



# **NVIDIA Quantum-2 Firmware Release Notes v31.2010.4102 LTS (2022 LTS U1)**

# Table of contents

Firmware Compatible Products	3
Changes and New Features	9
Bug Fixes in this Firmware Version	10
Known Issues	12
Changes and New Features History	14
Bug Fixes History	17

## **Note**

This is a long-term support (LTS) release. LTS is the practice of maintaining a software product for an extended period of time (up to three years) to help increase product stability. LTS releases include bug fixes and security patches.

## **Release Notes Update History**

<b>Revision</b>	<b>Date</b>	<b>Description</b>
1.0	November 17, 2022	Initial release of this release notes version.

## **Overview**

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

## **Firmware Download**

Please visit [nvidia.com/en-us/networking/](https://nvidia.com/en-us/networking/) → Support → Support → Firmware Download

---

# Firmware Compatible Products

These are the release notes for the NVIDIA Quantum™-2 firmware. This firmware complements the NVIDIA Quantum switch with a set of advanced features, allowing easy and remote management of the switch.

This firmware supports the following protocols:

- InfiniBand—SDR, EDR, HDR, NDR

## Supported Switch Systems

This firmware supports the devices listed in the table below:

Model Number	Description
QM9790	NVIDIA Quantum 2 based NDR InfiniBand Switch, 64 NDR ports, 32 OSFP ports, 2 Power Supplies (AC), Standard depth

## Firmware Interoperability

This firmware version has been validated to work against platforms with the following firmware and software versions.

HCA/Switch	Firmware Version
NVIDIA Quantum-2	31.2010.4102
NVIDIA Quantum	27.2010.4102
ConnectX-7	22.33.4030
ConnectX-6	20.33.1048
MFT	4.18.1-14

## Supported Cables

### Switch and HCAs InfiniBand Cable Connectivity Matrix

NVIDIA Quantum™ based switches and NVIDIA® ConnectX® 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
NDR	400Gb/s InfiniBand	4 lanes of 100Gb/s
NDR200	200Gb/s InfiniBand	2 lanes of 100Gb/s
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

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

### Switch-to-Switch Connectivity

NVIDIA Quantum-2 switches come with OSFP cages. NVIDIA Quantum and Switch-IB 2 switches come with QSFP cages. The connectivity matrix below are separated into multiple tables due to the above physical differences between the switches.

Switch	Switch	Cable			
		NDR Transceiver	NDR DAC/ACC	HDR DAC/AOC	EDR DAC/AOC

Switch	Switch	Cable			
NVIDIA Quantum-2	NVIDIA Quantum-2	2 × NDR	2 × NDR	N/A	N/A
NVIDIA Quantum-2	NVIDIA Quantum	N/A	N/A	2 × HDR	2 × EDR
NVIDIA Quantum-2	Switch-IB 2	N/A	N/A	N/A	2 × EDR

## HCA-to-Switch Connectivity

Switch		Adapter	Cable		
			HDR AOC	HDR DAC	HDR100 DAC/AOC (Copper Cables Only)
ConnectX-6 200Gb/s	NDR Switch	NVIDIA Quantum-2	2 × HDR	2 × HDR	4 × HDR100
ConnectX-6 100Gb/s		NVIDIA Quantum-2	N/A	2 × EDR	4 × HDR100
ConnectX-4/ ConnectX-5		NVIDIA Quantum-2	N/A	2 × EDR	N/A

## Supported Link Speed

The table below lists the current supported link speed.

Speed	Cable	Cable Length [meters]	Limitations
NDR	Optical	Up to 30	NDR optical cables support only NDR speed.
	Copper	Up to 2	

Speed	Cable	Cable Length [meters]	Limitations
HDR	Optical	Up to 30	HDR optical cables support only HDR speed.
	Copper	Up to 2	

## Validated and Supported Cables

Speed	Cable OPN #	Description
NDR	MCP7Y50-N001	NVIDIA passive copper splitter cable, IB NDR 800Gb/s to 4x200Gb/s, OSFP to 4xOSFP, 1m
NDR	MCP7Y50-N01A	NVIDIAssive copper splitter cable, IB NDR 800Gb/s to 4x200Gb/s, OSFP to 4xOSFP, 1.5m
NDR	MCP7Y50-N002	NVIDIAssive copper splitter cable, IB NDR 800Gb/s to 4x200Gb/s, OSFP to 4xOSFP, 2m
NDR	MCP7Y00-N001	NVIDIA passive copper splitter cable, IB NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP, 1m
NDR	MCP7Y00-N01A	NVIDIAssive copper splitter cable, IB NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP, 1.5m
NDR	MCP7Y00-N002	NVIDIAssive copper splitter cable, IB NDR 800Gb/s to 2x400Gb/s, OSFP to 2xOSFP, 2m
NDR	MMS4X00-NL*	NVIDIA twin port transceiver, 800Gbps, 2xNDR, OSFP, 2xMPO, 1310nm SMF, DR8, up to 30m
NDR	MCP4Y10-N00A	NVIDIA passive copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 0.5m
NDR	MCP4Y10-N00B	NVIDIA passive copper cable, IB NDR, up to 800Gb/s, OSFP, 0.75m
NDR	MCP4Y10-N001	NVIDIA passive copper cable, IB NDR, up to 800Gb/s, OSFP, 1m
NDR	MCP4Y10-N01A	NVIDIA passive Copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 1.5m
NDR	MCP4Y10-N002	NVIDIA passive copper cable, IB NDR, up to 800Gb/s, OSFP, 2m

Speed	Cable OPN #	Description
HDR	MCP7Y70-H001	NVIDIA passive copper splitter cable, IB HDR 400Gb/s to 4×100Gb/s, OSFP to 4xQSFP56, 1m
HDR	MCP7Y70-H01A	NVIDIA passive copper splitter cable, IB HDR 400Gb/s to 4×100Gb/s, OSFP to 4xQSFP56, 1.5m
HDR	MCP7Y70-H002	NVIDIA passive copper splitter cable, IB HDR 400Gb/s to 4×100Gb/s, OSFP to 4xQSFP56, 2m
HDR	MFA7U10-H003**	NVIDIA active fiber splitter cable, IB HDR, 400Gb/s to 2×200Gb/s, OSFP to 2xQSFP56, 3m
HDR	MFA7U10-H005**	NVIDIA active fiber splitter cable, IB HDR, 400Gb/s to 2×200Gb/s, OSFP to 2xQSFP56, 5m
HDR	MFA7U10-H010**	NVIDIA active fiber splitter cable, IB HDR, 400Gb/s to 2×200Gb/s, OSFP to 2xQSFP56, 10m
HDR	MFA7U10-H015**	NVIDIA active fiber splitter cable, IB HDR, 400Gb/s to 2×200Gb/s, OSFP to 2xQSFP56, 15m
HDR	MFA7U10-H020**	NVIDIA active fiber splitter cable, IB HDR, 400Gb/s to 2×200Gb/s, OSFP to 2xQSFP56, 20m
HDR	MFA7U10-H030**	NVIDIA active fiber splitter cable, IB HDR, 400Gb/s to 2×200Gb/s, OSFP to 2xQSFP56, 30m
HDR	MCP7Y60-H001	NVIDIA passive copper splitter cable, IB HDR 400Gb/s to 2×200Gb/s, OSFP to 2xQSFP56, 1m
HDR	MCP7Y60-H01A	NVIDIA passive copper splitter cable, IB HDR 400Gb/s to 2×200Gb/s, OSFP to 2xQSFP56, 1.5m
EDR	MCP7Y60-H002	NVIDIA passive copper splitter cable, IB HDR 400Gb/s to 2×200Gb/s, OSFP to 2xQSFP56, 2m

### **Warning**

\*The minimal required firmware version for MMS4X00-NL-QP1 cable is 45.110.234. \*\*The minimal required firmware version for MFA7U10-HOMFA7U10-H0xx is 40.120.327 .

## **Firmware Upgrade**

Firmware upgrade may be performed directly from any previous version to this version. To upgrade firmware, please refer to the NVIDIA Firmware Tools (MFT) package at [network.nvidia.com/products/adapter-software/firmware-tools/](https://network.nvidia.com/products/adapter-software/firmware-tools/)

## **PRM Revision Compatibility**

This firmware version complies with the NVIDIA Switches Programmer's Reference Manual (PRM), Rev 1.40 or later.

---

# Changes and New Features

Keyword	Description
General	Stability improvements.
General	See <a href="#">Bug fixes</a> .

---

# Bug Fixes in this Firmware Version

The following table provides a list of bugs fixed in this version. For a list of bug fixed from previous versions, see [Bug Fixes History](#).

Internal Ref.	Issues
3199650	<b>Description:</b> A physical link failure between switches while a SHARP job is running and utilizing the link can cause one of the switches to become invalid for further SHARP jobs. This can result in either a "No resource" response for new SHARP job requests or in jobs getting stuck. The bug fix requires SHARP version 3.2.
	<b>Keywords:</b> SHARP
	<b>Discovered in Version:</b> 31.2010.4010
	<b>Fixed in Version:</b> 31.2010.4102
3245821	<b>Description:</b> In case of an AR group table set request, the ARN mask is flushed for group that has an active pFRN timer.
	<b>Keywords:</b> PFRN
	<b>Discovered in Version:</b> 31.2010.4010
	<b>Fixed in Version:</b> 31.2010.4102
3253717	<b>Description:</b> mask_force_clear_timeout timer in pFRN feature was not functional (the mask was not cleared when the timer expired).
	<b>Keywords:</b> PFRN
	<b>Discovered in Version:</b> 31.2010.4010
	<b>Fixed in Version:</b> 31.2010.4102

Internal Ref.	Issues
3242209	<b>Description:</b> Set PFRN mad did not return error on wrong inputs in mask_clear_timer and mask_force_clear_timer fields.
	<b>Keywords:</b> PFRN
	<b>Discovered in Version:</b> 31.2010.4010
	<b>Fixed in Version:</b> 31.2010.4102

---

# Known Issues

The following sections describe known issues in firmware releases and possible workarounds.

Internal Ref.	Issue
3261861	<b>Description:</b> Connecting an HDR device to an NDR device with Optical cables longer than 30m causes degradation in the bandwidth.
	<b>Keywords:</b> HDR-to-NDR
	<b>Workaround:</b> N/A
	<b>Discovered in Version:</b> 31.2010.4102
2974424	<b>Description:</b> Currently, on cables that perform polarity inversion there is no link up.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Cables, Polarity Inversion
	<b>Discovered in Version:</b> 31.2010.3118
2922333	<b>Description:</b> In some cases, MMS4X00-NL 1.2 may have low BER.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Signal Integrity
	<b>Discovered in Version:</b> 31.2010.2110
2838195	<b>Description:</b> Using NDR speed with Optical Transceivers causes bandwidth to be 350Gb/s instead of 400Gb/s in small packets.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Optical Transceivers
	<b>Discovered in Version:</b> 31.2010.1310

Internal Ref.	Issue
2834238	<b>Description:</b> When using Optical Transceiver, toggling a port in a cage may toggle the adjacent port in the cage.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Optical Transceivers, Port Toggle
	<b>Discovered in Version:</b> 31.2010.1310
955641	<b>Description:</b> VL_HIGH_LIMIT is not affecting the VL arbiter as expected.
	<b>Workaround:</b> Arbitration table should be set using only the low priority VL arbitration table.
	<b>Keywords:</b> VL Arbitration
	<b>Discovered in Version:</b> 31.2010.1310
1249608	<b>Description:</b> Configuring weight "0" for VL, results in unexpected behavior.
	<b>Workaround:</b> Arbitration table should be configured with weights other than "0".
	<b>Keywords:</b> VL Arbitration
	<b>Discovered in Version:</b> 31.2010.1310
2057793	<b>Description:</b> Congestion profiles in VS-MAD PortProfileSetting support only fixed mode. Percentage mode is not supported.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand Congestion Control
	<b>Discovered in Version:</b> 31.2010.1310
Unsupported Features	Features that are not yet supported in the NVIDIA Quantum-2 systems: <ul style="list-style-type: none"> <li>• IB Router</li> </ul>

---

# Changes and New Features History

Keyword	Description
31.2010.4010	
pFRN Collector	Added support for mirroring of PFRN packets towards UFM entity (collector) in the subnet.
PKEY Filter for Multicast	Added support for MulticastPKeyTrapSuppression (PKEY mismatch filtering).
Congestion Control Updates	Added support for 1kb granularity for the port congestion profiles.
SL-to-VL Mapping	Added switch support for port mask optimization of SL-to-VL Mapping Table configuration.
General	See <a href="#">Bug fixes</a> .

Keyword	Description
31.2010.3118	
Hash-Based Routing	Enabled the reordering of sensitive traffic to load balance on multiple ports by using Hash-Based Routing.
General	See <a href="#">Bug fixes</a> .

Keyword	Description
31.2010.3004	
Counters	Added support for PortXmitWaitVLExtended counters.

Keyword	Description
pFRN	Added support for pFRN (Proactive Fault Routing Notification) which allows for dynamic link failure detection and route correction for topologies based on Adaptive Routing
General	See <a href="#">Bug fixes</a> .

Keyword	Description
31.2010.2300	
SHARPV3	Added GA-level support for aggregation jobs to run over parallel links.
General	<a href="#">Bug fixes</a> .

Keyword	Description
31.2010.2246	
SHARPV3	Added beta-level support for aggregation jobs to run over parallel links.
General	<a href="#">Bug fixes</a> .

Keyword	Description
31.2010.2110	
SHARPV3	Added GA-level support for SHARPV3 on Quantum-2 systems.
General	<a href="#">Bug fixes</a> .

Keyword	Description
31.2010.2036	
Systems	Added power and system monitoring optimizations.

<b>Keyword</b>	<b>Description</b>
SHARP V3	Added beta-level support for SHARP V3 on Quantum-2 systems.
Hash Based Forwarding	Added alpha-level support for Hash Based Forwarding Routing Capability.
Security	Added security enhancements to QM9790 system.

<b>Keyword</b>	<b>Description</b>
31.2010.1310	
Systems	Added GA-level support for NVIDIA Quantum-2-based switch QM9790.
Congestion Control	Added ES-level support for congestion control class key.
Vendor Key	Added ES-level support for vendor class key.
Hierarchy Information	Added support for Hierarchy Information mad.
Remote Debug Token	Added support for Remote Debug Token.
NVIDIA® Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™	Added GA-level support for SHARPV2 for NVIDIA Quantum-2 systems.
Counters	Added support for PortVLXmitFlowCtlUpdateErrors counters.
Security	Added support for Secure Firmware and Secure Firmware Boot on NVIDIA Quantum-2 systems.

---

# Bug Fixes History

The following table provides a list of bugs fixed in previous versions. For a list of bug fixed from the current version, see [Bug Fixes](#).

Internal Ref.	Issues
3174239	<b>Description:</b> On rare occasions, traps were not properly repressed, which caused redundant traps to be sent multiple times.
	<b>Keywords:</b> Traps
	<b>Discovered in Version:</b> 31.2010.3118
	<b>Fixed in Version:</b> 31.2010.4010
3002314	<b>Description:</b> On rare occasion, when port is configured to mloop toggle may cause link to not rise.
	<b>Keywords:</b> Optic in Mloop
	<b>Discovered in Version:</b> 31.2010.2110
	<b>Fixed in Version:</b> 31.2010.3118
3127727	<b>Description:</b> On rare occasion, when egress port is split to two, the egress port may get stuck due to wrong Fast Path configuration.
	<b>Keywords:</b> Switch Hang, Fast Path, Split
	<b>Discovered in Version:</b> 31.2010.3004
	<b>Fixed in Version:</b> 31.2010.3118
3082569	<b>Description:</b> In some traffic patterns involving small packets, the PortRcvErrors counter may mistakenly count events of local physical errors due to an internal flow in the hardware that involves link packets.
	<b>Keywords:</b> Counters
	<b>Discovered in Version:</b> 31.2010.2246
	<b>Fixed in Version:</b> 31.2010.3004

Internal Ref.	Issues
3085427	<b>Description:</b> On rare occasions, SHARP semaphore may remain locked on a port following an event of a port link down or an application crash.
	<b>Keywords:</b> SHARPV3
	<b>Discovered in Version:</b> 31.2010.2036
	<b>Fixed in Version:</b> 31.2010.3004
3011581	<b>Description:</b> On rare occasions, job failures with SharpError trap may be experienced as a result of previous jobs that have failed.
	<b>Keywords:</b> SHARPV3
	<b>Discovered in Version:</b> 31.2010.2036
	<b>Fixed in Version:</b> 31.2010.3004
3000602	<b>Description:</b> After disconnecting MMS4X00-NL* cable and connecting Ultron cable to the same port, ports fails to link up.
	<b>Keywords:</b> Cables
	<b>Discovered in Version:</b> 31.2010.2110
	<b>Fixed in Version:</b> 31.2010.2300
3060122	<b>Description:</b> In the event of link fault of a link between root switch and non-root switch during the run of a job, the next job run on the non-root switch may fail.
	<b>Keywords:</b> SHARPV3
	<b>Discovered in Version:</b> 31.2010.2036
	<b>Fixed in Version:</b> 31.2010.2300
2923464	<b>Description:</b> When using MMS4X00-NL Optical module, on rare occasions port that is in NDR speed may get stuck and stay in Polling state.
	<b>Keywords:</b> NDR, Optical Module
	<b>Discovered in Version:</b> 31.2010.1404
	<b>Fixed in Version:</b> 31.2010.2246

Internal Ref.	Issues
2859363	<b>Description:</b> When using NVIDIA Quantum-2 systems in Auto-Neg mode, NDR speed in one lane (1x) is not supported.
	<b>Keywords:</b> Auto-Negotiation
	<b>Discovered in Version:</b> 31.2010.1310
	<b>Fixed in Version:</b> 31.2010.2246
3033131	<b>Description:</b> The number of flows changed from 2 to 1, as intended.
	<b>Keywords:</b> SHARPV3
	<b>Discovered in Version:</b> 31.2010.2110
	<b>Fixed in Version:</b> 31.2010.2246
2972388	<b>Description:</b> Running of concurrent jobs may lead to states where jobs unexpectedly terminate or get stuck.
	<b>Keywords:</b> SHARPV3
	<b>Discovered in Version:</b> 31.2010.2036
	<b>Fixed in Version:</b> 31.2010.2110
2982113	<b>Description:</b> On rare occasions, job resource cleanup may fail.
	<b>Keywords:</b> SHARPV3
	<b>Discovered in Version:</b> 31.2010.2036
	<b>Fixed in Version:</b> 31.2010.2110
2971339	<b>Description:</b> During high load scenarios, performance degradation may be experienced.
	<b>Keywords:</b> SHARPV3
	<b>Discovered in Version:</b> 31.2010.2036
	<b>Fixed in Version:</b> 31.2010.2110

Internal Ref.	Issues
2849215	<b>Description:</b> On NVIDIA Quantum-2 switches, when working with MFA7U10-H0xx cables, if one of the ports in a cage is disabled at the time of initialization by user configuration, reenabling the port will require toggling the link (i.e. enable → disable → enable).
	<b>Keywords:</b> NVIDIA Quantum-2, Cables
	<b>Discovered in Version:</b> 31.2010.1310
	<b>Fixed in Version:</b> 31.2010.2036
2890632	<b>Description:</b> On NVIDIA Quantum-2 systems, changing the Optical module rate was not allowed.
	<b>Keywords:</b> Optical Modules
	<b>Discovered in Version:</b> 31.2010.1310
	<b>Fixed in Version:</b> 31.2010.2036
2885798	<b>Description:</b> In NVIDIA Quantum-2 systems, effective errors may occur with short Copper cable MCP4Y10-N00B.
	<b>Workaround:</b> N/A
	<b>Discovered in Version:</b> 31.2010.1310
	<b>Fixed in Version:</b> 31.2010.2036
2910161	<b>Description:</b> In auto-negotiation flow, using copper cables when toggling both port's sides may cause the port to get stuck on rare occasions.
	<b>Keywords:</b> Auto-Negotiation, Copper Cables
	<b>Discovered in Version:</b> 31.2010.1310
	<b>Fixed in Version:</b> 31.2010.2036

## 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. NVIDIA Corporation ("NVIDIA") makes no 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 NON-INFRINGEMENT, 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 and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© Copyright 2023, NVIDIA. PDF Generated on 01/15/2026