NGC-Ready Systems

Documentation

DA-09419-001 _v09 | January 2024

Table of Contents

Chapter 1. List of NGC-Ready Servers

NGC software runs on a wide variety of NVIDIA GPU-accelerated platforms, including on-premises NGC-Ready and NGC-Ready for Edge servers, NVIDIA DGX™ Systems, workstations with NVIDIA TITAN and NVIDIA Quadro® GPUs, and leading cloud platforms.

The following lists the 3rd-party systems that have been validated by NVIDIA as "NGC-Ready". NGC-Ready servers have passed an extensive suite of tests that validate their ability to deliver high performance running NGC containers. NGC-Ready system validation includes tests of:

- Single and multi-GPU Deep Learning training using TensorFlow, PyTorch and NVIDIA DeepStream Transfer Learning Toolkit
- High volume, low latency inference using NVIDIA TensorRT, TensorRT Inference Server, and DeepStream
- Data Science using RAPIDS and XGBoost
- Application development using the CUDA Toolkit.

NGC-Ready for Edge servers have, in addition to passing the NGC-Ready tests, demonstrated their ability to support the NVIDIA EGX platform that uses the industry standards of TPM for hardware-based key management and Redfish for remote systems management.

NVIDIA AI Enterprise Compatible Systems

The NGC-Ready servers listed below as NVIDIA AI Enterprise Compatible are only supported by NVIDIA AI Enterprise when they are equipped with NVIDIA V100 or T4 GPUs. NVIDIA AI Enterprise Compatible systems supporting newer generations of NVIDIA GPUs are listed on the NVIDIA-Certified systems web site. See the NVIDIA AI Enterprise technical documentation for more information about supported hardware and software platforms.

NGC-Ready Servers

Manufacturer	Server	Supported NVIDIA GPU	NVIDIA AI Enterprise Compatible[1]
Aetina	AIS-D422-A1	NVIDIA T4	Yes

Manufacturer	Server	Supported NVIDIA GPU	NVIDIA AI Enterprise Compatible[1]
ADLINK	MECS-7210	NVIDIA V100 for PCIe	Yes
Advantech	SKY-6100	NVIDIA T4	Yes
Altos	BrainSphere P550 F4	NVIDIA T4	Yes
Aparna Systems	GX2	NVIDIA T4	Yes
AsRock Rack	1U2FH-4L/C622	NVIDIA T4	Yes
AsRock Rack	1U2G-EPYC/2T	NVIDIA T4	Yes
AsRock Rack	2U2G C622	NVIDIA T4	Yes
AsRock Rack	3U8G+/C621	NVIDIA V100 for for PCIe	Yes
ASUS	ESC4000 G4	NVIDIA T4	Yes
ASUS	ESC8000 G4	NVIDIA V100 for PCIe, NVIDIA Quadro RTX 8000	Yes
ATOS	BullSequana Edge	NVIDIA T4	Yes
ATOS	BullSequana S200	NVIDIA T4	Yes
ATOS	BullSequana X112	NVIDIA V100 for NVLINK	Yes
Cisco	UCS C240 M5	NVIDIA T4	Yes
Cisco	UCS C480 ML M5	NVIDIA V100 for NVLINK	Yes
Cray	CS Storm NX	NVIDIA V100 for NVLINK	Yes
Dell EMC	DSS 8440	NVIDIA T4, NVIDIA V100 for PCIe, NVIDIA Quadro RTX 8000	Yes
Dell EMC	PowerEdge C4140	NVIDIA V100 for NVLINK, NVIDIA V100 for PCIe	Yes
Dell EMC	PowerEdge R6525	NVIDIA T4,	Yes
Dell EMC	PowerEdge R640	NVIDIA T4	Yes
Dell EMC	PowerEdge R740	NVIDIA T4, NVIDIA V100 for PCIe, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	Yes
Dell EMC	PowerEdge R6515	NVIDIA T4	Yes
Dell EMC	PowerEdge R6525	NVIDIA T4	Yes
Dell EMC	PowerEdge R7425	NVIDIA T4	Yes
Dell EMC	PowerEdge R7515	NVIDIA T4, NVIDIA V100 for PCIe	Yes

Manufacturer	Server	Supported NVIDIA GPU	NVIDIA AI Enterprise Compatible[1]
Dell EMC	PowerEdge R7525	NVIDIA T4, NVIDIA V100 for PCIe, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	Yes
Dell EMC	PowerEdge R840	NVIDIA V100 for PCIe	Yes
Dell EMC	PowerEdge R940XA	NVIDIA V100 for PCIe	Yes
Dell EMC	PowerEdge T640	NVIDIA V100 for PCIe, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	Yes
Dell EMC	PowerEdge XE2420	NVIDIA T4, NVIDIA V100 for PCIe, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	Yes
Dell EMC	PowerEdge XE7440	NVIDIA T4	Yes
Fujitsu	PRIMERGY RX2530 M5	NVIDIA T4	Yes
Fujitsu	PRIMERGY RX2540 M5	NVIDIA T4, NVIDIA V100 for PCIe	Yes
Fujitsu	PRIMERGY GX2570 M5	NVIDIA V100 for NVLink	Yes
Huawei	FusionServer G560 V5	NVIDIA T4, NVIDIA V100 for NVLINK	Yes
GDEP	Inference BOX	NVIDIA T4	Yes
GIGABYTE	E251-U70	NVIDIA T4	Yes
GIGABYTE	G191-H44	NVIDIA T4, NVIDIA V100 for PCIe, NVIDIA Quadro RTX 8000, NVIDIA Quadro RTX 6000	Yes
GIGABYTE	G241-G40	NVIDIA T4, NVIDIA V100 for PCIe, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	Yes
GIGABYTE	G242-Z10	NVIDIA T4, NVIDIA V100 for PCIe	Yes
GIGABYTE	G242-Z11	NVIDIA T4, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	Yes
GIGABYTE	G291-280	NVIDIA V100 for PCIe	Yes

Manufacturer	Server	Supported NVIDIA GPU	NVIDIA AI Enterprise Compatible[1]
GIGABYTE	G481-HA0	NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	
GIGABYTE	R281-G30	NVIDIA T4, NVIDIA V100 for PCIe, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	Yes
GIGABYTE	R282-Z93	NVIDIA T4, NVIDIA V100 for PCIe, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	Yes
GIGABYTE	R292-4S0	NVIDIA T4	Yes
GIGABYTE	W42G-P08R	NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	
H3C	UniServer R4900 G3 Server	NVIDIA V100 for PCIe	Yes
H3C	UniServer R5300 G3	NVIDIA T4, NVIDIA V100 for PCIe	Yes
Hewlett Packard Enterprise	Apollo 2000 / ProLiant XL190R Gen10	NVIDIA T4	Yes
Hewlett Packard Enterprise	Apollo 6500 / ProLiant XL270D Gen10	NVIDIA V100 for NVLINK	Yes
Hewlett Packard Enterprise	Edgeline EL1000	NVIDIA T4	Yes
Hewlett Packard Enterprise	Edgeline EL4000	NVIDIA T4	Yes
Hewlett Packard Enterprise	ProLiant DL380	NVIDIA T4	Yes
Hewlett Packard Enterprise	Edgeline EL8000 / ProLiant e910	NVIDIA T4	Yes
Inspur	NE5260M5	NVIDIA T4	Yes
Inspur	NF5280M5	NVIDIA T4	Yes
Inspur	NF5468M5	NVIDIA V100 for NVLINK, NVIDIA T4	Yes
Inspur	NF5488M5	NVIDIA V100 for NVLINK	Yes

Manufacturer	Server	Supported NVIDIA GPU	NVIDIA AI Enterprise Compatible[1]
Inventec	E850G4	NVIDIA T4	Yes
Lanner Electronics	FX-3420	NVIDIA T4	Yes
Lanner Electronics	LEC-2290E	NVIDIA T4	Yes
Lenovo	ThinkSystem SE350	NVIDIA T4	Yes
Lenovo	ThinkSystem SR670	NVIDIA T4	Yes
MiTAC Rack	2U/Firestone	NVIDIA T4	Yes
NetApp	HCI H615C	NVIDIA T4	Yes
PNY	PNYEXI244000T4-110	NVIDIA T4	Yes
PNY	PNYSRA14-SERIES-100	NVIDIA T4	Yes
PNY	PNYSRA22-SERIES-100	NVIDIA Quadro RTX 8000	
PNY	PNYSRA48-SERIES-100	NVIDIA Quadro RTX 8000	
PNY	PNYSRA222RTX6K-110	NVIDIA Quadro RTX 6000	
QCT	QuantaGrid D43K-1U	NVIDIA T4	Yes
QCT	QuantaGrid D52BV-2U	NVIDIA T4, NVIDIA V100 for PCIe, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	Yes
QCT	QuantaGrid D52G-4U	NVIDIA V100 for NVLINK, NVIDIA V100 for PCIe, NVIDIA Quadro RTX 8000	Yes
QCT	QuantaGrid D52Y-2U	NVIDIA T4, NVIDIA V100 for PCIe	Yes
QCT	QuantaGrid SD2H-1U	NVIDIA T4	Yes
QCT	QuantaGrid S43KL-1U	NVIDIA T4	Yes
Sugon	X760-G30	NVIDIA T4	
Sugon	X795-G30	NVIDIA V100 for NVLINK	

Manufacturer	Server	Supported NVIDIA GPU	NVIDIA AI Enterprise Compatible[1]
Supermicro	SYS-1019D-FHN13TP	NVIDIA T4	Yes
Supermicro	SYS-1019P-FHN2T	NVIDIA T4	Yes
Supermicro	SYS-1019P-WTR	NVIDIA T4	Yes
Supermicro	SYS-1029U-TRT	NVIDIA T4	Yes
Supermicro	SYS-2029GP-TR	NVIDIA V100 - PCIe, NVIDIA T4	Yes
Supermicro	SYS-4029GP-TRT2	NVIDIA V100 - PCIe	Yes
Supermicro	SYS-4029GP-TVRT	NVIDIA V100 for NVLINK, NVIDIA T4	Yes
Supermicro	SYS-5019D-FN8TP	NVIDIA T4	Yes
Supermicro	SYS-5039MD8-H8TNR	NVIDIA T4	Yes
Supermicro	SYS-5039MD18-H8TNR	NVIDIA T4	Yes
Supermicro	SYS-5039MP-H8TNR	NVIDIA T4	Yes
Supermicro	SYS-7049GP-TRT	NVIDIA V100 for PCIe, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000	Yes
Tyan	Thunder CX GT24E-B5556.	NVIDIA T4	Yes
Tyan	Thunder HX FT83-B7119	NVIDIA T4	Yes
Tyan	Thunder TN76-B7102	NVIDIA T4	Yes
Tyan	TRANSPORT HX TN83B8251	NVIDIA T4	Yes
Wiwynn	EP100	NVIDIA T4	Yes
Wiwynn	SV310G3	NVIDIA T4	Yes

Note:

1. Only NGC-Ready Servers systems equipped with NVIDIA V100 or T4 GPUs are NVIDIA AI Enterprise Compatible. NVIDIA-Certified systems that are NVIDIA AI Enterprise

Compatible are listed on the NVIDIA-Certified systems documentation site. See the NVIDIA AI Enterprise technical documentation for more information.

NGC-Ready Data Science Workstations

Manufacturer	Workstation	Supported NVIDIA GPU
APY	AI Lx	NVIDIA Quadro GV100, RTX 6000 and RTX 8000
APY	AI Lx2 G2	NVIDIA Quadro GV100, RTX 6000 and RTX 8000
ASUS	E900 G4	NVIDIA Quadro RTX 6000 and RTX 8000
ASUS	Pro E800 G4	NVIDIA Quadro RTX 6000 and RTX 8000
Azken Muga	W45 Data Science	NVIDIA Quadro GV100, RTX 6000, and RTX 8000
BOXX	Apexx W3	NVIDIA Quadro GV100, RTX 6000 and RTX 8000
BOXX	Apexx D4	NVIDIA Quadro GV100, RTX 6000 and RTX 8000
Colfax	SXT9700	NVIDIA Quadro RTX 8000
Dell	Precision T5820 Tower	NVIDIA Quadro GV100, RTX 6000, and RTX 8000
Dell	Precision T7920 Tower	NVIDIA Quadro GV100, RTX 6000, and RTX 8000
Dell	Precision 7920 Rack	NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000
Delta Computer	Data-Science-Workstation	NVIDIA Quadro GV100, RTX 6000, and RTX 8000
Exxact	TWS-1686525-NDS	NVIDIA Quadro RTX 6000
Exxact	Valence VWS-1735800-NDS	NVIDIA Quadro RTX 6000
Forsite	DSWS	NVIDIA Quadro GV100, RTX 6000, and RTX 8000
Forsite	DSWS PRO	NVIDIA Quadro GV100, RTX 6000, and RTX 8000
GDEP	Deep Learning Box	NVIDIA Quadro RTX 6000 and RTX 8000
GDEP	Deep Learning Station	NVIDIA Quadro GV100, RTX 6000 and RTX 8000

Manufacturer	Workstation	Supported NVIDIA GPU
HP	Z2 Tower G5 Workstation	NVIDIA Quadro RTX 5000, RTX 6000, RTX 8000
HP	Z4 G4	NVIDIA Quadro GV100, RTX 6000 and RTX 8000
HP	Z6 G4 Workstation	NVIDIA Quadro GV100, RTX 5000, RTX 6000, RTX 8000
HP	Z8 G4	NVIDIA Quadro GV100, RTX 6000 and RTX 8000
HP	ZCentral 4R Workstation	NVIDIA Quadro RTX 6000, RTX 8000
Image et Technologie	Quadro RTX Workstation	NVIDIA Quadro RTX 6000 and RTX 8000
Leadtek	WinFast WS2030	NVIDIA Quadro RTX 6000, RTX 8000
Leadtek	WinFast WS930	NVIDIA Quadro RTX 6000, RTX 8000
Leadtek	WS830	NVIDIA Quadro GV100, RTX 6000 and RTX 8000
Leadtek	WS1030	NVIDIA Quadro GV100, RTX 6000 and RTX 8000
Lenovo	ThinkStation P520 Tower	NVIDIA Quadro GV100, RTX 6000 and RTX 8000
Lenovo	ThinkStation P620 Tower	NVIDIA Quadro RTX 6000
Lenovo	ThinkStation P920 Tower	NVIDIA Quadro GV100, RTX 6000 and RTX 8000
Microway	Data Science WhisperStation	NVIDIA Quadro RTX 6000
Nextron	ScienceStation	NVIDIA Quadro RTX 6000 and RTX 8000
One Stop Systems	OSS-DSPRO	NVIDIA Quadro RTX 6000
RAVE	RAVE-DSW	NVIDIA Quadro RTX 6000
Scan	3XS Data Science Workstation G1000X	NVIDIA Quadro GV100, RTX 6000, and RTX 8000
Scan	3XS Data Science Workstation G2000X	NVIDIA Quadro GV100, RTX 6000, and RTX 8000
sysGen	devCUBE Data Science Workstation	NVIDIA Quadro GV100, RTX 6000, and RTX 8000

NGC-Ready Data Science Mobile Workstations

Manufacturer	Workstation	Supported NVIDIA GPU
Dell	Precision 7540 Mobile Workstation	NVIDIA Quadro RTX 5000
Dell	Precision 7740 Mobile Workstation	NVIDIA Quadro RTX 5000
Dell	Precision 7550 Mobile Workstation	NVIDIA Quadro RTX 5000
Dell	Precision 7750 Mobile Workstation	NVIDIA Quadro RTX 5000
HP	ZBook 17 G6 Mobile Workstation	NVIDIA Quadro RTX 5000
HP	ZBook Fury 17 Mobile Workstation	NVIDIA Quadro RTX 5000
HP	ZBook Studio G7 Mobile Workstation	NVIDIA Quadro RTX 5000
Lenovo	ThinkPad P53 Mobile Workstation	NVIDIA Quadro RTX 5000
Lenovo	ThinkPad P73 Mobile Workstation	NVIDIA Quadro RTX 5000
Lenovo	ThinkPad P15 Mobile Workstation	NVIDIA Quadro RTX 5000
Lenovo	ThinkPad P17 Mobile Workstation	NVIDIA Quadro RTX 5000

Chapter 2. List of NGC-Ready for Edge Systems

NGC-Ready system validation includes tests of:

- Single and multi-GPU Deep Learning training using TensorFlow, PyTorch and NVIDIA DeepStream Transfer Learning Toolkit
- High volume, low latency inference using NVIDIA TensorRT, TensorRT Inference Server, and DeepStream
- Data Science using RAPIDS and XGBoost
- Application development using the CUDA Toolkit.

NGC-Ready for Edge servers have, in addition to passing the NGC-Ready tests, demonstrated their ability to support the NVIDIA EGX platform that uses the industry standards of TPM for hardware-based key management and IPMI for remote systems management.

The following lists the 3rd-party systems that have been validated by NVIDIA as "NGC-Ready for Edge".

NGC-Ready for Edge Servers

Manufacturer	Server	Supported NVIDIA GPU
Advantech	SKY-6100	NVIDIA T4
Altos	BrainSphere P550 F4	NVIDIA T4
AsRock Rack	1U2FH-4L/C622	NVIDIA T4
AsRock Rack	1U2G-EPYC/2T	NVIDIA T4
AsRock Rack	2U2G C622	NVIDIA T4
ATOS	BullSequana Edge	NVIDIA T4
Dell EMC	PowerEdge R640	NVIDIA T4
Dell EMC	PowerEdge R740	NVIDIA T4, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000
Dell EMC	PowerEdge R6515	NVIDIA T4

Manufacturer	Server	Supported NVIDIA GPU
Dell EMC	PowerEdge R6525	NVIDIA T4
Dell EMC	PowerEdge R7515	NVIDIA T4
Dell EMC	PowerEdge R7425	NVIDIA T4
Dell EMC	PowerEdge XE2420	NVIDIA T4, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000
Fujitsu	PRIMERGY RX2530 M5	NVIDIA T4
Fujitsu	PRIMERGY RX2540 M5	NVIDIA T4
GIGABYTE	E251-U70	NVIDIA T4
GIGABYTE	G191-H44	NVIDIA T4
GIGABYTE	G241-G40	NVIDIA T4, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000
GIGABYTE	G242-Z10	NVIDIA T4, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000
GIGABYTE	G242-Z11	NVIDIA T4, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000
GIGABYTE	G481-HA0	NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000
GIGABYTE	R281-G30	NVIDIA T4, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000
GIGABYTE	R282-Z93	NVIDIA T4, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000
GIGABYTE	R292-4S0	NVIDIA T4
H3C	R4900 G3	NVIDIA T4
H3C	R5300G3	NVIDIA T4
Hewlett Packard Enterprise	Edgeline EL1000	NVIDIA T4
Hewlett Packard Enterprise	Edgeline EL4000	NVIDIA T4
Hewlett Packard Enterprise	ProLiant DL380	NVIDIA T4
Hewlett Packard Enterprise	Edgeline EL8000 / ProLiant e910	NVIDIA T4
Inspur	NE5260M5	NVIDIA T4

Manufacturer	Server	Supported NVIDIA GPU
Inspur	NF5280M5	NVIDIA T4
Inventec	E850G4	NVIDIA T4
Lenovo	ThinkSystem SE350	NVIDIA T4
Lenovo	ThinkSystem SR670	NVIDIA T4
MiTAC Rack	2U/Firestone	NVIDIA T4
PNY	PNYEXI244000T4-110	NVIDIA T4
QCT	QuantaGrid D43K-1U	NVIDIA T4
QCT	QuantaGrid D52BV-2U	NVIDIA T4, NVIDIA Quadro RTX 6000, NVIDIA Quadro RTX 8000
QCT	QuantaGrid D52Y-2U	NVIDIA T4
QCT	QuantaGrid SD2H-1U	NVIDIA T4
QCT	QuantaGrid S43KL-1U	NVIDIA T4
PNY	PNYEXI244000T4-110	NVIDIA T4
Supermicro	SYS-1019D-FHN13TP	NVIDIA T4
Supermicro	SYS-1019P-FHN2T	NVIDIA T4
Supermicro	SYS-1019P-WTR	NVIDIA T4
Supermicro	SYS-2029GP-TR	NVIDIA T4
Supermicro	SYS-5019D-FN8TP	NVIDIA T4
Tyan	Thunder CX GT24E-B5556.	NVIDIA T4
Tyan	Thunder TN76-B7102	NVIDIA T4
Wiwynn	EP100	NVIDIA T4

Chapter 3. NGC-Ready Supported Software

NGC-Ready Testing Software Environment

NGC-Ready and NGC-Ready for Edge servers are tested using standardized software environments that have been demonstrated to provide the highest levels of stability and performance. These NGC-Ready and NGC-Ready for Edge servers were tested using the following software:

- Ubuntu 16.04, 18.04, or 20.04
- RHEL 7.5 or 7.6
- NVIDIA drivers
- Docker-ce
- NVIDIA Container Runtime

List of NGC-Ready Supported Software

The following lists the software supported by NVIDIA as "NGC-Ready".

- Containers published on NGC by NVIDIA including:
  - TensorFlow: https://ngc.nvidia.com/catalog/containers/nvidia:tensorflow
  - PyTorch: https://ngc.nvidia.com/catalog/containers/nvidia:pytorch
  - TensorRT: https://ngc.nvidia.com/catalog/containers/nvidia:tensorrt
  - Triton Inference Server: https://ngc.nvidia.com/catalog/containers/nvidia:tritonserver
  - CUDA: https://ngc.nvidia.com/catalog/containers/nvidia:cuda
  - NVCaffe: https://ngc.nvidia.com/catalog/containers/nvidia:caffe
- RAPIDS 0.5 and later releases: https://ngc.nvidia.com/catalog/containers/nvidia:rapidsai:rapidsai

Free online support for NGC is available at the NVIDIA Developer Forum.

Enterprise support subscriptions for NGC-Ready systems are available through NVIDIA AI Enterprise Support Services.

Containers published on NGC by 3rd parties are supported by their respective publishers.

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 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, BlueField, ConnectX, CUDA, GPUDirect, NVIDIA-Certified Systems, NVIDIA HGX, NVIDIA RTX, and TensorRT 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.

Arm

Arm, AMBA and Arm Powered are registered trademarks of Arm Limited. Cortex, MPCore and Mali are trademarks of Arm Limited. All other brands or product names are the property of their respective holders. "Arm" is used to represent Arm Holdings plc; its operating company Arm Limited; and the regional subsidiaries Arm Inc.; Arm KK; Arm Korea Limited.; Arm Taiwan Limited; Arm France SAS; Arm Consulting (Shanghai) Co. Ltd.; Arm Germany GmbH; Arm Embedded Technologies Pvt. Ltd.; Arm Norway, AS and Arm Sweden AB.

Copyright

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
www.nvidia.com