIDIA active fiber hybrid solution, ETH 100GbE to 2x50GbE, QSFP28 to 2xQSFP28, 5m

Speed	Cable OPN	Description
100GbE	MFA7A20-C010	NVIDIA active fiber hybrid solution, ETH 100GbE to 2x50GbE, QSFP28 to 2xQSFP28, 10m
100GbE	MFA7A20-C020	NVIDIA active fiber hybrid solution, ETH 100GbE to 2x50GbE, QSFP28 to 2xQSFP28, 20m
100GbE	MFA7A50-C003	NVIDIA active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 3m
100GbE	MFA7A50-C005	NVIDIA active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 5m
100GbE	MFA7A50-C010	NVIDIA active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 10m
100GbE	MFA7A50-C015	NVIDIA active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 15m
100GbE	MFA7A50-C020	NVIDIA active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 20m
100GbE	MFA7A50-C030	NVIDIA active fiber hybrid solution, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 30m
100GbE	MMA1B00-C100D	NVIDIA transceiver, 100GbE, QSFP28, MPO, 850nm, SR4, up to 100m, DDMI
100GbE	MMA1L10-CR	NVIDIA optical transceiver, 100GbE, QSFP28, LC-LC, 1310nm, LR4 up to 10km **Note:** Only revision A2 and above.
100GbE	MFA1A00-C001-TG	NVIDIA customized active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 1m
100GbE	MFA1A00-C002-TG	NVIDIA customized active fiber cable, ETH 100GbE, 100Gb/s, QSFP28, LSZH, 2m
100GbE	MFA1A00-C003-TG	NVIDIA customized active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 3m
100GbE	MFA1A00-C005-TG	NVIDIA customized active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 5m
100GbE	MFA1A00-C007-TG	NVIDIA customized active fiber cable, ETH 100GbE, 100Gb/s, QSFP28, LSZH, 7m
100GbE	MFA1A00-C010-TG	NVIDIA customized active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 10m
100GbE	MFA1A00-C015-TG	NVIDIA customized active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 15m
100GbE	MFA1A00-C020-TG	NVIDIA customized active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 20m
100GbE	MFA1A00-C030-TG	NVIDIA customized active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 30m
100GbE	MFA1A00-C050-TG	NVIDIA customized active fiber cable, ETH 100GbE, 100Gb/s, QSFP, LSZH, 50m
100GbE	MMA1L30-CM	NVIDIA optical module, 100GbE, 100Gb/s, QSFP28, LC-LC, 1310nm, CWDM4, up to 2km
100GbE	MMS1C10-CM	NVIDIA active optical module, 100Gb/s, QSFP, MPO, 1310nm, PSM4, up to 500m

3.3.13 Validated and Supported 200GbE Cables

Speed	Cable OPN	Description
200GE	MCP1650-V001E30	NVIDIA Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 1m, black pulltab, 30AWG
200GE	MCP1650-V002E26	NVIDIA Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 2m, black pulltab, 26AWG
200GE	MCP1650-V002E26_FF	NVIDIA Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 2m, black pulltab, 26AWG
200GE	MCP1650-V003E26	NVIDIA Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 3m, black pulltab, 26AWG
200GE	MCP1650-V00AE30	NVIDIA Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 0.5m, black pulltab, 30AWG
200GE	MCP1650-V01AE30	NVIDIA Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 1.5m, black pulltab, 30AWG
200GE	MCP1650-V02AE26	NVIDIA Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 2.5m, black pulltab, 26AWG
200GE	MCP7H50-V001R30	NVIDIA passive copper hybrid cable, 200GbE 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, colored, 1m, 30AWG
200GE	MCP7H50-V002R26	NVIDIA passive copper hybrid cable, 200GbE 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, colored, 2m, 26AWG
200GE	MCP7H50-V003R26	NVIDIA passive copper hybrid cable, 200GbE 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, colored, 3m, 26AWG
200GE	MCP7H50-V01AR30	NVIDIA passive copper hybrid cable, 200GbE 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, colored, 1.5m, 30AWG
200GE	MCP7H50-V02AR26	NVIDIA passive copper hybrid cable, 200GbE 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, colored, 2.5m, 26AWG
200GE	MCP7H70-V001R30	NVIDIA passive copper hybrid cable, 200GbE 200Gb/s to 4x50Gb/s, QSFP56 to 4xSFP56, colored, 1m, 30AWG
200GE	MCP7H70-V002R26	NVIDIA passive copper hybrid cable, 200GbE 200Gb/s to 4x50Gb/s, QSFP56 to 4xSFP56, colored, 2m, 26AWG
200GE	MCP7H70-V003R26	NVIDIA passive copper hybrid cable, 200GbE 200Gb/s to 4x50Gb/s, QSFP56 to 4x4SFP56, colored, 3m, 26AWG
200GE	MCP7H70-V01AR30	NVIDIA passive copper hybrid cable, 200GbE 200Gb/s to 4x50Gb/s, QSFP56 to 4xSFP56, colored, 1.5m, 30AWG
200GE	MCP7H70-V02AR26	NVIDIA passive copper hybrid cable, 200GbE 200Gb/s to4x50Gb/s, QSFP56 to 4xSFP56, colored, 2.5m, 26AWG
200GE	MFS1S00-V003E	NVIDIA active fiber cable, 200GbE, 200Gb/s, QSFP56, LSZH, black pulltab, 3m

Speed	Cable OPN	Description
200GE	MFS1S00-V005E	NVIDIA active fiber cable, 200GbE, 200Gb/s, QSFP56, LSZH, black pulltab, 5m
200GE	MFS1S00-V010E	NVIDIA active fiber cable, 200GbE, 200Gb/s, QSFP56, LSZH, black pulltab, 10m
200GE	MFS1S00-V015E	NVIDIA active fiber cable, 200GbE, 200Gb/s, QSFP56, LSZH, black pulltab, 15m
200GE	MFS1S00-V020E	NVIDIA active fiber cable, 200GbE, 200Gb/s, QSFP56, LSZH, black pulltab, 20m
200GE	MFS1S00-V030E	NVIDIA active fiber cable, 200GbE, 200Gb/s, QSFP56, LSZH, black pulltab, 30m
200GE	MFS1S00-V050E	NVIDIA active fiber cable, 200GbE, 200Gb/s, QSFP56, LSZH, black pulltab, 50m
200GE	MFS1S00-V100E	NVIDIA active fiber cable, 200GbE, 200Gb/s, QSFP56, LSZH, black pulltab, 100m
200GE	MCP1650-V00AE30	NVIDIA Passive Copper cable, 200GbE, 200Gb/s, QSFP56, LSZH, 0.5m, black pulltab, 30AWG
200GE	MMA1T00-VS	NVIDIA transceiver, 200GbE, up to 200Gb/s, QSFP56, MPO, 850nm, SR4, up to 100m
200GE	MFS1S50-V003E	NVIDIA active fiber splitter cable, 200GbE, 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, black pulltab, 3m
200GE	MFS1S50-V005E	NVIDIA active fiber splitter cable, 200GbE, 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, black pulltab, 5m
200GE	MFS1S50-V010E	NVIDIA active fiber splitter cable, 200GbE, 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, black pulltab, 10m
200GE	MFS1S50-V015E	NVIDIA active fiber splitter cable, 200GbE, 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, black pulltab, 15m
200GE	MFS1S50-V020E	NVIDIA active fiber splitter cable, 200GbE, 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, black pulltab, 20m
200GE	MFS1S50-V030E	NVIDIA active fiber splitter cable, 200GbE, 200Gb/s to 2x100Gb/s, QSFP56 to 2xQSFP56, LSZH, black pulltab, 30m

3.4 Supported 3rd Party Cables and Modules

Speed	Cable OPN	Description
10GbE	FTLX8571D3BCL-ME	10gb SFP 850nm Optic Transceiver
10GbE	SP7051-HP	HP-MethodeElec. 10GbE AOM
40GbE	2231254-2	Cisco 3m 40GbE copper

Speed	Cable OPN	Description
40GbE	AFBR-7QER15Z-CS1	Cisco 40GbE 15m AOC
40GbE	BN-QS-SP-CBL-5M	PASSIVE COPPER SPLITTER CABLE ETH 40GBE TO 4X10GBE 5M
40GbE	NDCCGJ-C402	15m (49ft) Avago AFBR-7QER15Z Compatible 40G QSFP+ Active Optical Cable
40GbE	QSFP-40G-SR-BD	Cisco 40GBASE-SR-BiDi, duplex MMF
100GbE	1AT-3Q4M01XX-12A	O-NET QSFP28 100G Active cable/module
100GbE	AQPMANQ4EDMA0784	QSFP28 100G SMF 500m Transceiver
100GbE	CAB-Q-Q-100G-3M	Passive 3 meter, QSFP+ to QSFP+ QSFP100 TWINAX 103.125Gbps-CR4
100GbE	CAB-Q-Q-100GbE-3M	Passive 3 meter , QSFP+ to QSFP+ QSFP100 TWINAX 103.125Gbps-CR4
100GbE	FCBN425QE1C30-C1	100GbE Quadwire® QSFP28 Active Optical Cable 30M
100GbE	FTLC1151RDPL	TRANSCIEVER 100GBE QSFP LR4
100GbE	FTLC9152RGPL	100G 100M QSFP28 SWDM4 OPT TRANS
100GbE	FTLC9555REPM3-E6	100m Parallel MMF 100GQSFP28Optical Transceiver
100GbE	NDAAFJ-C102	SF-NDAAFJ100G-005M
100GbE	QSFP-100G-AOC30M	30m (98ft) Cisco QSFP-100G-AOC30M Compatible 100G QSFP28 Active Optical Cable
100GbE	QSFP28-LR4-AJ	CISCO-PRE 100GbE LR4 QSFP28 Transceiver Module
100GbE	SFBR-89BDDZ-CS2	CISCO-PRE 100G AOM BiDi
100GbE	SQF1002L4LNC101P	Cisco-SUMITOMO 100GbE AOM

3.5 Tested Switches

3.5.1 Tested HDR / 200Gb/s Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
HDR	Quantum	MQM8700-xxx	40-port Managed Non-blocking HDR 200Gb/s InfiniBand Smart Switch	NVIDIA
HDR	Quantum	MQM8790-xxx	40-port Unmanaged, Non-blocking HDR 200Gb/s InfiniBand Smart Switch	NVIDIA

3.5.2 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

3.5.3 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
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

3.5.4 Tested 200GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
200GbE	Spectrum	MSN3700-XXXX	32 QSFP56 ports, 200GbE Open Ethernet Switch System	Mellanox

3.6 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.

4 Changes and New Features

4.1 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 the configuration on multi-host adapter cards, for the changes to take effect, PCIe restart must be simultaneously sent 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.

⚠ It is recommended to enable the "above 4G decoding" BIOS setting for features that require a large amount of PCIe resources (e.g., SR-IOV with numerous VFs, PCIe Emulated Switch, Large BAR Requests).

4.2 Changes and New Feature in this Firmware Version

Feature/Change	Description
20.31.1014	
Using NC-SI Commands for Debugging PCI Link Failures	Implemented a new NC-SI command `get_debug_info` 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.
NC-SI Speed Reporting	Updated the NC-SI speed reporting output to support 200GbE speed. Now when running the NC-SI command, the output presents 200GbE speed as well.

Feature/Change	Description
20.31.1014	
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 (NUM_VF_MSIX) affects the configured number of VFs (NUM_OF_VFS). 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 (MAX_TOTAL_MSIX).
Asymmetrical MSIX Configuration	This feature allows the device to be configured with a different number of MSIX vectors per physical PCI functions. To use this feature, please follow these steps: 1. Clear NUM_PF_MSIX_VALID to disable global symmetrical MSIX configuration. 2. Set PF_NUM_PF_MSIX_VALID to enable asymmetrical per Physical Function MSIX configuration. 3. Configure PF_NUM_PF_MSIX **per physical PCI function**. **Notes:** • When using this configuration, NUM_PF_MSIX (the symmetrical MSIX configuration field) will be ignored by the firmware although it will be present in the mlxconfig query. • The asymmetrical MSIX configuration and the legacy symmetrical MSIX configuration are mutually exclusive. When both are enabled, the firmware will work according to the symmetrical MSIX configuration. • Step #3 should be done on each of the physical PCI functions, otherwise, the device will use default configurations.
RDMA, NC-SI	Added support for RDMA partitioning and RDMA counters in IB mode.
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 funcion 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 firmare.
PCIe Eye Opening	Enabled measuring PCIe eye dynamic grading over PCIe Gen3 speed.
User Memory (UMEM)	Enabled UID 0 to create resources with UMEM.
Native IB Packets	Added support for receiving and sending native IB packets from/to the software (including all headers) via raw IBL2 QPs.
InfiniBand Packet Steering	Added support for RX RDMA NIC flow table on an IB port. Now the software can steer native IB packets to raw IB receive queues according to the DLID and the DQPN.
Steering	Added support for matching field `ipv4_ihl` in `create_flow_group` and `set_flow_table_entry` commands.
Bug Fixes	See *Bug Fixes in this Firmware Version* section.

4.3 Unsupported Features and Commands

4.3.1 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

4.3.2 Unsupported Commands

- QUERY_MAD_DEMUX
- SET_MAD_DEMUX
- CREATE_RQ - MEMORY_RQ_RMP
- MODIFY_LAG_ASYNC_EVENT

5 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
2648336	**Description:** Disabled the CNP counter “rp_cnp_ignored " (triggered by OOS (out-of-sequence)) when all ports are IB. **Note:** For mixed IB/ETH scenario, the behavior depends on the RoCE configuration, the counter on the IB port may still increase but will not affect the regular use.
	Keywords: CNP counter, IB
	Discovered in Version: 20.30.1004
	Fixed in Release: 20.31.1014
2178949	**Description:** Improved PortXmitWait IB counter accuracy.
	Keywords: Counters
	Discovered in Version: 20.30.1004
	Fixed in Release: 20.31.1014
2641734	**Description:** Fixed the rate select mechanism in QSFP modules.
	Keywords: Cables
	Discovered in Version: 20.30.1004
	Fixed in Release: 20.31.1014
2600783	**Description:** Fixed classification issues for "Passive" cables to be more robust.
	Keywords: Cables
	Discovered in Version: 20.30.1004
	Fixed in Release: 20.31.1014
2574322	**Description:** Fixed an issue that occasionally caused some performance issues related to RC QPs using E2E-credits (not connected to SRQ and doing send/receive traffic) when the ROCE_ACCL tx_window was enabled.
	Keywords: Bandwidth, performance
	Discovered in Version: 20.30.1004
	Fixed in Release: 20.31.1014
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 receive WQE that was being scattered.
	Keywords: Strided RQ, MTU
	Discovered in Version: 20.30.1004
	Fixed in Release: 20.31.1014
2569999	**Description:** Fixed a rare issue that caused RX pipe to hang.
	Keywords: RX pipe
	Discovered in Version: 20.30.1004
	Fixed in Release: 20.31.1014

Internal Ref.	Issue
2621704	**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.
	Keywords: Resource number size, free_4k page
	Discovered in Version: 20.30.1004
	Fixed in Release: 20.31.1014
2619161	**Description:** Initialized the rate table in the static configuration so it will be configured at the link-not-up scenarios.
	Keywords: RoCE, static configuration, rate table
	Discovered in Version: 20.30.1004
	Fixed in Release: 20.31.1014
2589430	**Description:** CRT_DCR with index larger than 1 << 21 can collide with the CRT_SW_RESERVED address.
	Keywords: DCR
	Discovered in Version: 20.30.1004
	Fixed in Release: 20.31.1014
2447160	**Description:** In InfiniBand non-virtualization system, due to a corrupted steering root, traffic fails after a warm reboot.
	Keywords: Steering, Traffic
	Discovered in Version: 20.30.1004
	Fixed in Release: 20.31.1014

6 Known Issues

For a list of older versions' Known Issues that are not listed in this chapter, please refer to the relevant firmware versions Release Notes in https://docs.mellanox.com/category/adapterfw.

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
2684071	**Description:** Changing the default host chaining buffer size or WQE size (HOST_CHAINING_DESCRIPTORS, HOST_CHAINING_TOTAL_BUFFER_SIZE) using NVconfig might result in driver initialization failure.
	Workaround: N/A
	Keywords: Host chaining
	Discovered in Version: 20.29.2002
2582094	**Description:** When performing a stress toggle test vs. IXIA, the IXIA side is not ready for few seconds.
	Workaround: Wait for 1 sec between running the down and up commands.
	Keywords: Auto-negotiation.
	Discovered in Version: 20.31.1014
2446583	**Description:** On rare occasions, when both network devices are NVIDIA, PAM4 link will raise with several effective errors. These errors will not affect traffic once the link is up.
	Workaround: Clear counters once the link is up
	Keywords: Effective errors
	Discovered in Version: 20.29.2002

Internal Ref.	Issue
2411542	**Description:** Multi-APP QoS is not supported when LAG is configured.
	Workaround: N/A
	Keywords: Multi-APP QoS, LAG
	Discovered in Version: 20.29.1016
SF 933911	**Description:** PXE boot will not function if the adapter card is connected to a NVIDIA Quantum™ based switch over an HDR active optical cable.
	Workaround: Set KEEP_LINK_UP_ON_BOOT configuration to enable via mlxconfig. For further information please contact Support.
	Keywords: IB, Link Speed , Link Down, SDR , optical cable
	Discovered in Version: 20.30.1004
2378593	**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.
	Workaround: Use full reset flow for firmware upgrade/downgrade.
	Keywords: Sub 1sec firmware update
	Discovered in Version: 20.29.1016
2396506	**Description:** On systems with high PCIe latency (2us or above), lower bandwidth may be experienced.
	Workaround: If such issue is observed: 1. Enable ZTT to overcome the high latency. Run: `mlxconfig -d <mst device> set ZERO_TOUCH_TUNING_ENABLE=1` 2. Reset or power cycle the firmware for change to take effect
	Keywords: Performance, ZTT
	Discovered in Version: 20.29.1016

2213356	**Description:** The following are the Steering Dump limitations:
	• 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:
	  ◦ 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

	Workaround: N/A
	Keywords: Steering Bump
	Discovered in Version: 20.29.1016
2365322	**Description:** When configuring adapter card's Level Scheduling, a QoS tree leaf (QUEUE_GROUP) configured with default rate_limit and default bw_share, may not obey the QoS restrictions imposed by any of the leaf's ancestors.
	Workaround: To prevent such a case, configure at least one of the following QoS attributes of a leaf: `max_average_bw` or `bw_share`

Internal Ref.	Issue
	Keywords: QoS
	Discovered in Version: 20.29.1016
2330700	**Description:** Effective BER may be observed when connecting to SN3700 switch system which can impact the link up time to be up to 97sec.
	Workaround: N/A
	Keywords: Effective BER
	Discovered in Version: 20.28.4000
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.
	Workaround: N/A
	Keywords: SR-IOV
	Discovered in Version: 20.28.1002
2239632	**Description:** EDR linkup time might take up to 50sec when using HDR optical cable.
	Workaround: N/A
	Keywords: Cables, EDR
	Discovered in Version: 20.28.1002
2199939	**Description:** High linkup time may be experienced when connecting to an H3C switch using 25GbE\50GbE\100GbE link speeds.
	Workaround: N/A
	Keywords: Linkup time, switch
	Discovered in Version: 20.28.1002
2145881	**Description:** FDR link is unstable when using an FDR cable in ports: #27-#34.
	Workaround: N/A
	Keywords: FDR, cables
	Discovered in Version: 20.27.6008
2149437	**Description:** When the SLTP configuration is wrongly set, the “Bad status” explanation will not be presented (only error indication) to the user.
	Workaround: N/A
	Keywords: SLTP configuration
	Discovered in Version: 20.27.6008
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: 20.27.1016
1796936	**Description:** 200GbE Optical cables in Auto-Negotiation mode work only in 200GbE speed.
	Workaround: N/A
	Keywords: Cables

Internal Ref.	Issue
	Discovered in Version: 20.27.1016
1959529	**Description:** When HDR Active Copper cables are used between Quantum switches, or between Quantum switch and ConnectX-6 HCA, the counter indicating 'Link Down' may have a value other than zero, after the first time the cable is connected. As this may happened only at the first time, it is recommend to clear the counters after the cluster is brought up.
	Workaround: Toggle the Active Copper or Optics cables as the switch performs a reset.
	Keywords: Cables, BER
	Discovered in Version: 20.27.1016
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: 20.27.1016
1959529	**Description:** Occasionally (up to 15% of connections), the link will go down when using ACC cables P/N: MCA1J00-H003E, MCA1J00-H004E and when connecting a Quantum switch to a Quantum switch.
	Workaround: N/A
	Keywords: Cables
	Discovered in Version: 20.27.1016
1997329	**Description:** Downgrading from firmware v20.26.4012 to firmware v20.26.1040 and lower is not supported on Windows OSes using the mlxfwmanager tool.
	Workaround: N/A
	Keywords: mlxfwmanager, firmware downgrade
	Discovered in Version: 20.26.4012
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: 20.26.1040
-	**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: 20.26.1040
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

Internal Ref.	Issue
	Discovered in Version: 20.26.1040
1796936	**Description:** HDR split cables support only HDR speed.
	Workaround: N/A
	Keywords: Link Speed, cables, Break-Out cables
	Discovered in Version: 20.26.1040
1911160	**Description:** When in loopback mode, the link is not raised when using Cisco 10GbE AOM.
	Workaround: N/A
	Keywords: Loopback mode, link up
	Discovered in Version: 20.26.1040
1733559	**Description:** The effective BER of ~ 1E-7 is expected when using ConnectX-6 adapter cards in 50GbE (PAM4) link speed and connecting to a Spectrum-2 SN3700 switch systems using copper split cable (100-->2x50).
	Workaround: N/A
	Keywords: BER, 50GbE, Spectrum-2
	Discovered in Version: 20.26.1040
1750460 / 2063991	**Description:** BER issues might occur when using ConnectX-6 adapter cards in 100GbE link speed, and connecting with and 3rd party switch systems.
	Workaround: N/A
	Keywords: BER, 100GbE, Spectrum-2
	Discovered in Version: 20.26.1040
1906389	**Description:** When using 100GbE link speed and connecting to a Cisco9000 switch, the link might take up to 2 min to raise.
	Workaround: N/A
	Keywords: Link speed
	Discovered in Version: 20.26.1040
1918749	**Description:** mlxlink tool displays a wrong speed when using ETH cables on ConnectX-6 adapter cards.
	Workaround: N/A
	Keywords: mlxlink
	Discovered in Version: 20.26.1040
1901198	**Description:** Firmware is not loaded on Multi-Host setups after reboot.
	Workaround: N/A
	Keywords: Firmware load, Multi-Host
	Discovered in Version: 20.26.1040
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

Internal Ref.	Issue
	Discovered in Version: 20.26.1040
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: 20.26.1040
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: 20.26.1040
-	**Description:** Coherent Accelerator Processor Interface (CAPI) in ConnectX-6 firmware v20.25.7020 and above has low test coverage, however, it has no known issues.
	Workaround: N/A
	Keywords: CAPI
	Discovered in Version: 20.25.7020
1563590	**Description:** LR4 modules are currently not supported.
	Workaround: N/A
	Keywords: Modules/Cables
	Discovered in Version: 20.25.6000
-	**Description:** HDR optical cables and Split cables support only HDR speed.
	Workaround: N/A
	Keywords: Link Speed, cables, Break-Out cables
	Discovered in Version: 20.25.6000
1755286	**Description:** Port speed may change to SDR spontaneously, without a clear reason.
	Workaround: Keep the "keep_ib_link_up" bit at 0 in NVconfig to make sure the port is raised with the correct speed.
	Keywords: SDR, port speed
	Discovered in Version: 20.25.2006
1778616	**Description:** If the flash memory is not cleared, link_maintenance can be wrongly disabled by the NV configuration.
	Workaround: N/A
	Keywords: Flash memory
	Discovered in Version: 20.25.2006
1774135	**Description:** PXE boot is not functional when connecting a splitter cable to the host.

Internal Ref.	Issue

	Workaround: Update the SM as follow:
	• MLNX_OFED SM:
	  • Set the default partition in the SM partitions.conf file as shown in the example below:
	  Default=0x7fff,ipoib,rate=5:ALL=full;
	  Note: "rate" must be set to "5" regardless to the other flags values.
	• MLNX-OS SM:
	  Run the following CLI commands:
	  `no ib sm`
	  `ib partition Default rate 5 ib sm`
	• UFM SM:
	  Use REST API to change default partition rate:
	  PUT https://<some IP>/ufmRest/resources/networks/management
	  `{`
	  `"qos_parameters": { "rate_limit": 900`
	  `}`
	  `}`
	  As a result, /opt/ufm/files/conf/opensm/partitions.conf will include the following line:
	  management=0x7fff,ipoib, sl=0,rate=5, defmember=full : ALL, ALL_SWITCHES=full,SELF=full;

	Keywords: PXE boot, splitter cable
	Discovered in Version: 20.25.2006
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: 20.25.2006
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: 20.25.2006
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: 20.25.2006
1774890	**Description:** If ConnectX-6 adapter card is connected to a Quantum based switch over an HDR fiber optic cable, the SDR speed will not function. This limitation causes PXE not to function when performing: • PXE boot, as PXE traffic utilizes SDR speed • PXE boot over IB, when using either an optic cable or a copper splitter
	Workaround: N/A

Internal Ref.	Issue
	Keywords: PXE, Quantum, SDR, EDR, HDR, cables
	Discovered in Version: 20.25.2006
1716334	**Description:** When mlxconfig.PF_BAR2_EN is enabled, configuring more than 255 PCI functions will raise an assert.
	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: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
-	**Description:** The supported length of HDR copper cables is currently up to 2M.
	Workaround: N/A
	Keywords: HDR cables
	Discovered in Version: 20.25.1500 **[Beta]**
-	**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: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**

Internal Ref.	Issue
-	**Description:** In some cases, the power consumption might be 10% higher than what is stated in the adapter cards User Manual.
	Workaround: Power consumption will be aligned with the User Manual statement in the next release
	Keywords: Power consumption
	Discovered in Version: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
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
	Discovered in Version: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
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

Internal Ref.	Issue
	Discovered in Version: 20.25.1500 **[Beta]**
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: .20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**
1171013	**Description:** Signature Handover Operations is not supported when FPP (Function-Per-Port) mode is disabled.
	Workaround: N/A
	Keywords: Signature Handover Operations, FPP
	Discovered in Version: 20.25.1500 **[Beta]**
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: 20.25.1500 **[Beta]**

7 PreBoot Drivers (FlexBoot/UEFI)

7.1 FlexBoot Changes and New Features

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

7.2 UEFI Changes and Major New Features

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

8 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_GLOBAL_MASK	ece_disable_mask		0x116
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		
	log_max_outstandng_wqe		
	NV_config.sr_enable (ConnectX-6 Dx and above)		
NV_IB_DC_CONF	LOG_DCR_HASH_TABLE_SIZE		0x190

Configuration	mlxconfig Parameter Name	Class	TLV ID
	DCR_LIFO_SIZE		
NV_VPI_LINK_TYPE	LINK_TYPE	PHYSICAL_PORT (2)	0x12
NV_ROCE_CC	ROCE_CC_PRIO_MASK		0x107
	ROCE_CC_ALGORITHM		
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

Configuration	mlxconfig Parameter Name	Class	TLV ID
	NUM_OF_TC		
	NUM_OF_PFC		
NV_MPFS_CONF	DUP_MAC_ACTION		0x196
	SRIOV_IB_ROUTING_MODE		
	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		

9 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
20.30.1004	
PAM4	Added support for PAM4 Auto Negotiation and Link Training in 200GbE link speed.
RoCE, Lossy, slow_restart_idle	Removed triggering unexpected internal CNPs for RoCE Lossy slow_restart_idle feature.
KR-Startup in Auto-Negotiation	Enabled KR-Startup in Auto-Negotiation mode for PAM4.
Performance: Steering	Added support for a new NV config mode `“icm_cache_mode_large_scale_steering”` t hat 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"`.
Active-State Power Management (ASPM)	Added support for power saving in L1 ASPM link state.
VF/VF-group rate-limiting	This new capability enables VF/VF-group rate-limiting while per-host rate-limiter is also applied.
Bug Fixes	See Bug Fixes History section.
20.29.2002	
Remote Loopback in NRZ	**[Beta]** Enabled remote (PMA) loopback in NRZ, Rx-to-Tx. **Note:** To use the PMA loopback, both sides should be in Force Mode (AN Disabled).
PAM4	PAM4 200GbE linkup time improvement, the linkup time is now sub 5 seconds.
Reserved QPN	**[Beta]** This capability allows the software to reserve a QPN that can be used to establish connection performed over RDMA_CM, and provide the software a unique QP number. Since RDMA_CM does not support DC, by using CREATE_QPN_RESERVED_OBJECT the software can reserve a QPN value from the firmware's managed QP number namespace range. This allows multiple software processes to hold a unique QPN value instead of using UD-QPs.
Bug Fixes	See Bug Fixes History section.
20.29.1016	
Cable Firmware Burning	**[Beta]** Added support for LinkX module burning via MFT toolset. The new capability enables direct firmware burning from the internal flash storage to reduce the bandwidth and accelerate the burning process, including burning several modules at a time.
Eye-Opening	**[Beta]** Eye-opening is supported only when using NRZ signal.
Multi-Application QoS per QP	Added the option to allow applications to build their own QoS tree over the NIC hierarchy by connecting QPs to responder/requestor Queue Groups.
NRZ Link Performance	Improved NRZ link performance (RX algorithm).

NRZ Link-Up Time	Improved NRZ link-up time (25G\50G\100G speeds).
Tx Sets	Enabled the options to control different Tx sets for the same attribute when connecting a Mellanox-Mellanox vs Mellanox to 3rd party HCA.
InfiniBand Support in RDE	Added "InfiniBand" properties set to the Network Device Function Redfish object.
Direct Packet Placement (DPP)	Added support for Direct Packet Placement (DPP). DPP is a receive side transport service in which the Ethernet packets are scattered to the memory according to a packet sequence number (PSN) carried by the packet, and not by their arrival order. To enable DPP offload, the software should create a special RQ by using the CREATE_RQ command, and set DPP relevant attributes.
HW Offloads Enablement on VF	Added trust level for VFs. Once the VF is trusted, it will get a set of trusted capabilities.
Enabling Adaptive-Routing (AR) for the Right SL via UCX	UCX can now enable AR by exposing Out-Of-Ordering bitmask per SL with `"ooo_per_sl"` field in the HCA_VPORT context. It can be also queried by running the `QUERY_HCA_VPORT_CONTEXT` command.
InfiniBand Congestion Control	Enhanced IB Congestion Control to support lower minimum rate. Now it uses destination-lid to classify flows to handle larger scale, and achieve better results in GPCNeT benchmark.
Steering Dump	Hardware steering dump output used for debugging and troubleshooting. Please see Known Issue 2213356 for its limitations.
20.28.4000	
PAM4	PAM4 link performance improvement.
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”.
Access Registries	`ignore_flow_level` is now enabled by the TRUST LEVEL access registry.
Pause Frames from VFs	**[Beta]** Enabled the capability to allow Virtual Functions to send Pause Frames packets.
Counters	Added support for the `cq_overrun` 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.
Bug Fixes	See Bug Fixes.
20.28.2006	
Sub Function (SF) BAR Size	Increased the minimum Sub Function (SF) BAR size from 128KB to 256KB. Due to the larger SF BAR size, for the same PF BAR2 size, which can be queried/modified by LOG_PF_BAR2_SIZE NV config, the firmware will support half of the SFs. To maintain the same amount of supported SFs, software needs to increase the LOG_PF_BAR2_SIZE NV config value by 1.
AES-XTS	AES_XTS is used to perform all disk encryption/decryption related flows in the NIC and reduce cost and overheads of the related FIPS certification.
GPUDirect in Virtualized Environment	Enabled a direct access to ATS from the NIC to GPU buffers using PCIe peer-to-peer transactions. To enable this capability, the “p2p_ordering_mode” parameter was added to the NV_PCI_CONF configuration. **Note:** When SECURE_ALL or SECURE_TRUST is configured, ATS and RO must be set identically. When SECURE_NONE is configured, ATS and RO may be set independently as the current firmware behavior allows.

Non-Volatile Configurations	Added a new Non-Volatile Configuration parameter to control VL15 buffer size (VL15_BUFFER_SIZE). **Note:** VL15 buffer size enlargement will decrease all other VLs buffers size.
NC-SI	Added a new NC-SI command (get_device_id) to report a unique device identifier.
NC-SI	Added new NC-SI commands (`get_lldp_nb, set_lldp_nb`) to query the current status of LLDP and to enable/disable it.
ROCE ACCL	Split the SlowRestart ROCE_ACCL into the following: • `slow-restart` - used to reduce rate on retransmission events • `slow-restart-after-idle` - used to reduce rate before first transmission after >1s without transmitting
ROCE ACCL	Enabled TX PSN window size configuration using LOG_TX_PSN_WINDOW NVconfig parameter. **Note:** Due to hardware limitations, `max log_tx_psn_win` value can be set 9.
Bug Fixes	See Bug Fixes.
20.28.2006	
Sub Function (SF) BAR Size	Increased the minimum Sub Function (SF) BAR size from 128KB to 256KB. Due to the larger SF BAR size, for the same PF BAR2 size, which can be queried/modified by LOG_PF_BAR2_SIZE NV config, the firmware will support half of the SFs. To maintain the same amount of supported SFs, software needs to increase the LOG_PF_BAR2_SIZE NV config value by 1.
AES-XTS	AES_XTS is used to perform all disk encryption/decryption related flows in the NIC and reduce cost and overheads of the related FIPS certification.
GPUDirect in Virtualized Environment	Enabled a direct access to ATS from the NIC to GPU buffers using PCIe peer-to-peer transactions. To enable this capability, the “p2p_ordering_mode” parameter was added to the NV_PCI_CONF configuration. **Note:** When SECURE_ALL or SECURE_TRUST is configured, ATS and RO must be set identically. When SECURE_NONE is configured, ATS and RO may be set independently as the current firmware behavior allows.
Non-Volatile Configurations	Added a new Non-Volatile Configuration parameter to control VL15 buffer size (VL15_BUFFER_SIZE). **Note:** VL15 buffer size enlargement will decrease all other VLs buffers size.
NC-SI	Added a new NC-SI command (get_device_id) to report a unique device identifier.
NC-SI	Added new NC-SI commands (`get_lldp_nb, set_lldp_nb`) to query the current status of LLDP and to enable/disable it.
ROCE ACCL	Split the SlowRestart ROCE_ACCL into the following: • `slow-restart` - used to reduce rate on retransmission events • `slow-restart-after-idle` - used to reduce rate before first transmission after >1s without transmitting
ROCE ACCL	Enabled TX PSN window size configuration using LOG_TX_PSN_WINDOW NVconfig parameter. **Note:** Due to hardware limitations, `max log_tx_psn_win` value can be set 9.
Bug Fixes	See Bug Fixes.
20.28.1002	
EDR Link in ConnectX-6 100Gb/s cards	EDR link speed is now supported when using ConnectX-6 100Gb/s HCA and connecting with HDR optical cables.

NC-SI 1.2 New Commands	Implemented the following new commands from NS-SI 1.2 specification: • Get IB Link Status • Get IB Statistics • Get PF Assignment
NC-SI	Added support for Virtual node GUID, and set & get address through the NC-SI commands.
Error Injection Port Level	Added the ability to inject iCRC/vCRC port level error using Port Transmit Error Register (PTER).
In-Node Sync	Added support for in-node sync.

IPoIB Virtualization Updates	Added the following IPoIB Virtualization updates:
	• Support for SX RDMA Flow Table type in IB port
	• Support for modifying header action in IB port
	• Support for new hairpin mode:
	  • IB-to-IB
	  • Eth-to-IB
	  • IB-to-Eth

MPFS Forwarding Packets Behavior	This new feature defines the forwarding behavior in MPFS for packets arriving from the network (uplink) with destination MAC address that does not appear in the MPFS FDB. The new feature is configured by a new NV configuration (UNKNOWN_UPLINK_MAC_FLOOD) which when enabled, floods all local MPFS ports with these packets, otherwise drops these packets.
Hardware Tag Matching	Increased the maximum XRQ number to 512.
Non-Volatile Configurations (NVCONFIG)	Added the following new mlxconfig parameters to the Non-Volatile Configurations section. • `log_max_outstandng_wqe` • `ece_disable_mask`
Bug Fixes	See Bug Fixes.

9.1 Customer-Affecting Changes

Feature/ Change	Description
20.27.1016	
Link Protocol	Due to in a change in link protocol in 100GbE and 200GbE adapter cards (from PAM4 to NRZ), the link may not come up on certain configurations. For limitations related to this change, see issue 2094355.

10 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 Release Notes in https://docs.mellanox.com/category/adapterfw.

Internal Ref.	Issue
2507096	**Description:** Removed the option to create unnecessary internal CNP operation for the Lossy ADP retransmission feature.
	Keywords: RoCE, Lossy, Adp_retrans
	Discovered in Version: 20.29.1016
	Fixed in Release: 20.30.1004
2447334	**Description:** Fixed an issue related to unused port LEDs when no cable is connected to the adapter card.
	Keywords: Cables, LEDs
	Discovered in Version: 20.29.1016
	Fixed in Release: 20.30.1004
2444837	**Description:** Set the cap to 0 for high index functions to avoid too many parallel VF NODNIC functions.
	Keywords: NODNIC, VF, ETH PXE
	Discovered in Version: 20.29.1016
	Fixed in Release: 20.30.1004
2292837	**Description:** Added to firmware the option to have an adaptive VL mapping based on the configuration of the remote peer.
	Keywords: adaptive VL mapping
	Discovered in Version: 20.29.1016
	Fixed in Release: 20.29.2002
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: 20.29.1016
	Fixed in Release: 20.29.2002
2434292	**Description:** Fixed an issue with the MFS1S00-HxxxE cable that resulted in invalid firmware version display.
	Keywords: Cables
	Discovered in Version: 20.29.1016
	Fixed in Release: 20.29.2002
2410395	**Description:** Fixed an issue that prevented a SFP28 cable from linking up in a 25GbE speed.
	Keywords: Cables
	Discovered in Version: 20.29.1016

Internal Ref.	Issue
	Fixed in Release: 20.29.2002
2385117	**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. **Note:** In case of compatibility issues with ab 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).
	Keywords: MSIX
	Discovered in Version: 20.29.1016
	Fixed in Release: 20.29.2002
2360496	**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.
	Keywords: RoCE, Congestion control, DCQCN
	Discovered in Version: 20.28.1002
	Fixed in Release: 20.29.1016
2200824	**Description:** Fixed an issue that prevented VXLAN packets with svlan/cvlan tag from being matched.
	Keywords: VXLAN
	Discovered in Version: 20.28.4000
	Fixed in Release: 20.29.1016
2355328	**Description:** Fixed an issue that caused the `eth_wqe_too_small` counter to count ODP page faults.
	Keywords: Counters
	Discovered in Version: 20.28.4000
	Fixed in Release: 20.29.1016
2281266	**Description:** Fixed an issue related to raising 100GbE link on ConnectX-6 VPI 100Gb/s adapter cards.
	Keywords: Link up
	Discovered in Version: 20.28.2006
	Fixed in Release: 20.29.1016
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: 20.28.1002
	Fixed in Release: 20.29.1016

Internal Ref.	Issue
2282225/2241765	**Description:** Fixed an issue that resulted in low performance after enabling the RoCE Accelerator capability. **Note:** The fix is available when all ports are set as Ethernet.
	Keywords: Performance, RoCE
	Discovered in Version: 20.28.1002
	Fixed in Release: 20.29.1016
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: 20.26.1040
	Fixed in Release: 20.29.1016
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: 20.28.1002
	Fixed in Release: 20.28.4000
2200443	**Description:** On very rare occasions, a raw BER of 10e-12 might be experienced.
	Keywords: Raw BER
	Discovered in Version: 20.28.1002
	Fixed in Release: 20.28.4000
2244412	**Description:** ConnectX-6 Lx does not support phyless reset.
	Keywords: Phyless reset
	Discovered in Version: 20.28.1002
	Fixed in Release: 20.28.4000
2215104	**Description:** Updated the following Mellanox OEM NC-SI commands to fix an issue that caused the "Port swap" capability not to function properly: • Get Temperature • Get Module Serial Data • Set Module Serial Data
	Keywords: Port swap
	Discovered in Version: 20.28.1002
	Fixed in Release: 20.28.2006
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: 20.28.1002
	Fixed in Release: 20.28.2006

Internal Ref.	Issue
2073222	**Description:** In rare cases, HDR active copper cable link up time might be higher than expected (up to 2 minutes).
	Keywords: Cables
	Discovered in Version: 20.27.1016
	Fixed in Release: 20.28.1002
2149674	**Description:** Fixed an issue that caused packets to get stuck when the Rate Limiter was enabled.
	Keywords: Rate Limiter
	Discovered in Version: 20.27.6008
	Fixed in Release: 20.28.1002
2197232	**Description:** Active SHARP SAT QPs (QP with packet-based e2e credits) can break the live-FW-patch flow and result in firmware getting stuck.
	Keywords: SHARP SAT QPs
	Discovered in Version: 20.27.6008
	Fixed in Release: 20.28.1002
2113608	**Description:** Fixed an issue that prevented a QP with ATS buffer from being using by the NVMF offload.
	Keywords: ATS, NVMF offload
	Discovered in Version: 20.27.6008
	Fixed in Release: 20.28.1002

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.