



NVIDIA-Certified Systems

Documentation

Table of Contents

Chapter 1. NVIDIA-Certified Systems.....	1
Chapter 2. NVIDIA-Certified Systems Testing.....	3
Chapter 3. NVIDIA AI Enterprise Compatible Systems.....	4
Chapter 4. List of NVIDIA-Certified Systems - OVX Servers.....	5
Chapter 5. List of NVIDIA-Certified Systems - Data Center Servers.....	6
Chapter 6. List of NVIDIA-Certified Systems - Workstations.....	16
Chapter 7. List of NVIDIA-Certified Systems - Mobile Workstations.....	19
Chapter 8. List of NVIDIA-Certified Systems - High-Density VDI Servers.....	20
Chapter 9. List of NVIDIA-Certified Systems - Edge Systems.....	22
Chapter 10. NVIDIA-Certified Systems Supported Software.....	27

Chapter 1. NVIDIA-Certified Systems

The NVIDIA-Certified Systems program has assembled the industry's most complete set of accelerated workload performance tests to help its partners deliver the highest performing systems. NVIDIA-Certified Systems are tested with the most powerful enterprise NVIDIA GPUs and networking and are evaluated by NVIDIA engineers for performance, functionality, scalability, and security. NVIDIA-Certified Systems have been proven to deliver predictable performance and enable enterprises to quickly deploy optimized platforms for AI, Data Analytics, HPC, high-density VDI, and other accelerated workloads in the data center, at the Edge, and on the desktop.

NVIDIA has expanded the NVIDIA-Certified Systems program beyond servers designed for the data center to include GPU-powered workstations, high-density VDI systems, and Edge devices. NVIDIA-Certified systems for the data center are tested both as single nodes and in a 2-node configuration. Workstations, high-density VDI, and Edge systems in the NVIDIA-Certified systems program are evaluated on their standalone performance with the NVIDIA GPUs within a single system.

Guidelines for configuring NVIDIA-Certified Systems to achieve best results for running a range of accelerated computing workloads can be found [here](#).

The NVIDIA NGC Catalog is the hub for GPU-accelerated and network-optimized software for AI and other compute-intensive workloads. It simplifies deployments and shortens time-to-solution with curated containers, pre-trained models, resources, SDKs, and Helm charts. NGC software runs on a wide variety of NVIDIA GPU-accelerated platforms, including [NVIDIA DGX™ Systems](#), on-premises servers from NVIDIA partners, and leading cloud platforms. NVIDIA-Certified Systems are tested using software from the NGC Catalog, and customers can choose to purchase enterprise-grade support through [NVIDIA AI Enterprise Support Services](#). NVIDIA AI Enterprise Support Services provides direct access to NVIDIA subject matter experts who can quickly address software issues to minimize system downtime and maximize system utilization and user productivity.

NVIDIA-Certified systems for the data center are tested with NVIDIA networking but can be purchased and deployed by customers with their choice of networking. NVIDIA partners resell NVIDIA AI Enterprise Support Services. These services are available for NVIDIA-Certified systems configured with any network adaptors and deployed in any data center.

Visit the following pages for additional information:

- ▶ [NVIDIA-Certified Systems](#)

- ▶ [NVIDIA-Certified Systems Configuration Guide \(http://docs.nvidia.com/ngc/ngc-deploy-on-premises/nvidia-certified-configuration-guide/index.html\)](http://docs.nvidia.com/ngc/ngc-deploy-on-premises/nvidia-certified-configuration-guide/index.html)
- ▶ [NGC Catalog](#)
- ▶ [NVIDIA AI Enterprise Software](#)
- ▶ [NVIDIA AI Enterprise Support Services](#)
- ▶ [NGC-Ready Servers](#)

Chapter 2. NVIDIA-Certified Systems Testing

NVIDIA-Certified Systems have successfully completed a rigorous suite of functional and performance tests. NVIDIA-Certified systems for the data center are tested both as single nodes and in a 2-node configuration. Workstations, high-density VDI systems, and Edge devices in the NVIDIA-Certified systems program are evaluated on their standalone performance with the NVIDIA GPUs within a single system.

NVIDIA-Certified Systems testing of Data Center Servers, Workstations, High-Density VDI systems, and Edge devices includes measures of:

- ▶ Single and multi-GPU Deep Learning training performance using TensorFlow and PyTorch
- ▶ High volume, low latency inference using NVIDIA TensorRT and TRITON
- ▶ GPU-Accelerated Data Analytics & Machine Learning using RAPIDS
- ▶ Application development using the NVIDIA CUDA Toolkit and the NVIDIA HPC SDK

NVIDIA-Certified Systems - Data Center Servers are tested for:

- ▶ Multi-Node Deep Learning training performance
- ▶ High bandwidth, low latency networking, and accelerated packet processing
- ▶ System-level security and hardware-based key management

Chapter 3. NVIDIA AI Enterprise Compatible Systems

NVIDIA AI Enterprise, version 2.0 and later, supports bare metal and virtualized deployments. All NVIDIA-Certified Data Center Servers and NGC-Ready servers with eligible NVIDIA GPUs are NVIDIA AI Enterprise Compatible for bare metal deployments. The NVIDIA-Certified systems that NVIDIA has validated for deployment in a VMware vSphere environment are listed below.

The NVIDIA GPUs supported by VMware vSphere on each partner system are listed on the [VMWare Compatibility Guide](#). For bare metal deployments with Red Hat Enterprise Linux, refer to the [Red Hat Certified Hardware](#) page for supported systems.

NGC-Ready systems that are NVIDIA AI Enterprise Compatible are listed on the [NGC-Ready systems documentation site](#).

See the [NVIDIA AI Enterprise technical documentation](#) for more information on supported systems.

Chapter 4. List of NVIDIA-Certified Systems - OVX Servers

The following systems have been validated by NVIDIA as NVIDIA-Certified Systems – OVX Servers with the supported NVIDIA GPUs listed in the following table. OVX Servers are also tested for Zero Trust deployments using Bluefield-3 3220.

Partner	OVX Server	Supported NVIDIA GPUs ¹	Zero Trust Compatible ²
HPE	Proliant DL385 Gen11	L40S, L40	Yes



Note:

1. NVIDIA-Certified systems are tested with each supported NVIDIA GPU to ensure that they can work together to provide the highest level of performance and reliability.
2. Zero-trust mode is a specialization of DPU mode which implements an additional layer of security where the host system administrator is prevented from accessing the DPU from the host. Once zero-trust mode is enabled, the data center administrator should control the DPU entirely through the Arm cores and/or BMC connection instead of through the host.

Chapter 5. List of NVIDIA-Certified Systems - Data Center Servers

The following systems have been validated by NVIDIA as NVIDIA-Certified Systems - Data Center Servers with the supported NVIDIA GPUs listed in the following table. The GPUs listed below are certified for Data Center Servers for every generally available GPU version. For example, every system validated for the A100 GPU is certified for both the A100 40 GB and A100 80 GB GPU versions.

Partner	Data Center Server	Supported NVIDIA GPUs ¹	NVIDIA AI Enterprise Compatible ²
Aetina	AIS-D422-A1	L4, A2, T4	Bare Metal
Advantech	SKY-640V2	A100, A30	Bare Metal
Altos Computing	BrainSphere R685 F5	RTX A6000, A40	Bare Metal
Asus	ESC8000-E11	L4	Bare Metal
AMAX	AceleMax DGS- 140W	A100, A30, A40	Bare Metal
AMAX	AceleMax DGS- 214A	A100, A30, A40	Bare Metal
AMAX	AceleMax DGS- 224AS	HGX A100 - 4 GPU	Bare Metal
AMAX	AceleMax DGS- 260W	A100, A30, A40	Bare Metal
AMAX	AceleMax DGS- 410W	A100, A30, A40	Bare Metal
AMAX	AceleMax DGS- 428A	A100, A30, A40	Bare Metal
AMAX	AceleMax DGS- 428AS	HGX A100 - 8 GPU	Bare Metal
AMAX	AceleMax DGS- 428WS	HGX A100 - 8 GPU	Bare Metal
AMAX	AceleMax DL- E440W	A100, A30, A40	Bare Metal
ASRock Rack	1U4G-ROME	A100, A30	Bare Metal
ASRock Rack	2U-Open19-m3.xlarge.x86	A100	Bare Metal
ASRock Rack	2U-Open19-n3.xlarge.x86	A100	Bare Metal
ASRock Rack	2U4G-ROME/2T	A100, A30	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs ¹	NVIDIA AI Enterprise Compatible ²
ASRock Rack	4U8G-ROME2/2T	A100	Bare Metal
ASRock Rack	4U10G-ICX2/2T	A100	Bare Metal
ASRock Rack	4U10G-ROME2/2T	A100	Bare Metal
ASUS	ESC4000-E10 / ESC4000-E10S	A100	Bare Metal, vSphere ³
ASUS	ESC4000A-E10	A100, A30, A40, RTX A6000	Bare Metal, vSphere ³
ASUS	ESC8000A-E11	A100	Bare Metal
ASUS	ESC8000A-E11	A40, A30	Bare Metal, vSphere ³
ASUS	ESC8000A-E12	H100	Bare Metal
ASUS	ESC N8-E11	HGX H100 - 8 GPU	Bare Metal
Atipa Technologies	Altezza SX227-32G3	A100, A40	Bare Metal
Atipa Technologies	Altezza SX427-32G10	A100, A40	Bare Metal
Atipa Technologies	Altezza SX427-32G8	A100, A40	Bare Metal
Atipa Technologies	Altezza SX427-32G8SXM4	HGX A100 - 8 GPU	Bare Metal
Atipa Technologies	Procyon SE218-8G4	A100, T4	Bare Metal
Atipa Technologies	Procyon SE218-8G8	A100, A30	Bare Metal
Atipa Technologies	Procyon SE228-16G4SXM4	HGX A100 - 4 GPU	Bare Metal
Atipa Technologies	Procyon SE228-16G8	A100, A40, A30	Bare Metal
Atipa Technologies	Procyon SE228-32G3	A100	Bare Metal
Atipa Technologies	Procyon SE428-32G10	A100, A30	Bare Metal
Atipa Technologies	Procyon SE428-32G8	A100, A30, A40	Bare Metal
Atipa Technologies	Procyon SE428-32G8SXM4	HGX A100 - 8 GPU	Bare Metal
ATOS	BullSequana SA20G	A100	Bare Metal, vSphere ³
ATOS	BullSequana X410-A5	A100	Bare Metal
Boston	ANNA Ampere L1	HGX A100 - 4 GPU	Bare Metal
Boston	ANNA Ampere M1	A100, A40, A30	Bare Metal
Boston	ANNA Ampere S2	A100, A40	Bare Metal
Boston	ANNA Ampere XL1	A100, A40, A30	Bare Metal
BOXX	RAXX P4G	A40	Bare Metal
Cisco	UCS C240 M6 Rack Server⁴	A100, A30	Bare Metal, vSphere ³
Cisco	UCS C240 M7⁴	H100	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs ¹	NVIDIA AI Enterprise Compatible ²
Cisco	UCS-X⁴	A100	Bare Metal
Colfax	Colfax CX2460s-EK8 2U Rackmount Server	HGX A100 - 4 GPU	Bare Metal
Colfax	Colfax CX41060s-XK8 4U Rackmount Server	A100, A30, A40	Bare Metal
Colfax	Colfax CX4860s-EK8 4U Rackmount Server	A100, A30, A40	Bare Metal
Colfax	Colfax CX4880s-EK8 4U Rackmount Server	HGX A100 - 8 GPU	Bare Metal
Colfax	Colfax CX4880s-XK8 4U Rackmount Server	HGX A100 - 8 GPU	Bare Metal
Colfax	Colfax ProEdge SXP9000 Workstation	A100, A30, A40	Bare Metal
DataON	AZS-6208G	A40, A30, A2	Bare Metal
Dell Technologies	DSS8440	A100, A40, T4	Bare Metal, vSphere ³
Dell Technologies	PowerEdge R650	L4, T4	Bare Metal
Dell Technologies	PowerEdge R660 Rack Server	L4	Bare Metal
Dell Technologies	PowerEdge R6615 Rack Server	L4	Bare Metal
Dell Technologies	PowerEdge R6625 Rack Server	L4	Bare Metal
Dell Technologies	PowerEdge R740/R740xd	A100, A40, A30, A10, T4	Bare Metal, vSphere ³
Dell Technologies	PowerEdge R750 Rack Server	L4, L40, H100, A100, A40, A30, A10, T4	Bare Metal, vSphere ³
Dell Technologies	PowerEdge R750XA Rack Server	L4, H100, A100, A40, A30, A10, T4	Bare Metal
Dell Technologies	PowerEdge R7515 Rack Server	L4, A30	Bare Metal
Dell Technologies	PowerEdge R7525 Rack Server	L40, L4, H100, A100, A40, A30, A10, T4	Bare Metal
Dell Technologies	PowerEdge R760 Rack Server	L4, L40, H100, A100, A40, A30	Bare Metal
Dell Technologies	PowerEdge R760xa Rack Server	L40, L4	Bare Metal
Dell Technologies	PowerEdge R760xd2 Rack Server	L4	Bare Metal
Dell Technologies	PowerEdge R7615 Rack Server	L40, A100, A40, A30	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs ¹	NVIDIA AI Enterprise Compatible ²
Dell Technologies	PowerEdge R7625 Rack Server	L4, L40, A100, A40, A30	Bare Metal
Dell Technologies	PowerEdge R960 Server Node	A100	Bare Metal
Dell Technologies	PowerEdge XE9680 Rack Server	HGX H100 - 8GPU, HGX A100-SXM4	Bare Metal
Dell Technologies	PowerEdge XE8545 Rack Server	HGX A100 - 4 GPU	Bare Metal, vSphere ³
Dell Technologies	PowerEdge XR11 Rack Server	L4, T4	Bare Metal
Dell Technologies	PowerEdge XR12 Rack Server	L4, A100, A40, A30, A10, T4	Bare Metal
Dell Technologies	PowerEdge XR5610 Rack Server	L4	Bare Metal
Dell Technologies	PowerEdge XR7620 Rack Server	L4, A100, A30	Bare Metal
Dell Technologies	PowerEdge XR8620t	L4	Bare Metal
Dell Technologies	VxRail V570 / V570F	A100, A40, A30, T4	Bare Metal, vSphere ³
Dell Technologies	VxRail V670F	A100, A40, A30, T4	Bare Metal, vSphere ³
Exxact	TensorEX TS2-185671979	A100, A30, A40	Bare Metal
Exxact	TensorEX TS2-169843133	A100	Bare Metal
Exxact	TensorEX TS2-150341732	A100, A30	Bare Metal
Exxact	TensorEX TS4-194492555	A100, A30, A40	Bare Metal
Exxact	TensorEX TS2-184920243	A100	Bare Metal
Exxact	TensorEX TS2-171138844	HGX A100 - 4 GPU	Bare Metal
Exxact	TensorEX TS4-195183185	HGX A100 - 8 GPU	Bare Metal
Exxact	TensorEX TS4-173535991	A100, A30, A40	Bare Metal
Exxact	TensorEX TS1-147097335	A100, A30, A40	Bare Metal
Exxact	TensorEX TS2-144629457	A100, A30, A40	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs ¹	NVIDIA AI Enterprise Compatible ²
Exxact	TensorEX TS4-168747704	HGX A100 - 8 GPU	Bare Metal
Exxact	TensorEX TS4-133524070	A100, A30, A40	Bare Metal
Exxact	TensorEX TWS-115999024	A100, A30, A40	Bare Metal
Fujitsu	PRIMERGY GX2460 M1	A100, A30	Bare Metal, vSphere ³
Fujitsu	PRIMERGY GX2560 M7	HGX H100 - 4 GPU	Bare Metal
Fujitsu	PRIMERGY GX2570 M6	HGX A100 - 8 GPU	Bare Metal
Fujitsu	PRIMERGY RX2450 M1	A100, A40, A30, RTX A6000	Bare Metal
Fujitsu	PRIMERGY RX2540 M5	A100, T4	Bare Metal, vSphere ³
Fujitsu	PRIMERGY RX2540 M6	L40, L4	Bare Metal
Fujitsu	PRIMERGY RX2540 M7	L40, L4, RTX A6000, H100, A100, A40, A30	Bare Metal
GIGABYTE	E152-ZEO	A100	Bare Metal
GIGABYTE	E162-220	A100, A40	Bare Metal
GIGABYTE	G242-Z11	A100, T4	Bare Metal
GIGABYTE	G242-Z12	A100, A40	Bare Metal
GIGABYTE	G262-ZR0	HGX A100 - 4 GPU	Bare Metal
GIGABYTE	G292-Z20	NVIDIA A100 for PCIe	Bare Metal
GIGABYTE	G292-Z24	A100, A30	Bare Metal
GIGABYTE	G292-Z40	A100, A40, A30	Bare Metal
GIGABYTE	G292-Z44	A100, A40, A10	Bare Metal
GIGABYTE	G292-280	A100, A40	Bare Metal
GIGABYTE	G293-S40-A	L4	Bare Metal
GIGABYTE	G293-S42-A	L40S	Bare Metal
GIGABYTE	G293-S43	L4	Bare Metal
GIGABYTE	G363-ZR0-A	H100	Bare Metal
GIGABYTE	G482-Z54	A100, A30, A40	Bare Metal
GIGABYTE	G492-HA0	A100, A40	Bare Metal
GIGABYTE	G492-H80	A100, A40	Bare Metal
GIGABYTE	G492-ID0	HGX A100 - 8 GPU	Bare Metal
GIGABYTE	G492-ZD2	HGX A100 - 8 GPU	Bare Metal
GIGABYTE	G492-Z51	A100, A30	Bare Metal
GIGABYTE	G493-SB0-A	L40, L4	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs ¹	NVIDIA AI Enterprise Compatible ²
GIGABYTE	G493-SB3-A	L40S	Bare Metal
GIGABYTE	G493-ZB1-A	L40	Bare Metal
GIGABYTE	G593-SD0-A	HGX H100 - 8 GPU	Bare Metal
GIGABYTE	R262-ZA2	A100, RTX A6000	Bare Metal
GIGABYTE	R281-G30	A100, T4	Bare Metal
GIGABYTE	R282-G30	A100, A40	Bare Metal
GIGABYTE	R282-Z93	A100, A40	Bare Metal, vSphere ³
GIGABYTE	R282-Z96	A100, A10	Bare Metal
GIGABYTE	R283-S93-A	L40	Bare Metal
H3C	UniServer R4900 G5	A100, A30, A10	Bare Metal
H3C	UniServer R5300 G5	A100, A40, A30	Bare Metal, vSphere ³
H3C	UniServer R5500 G5 AMD	HGX A100 - 8 GPU	Bare Metal
H3C	UniServer R5500 G5 Intel	HGX A100 - 8 GPU	Bare Metal
Hitachi	HA8000V/DL360 Gen10	T4	Bare Metal
Hitachi	HA8000V/DL360 Gen10 Plus	A100, T4	Bare Metal
Hitachi	HA8000V/DL380 Gen10 Plus	A100, A40, A30, A10, T4	Bare Metal, vSphere ³
Hitachi Vantara	Hitachi Advanced Server DS120 G2	T4	Bare Metal
Hitachi Vantara	Hitachi Advanced Server DS220 G2	A100, A30, T4	Bare Metal, vSphere ³
HPE	ProLiant DL360 Gen10 Plus	T4	Bare Metal
HPE	ProLiant DL380 Gen10	A100, T4	Bare Metal, vSphere ³
HPE	ProLiant DL380 Gen10 Plus	A100, A40, A30, A10, T4	Bare Metal, vSphere ³
HPE	ProLiant DL380a Gen11	H100	Bare Metal
HPE	ProLiant DL385 Gen10 Plus V2	A100, A40, T4	Bare Metal
HPE	Apollo 2000 (XL290 Gen10 Plus)	A100, T4	Bare Metal
HPE	Apollo 6500 (XL270d Gen10)	A100	Bare Metal, vSphere ³

Partner	Data Center Server	Supported NVIDIA GPUs ¹	NVIDIA AI Enterprise Compatible ²
HPE	Apollo 6500 (XL645d Gen10 Plus)	A100, A40, A10, HGX A100 - 4 GPU	Bare Metal, vSphere ³
HPE	Apollo 6500 (XL675d Gen10 Plus)	A100, A40, A10, HGX A100 - 8 GPU	Bare Metal, vSphere ³
HPE	Cray XD665 Gen11⁴	HGX H100 - 4 GPU	Bare Metal
HPE	Superdome Flex 280	A100	Bare Metal
Inspur Electronic Information	NF5280M6	A100, A30, A10, T4	Bare Metal, vSphere ³
Inspur Electronic Information	NF5468A5	A100, A40, A30, T4	Bare Metal, vSphere ³
Inspur Electronic Information	NF5468M5	A100, A40, A30, T4	Bare Metal
Inspur Electronic Information	NF5468M6	A100, A40, A30, T4	Bare Metal, vSphere ³
Inspur Electronic Information	NF5488A5	HGX A100 - 8 GPU	Bare Metal
Inspur Electronic Information	NF5688M6	HGX A100 - 8 GPU	Bare Metal
Koi Computers	XG-EA21-2UDPE10	A100, A40	Bare Metal
Koi Computers	XG-EG21-2UDP893	A100, A40	Bare Metal
Koi Computers	XG-EG21-2UDP896	A100	Bare Metal
Koi Computers	XG-EG21-2UDP940	A100, A30	Bare Metal
Koi Computers	XG-EG21-2UDP944	A100	Bare Metal
Koi Computers	XG-EG21-2UUP411	A100	Bare Metal
Koi Computers	XG-EG21-2UUP920	A100	Bare Metal
Koi Computers	XG-EG21-2UUP924	A100, A30	Bare Metal
Koi Computers	XG-EG21-4UDP854	A100, A30	Bare Metal
Koi Computers	XG-EG21-4UDP951	A100, A30	Bare Metal
Koi Computers	XG-XA21-2UDPE10 / XG-XA21-2UDP10S	A100	Bare Metal
Koi Computers	XG-XG21-2UDP830	A100, A40	Bare Metal
Koi Computers	XG-XG21-4UDP980	A100	Bare Metal
Koi Computers	XG-XG21-4UDP9A0	A40	Bare Metal
Lambda	Hyperplane 8-A100	HGX A100 - 8 GPU	Bare Metal
Lambda	Hyperplane 4-A100	HGX A100 - 4 GPU	Bare Metal
Lambda	Scalar 8 AMD g	A100, A30	Bare Metal
Lambda	Scalar 2U	A100	Bare Metal
Lambda	Scalar 8 AMD	A100, A30, A40	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs ¹	NVIDIA AI Enterprise Compatible ²
Lanner Electronics	FX-3420	A30	Bare Metal
Leadtek	WinFast GS2040T	A100, RTX A6000	Bare Metal
Leadtek	WinFast GS2045T	A100	Bare Metal
Leadtek	WinFast GS4840	A100	Bare Metal
Leadtek	WinFast GS4845	A100, RTX A6000	Bare Metal
Leadtek	WinFast GS4850T	RTX 6000 Ada	Bare Metal
Leadtek	WinFast WS2040	A100, RTX A6000	Bare Metal
Lenovo	SD665-N-V3	HGX H100 – 4 GPU	Bare Metal
Lenovo	SR670 V2	H100, L40, L4, A100, A40, A30, T4	Bare Metal, vSphere
Lenovo	ThinkSystem SR630 V2	L4	Bare Metal
Lenovo	ThinkSystem SR630 V3 Rack Server	L4	Bare Metal
Lenovo	ThinkSystem SR650 V2 Rack Server	L40, L4, RTX A6000, A100, A40, A30	Bare Metal
Lenovo	ThinkSystem SR650 V3 Rack Server	H100, L40, L4, A100, A40, RTX A6000	Bare Metal
Lenovo	ThinkSystem SR655 V3 Server	H100, RTX A6000, RTX A4500, A100, A40, A2	Bare Metal
Lenovo	ThinkSystem SR665	L40, L4, A100, A40, A30	Bare Metal, vSphere ³
Lenovo	ThinkSystem SR665 V3 Server	L40S, L40	Bare Metal
Lenovo	ThinkSystem SR670	A100, A40	Bare Metal
Lenovo	ThinkSystem SR670 V2	L40, L4, A100, A40, A30, T4	Bare Metal, vSphere ³
Lenovo	ThinkSystem ST650 V2 Server	L4, RTX A6000	Bare Metal
NEC	Express5800/R120h-2M	A100, T4	Bare Metal, vSphere ³
NEC	Express5800/R120i-2M	A100, A40	Bare Metal, vSphere ³
NEC	NX7700x/A50102E-2 v2	A100, T4	Bare Metal, vSphere ³
NEC	NX7700x/A6010E-2	A100, A40	Bare Metal, vSphere ³
Nettrix	X620 G40	A100, A30	Bare Metal
Nettrix	X640 G40	A100, A40, A30	Bare Metal
Nettrix	X660 G45	HGX A100 - 8 GPU	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs ¹	NVIDIA AI Enterprise Compatible ²
Nutanix/Dell Technologies	XC7525-24	A100	Bare Metal, Nutanix
Nutanix/Dell Technologies	XC750xa	A100	Bare Metal, Nutanix
Nutanix/Dell Technologies	XC750-24	A100	Bare Metal, Nutanix
Nutanix/HPE	ProLiant DX380 Gen10 Plus	A100	Bare Metal, Nutanix
Nutanix/HPE	DX385 Gen10 Plus v2	A100	Bare Metal, Nutanix
Nutanix/SuperMicro	NX-3155GN-G8	A100	Bare Metal, Nutanix
Nutanix/Lenovo	HX5531	A100	Bare Metal, Nutanix
Nutanix/Lenovo	HX7531	A100	Bare Metal, Nutanix
Nutanix/Lenovo	HX650 V3 IS/CN	A100	Bare Metal, Nutanix
One Stop Systems	EOS 4a / EB4400/ SDS 4a / EB4400	A100	Bare Metal
QCT	QuantaGrid D43N-3U	A100, A30, A10, RTX A6000	Bare Metal, vSphere ³
QCT	QuantaGrid D53XQ-2U	A100, A30, T4	Bare Metal
QCT	QuantaGrid D74H-7U	HGX H100	Bare Metal
QCT	QuantaPlex T43Z-2U	T4	Bare Metal
Supermicro	A+ Server 2114GT-DNR	A100, A40, A10	Bare Metal
Supermicro	A+ Server AS-2124GQ-NART	HGX A100 - 4 GPU	Bare Metal, vSphere ³
Supermicro	A+ Server AS-4124GO-NART	HGX A100 - 8 GPU	Bare Metal
Supermicro	A+ Server AS-4124GS-TNR	A100, A30, A40	Bare Metal, vSphere ³
Supermicro	CloudDC SuperServer SYS-620C-TN12R	RTX A6000	Bare Metal
Supermicro	GPU Blade SBA-4119SG	A100, A40, A10, T4	Bare Metal
Supermicro	GPU SuperServer SYS-120GQ-TNRT	A100, A40, A30, A10, T4	Bare Metal
Supermicro	GPU SuperServer SYS-210GP-DNR	A100, A40, A30	Bare Metal
Supermicro	GPU SuperServer SYS-220GP-TNR	A100, A40, A30, A10, T4	Bare Metal, vSphere ³
Supermicro	GPU SuperServer SYS-220GQ-TNAR+	HGX A100 - 4 GPU	Bare Metal, vSphere ³
Supermicro	GPU SuperServer SYS-420GP-TNAR	HGX A100 - 8 GPU	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs ¹	NVIDIA AI Enterprise Compatible ²
Supermicro	GPU SuperServer SYS-420GP-TNR	A100, A40, A30, A10, T4	Bare Metal
Supermicro	GPU SuperServer SYS-421GU-TNXR	HGX H100	Bare Metal
Supermicro	GPU SuperWorkstation SYS-740GP-TNRT	A100, A40, A30, A10, T4	Bare Metal
Supermicro	SYS-751GE-TNRT	A100	Bare Metal
Supermicro	SYS-821GE-TNHR	HGX H100 - 8 GPU	Bare Metal
Supermicro	Ultra SuperServer SYS-120U-TNR	A100, A30, T4	Bare Metal, vSphere ³
Supermicro	Ultra SuperServer SYS-220U-TNR	A100, A40, A30, A10, T4	Bare Metal, vSphere ³
xFusion	FusionServer 2288H V6	A100, A30	Bare Metal

**Note:**

1. NVIDIA-Certified systems are tested with each supported NVIDIA GPU to ensure that they can work together to provide the highest level of performance and reliability. The supported GPUs are certified for Data Center Servers for every generally available GPU version. For example, every system validated for the A100 GPU is certified for both the A100 40 GB and A100 80 GB GPU versions.
2. Only systems equipped with eligible NVIDIA GPUs are NVIDIA AI Enterprise Compatible. See the [NVIDIA AI Enterprise technical documentation](#) for the list of eligible GPUs and more information.
3. The specific NVIDIA GPUs that are supported by vSphere on each partner system are listed on the [VMWare Compatibility Guide](#).
4. Modular system with built-in OEM networking.

Chapter 6. List of NVIDIA-Certified Systems - Workstations

The following systems have been validated by NVIDIA as NVIDIA-Certified Systems - Workstations.

NVIDIA-Certified workstations provide the hardware foundation of [Data Science Workstations](#) that combine optimized hardware with a data science software stack built on NVIDIA CUDA-X AI. They utilize enterprise-class GPUs to accelerate data science workloads, and are validated for optimal performance, reliability, and compatibility with NVIDIA software and services.

Partner	Workstation	Supported NVIDIA GPUs ¹
Advantech	ACP-4340	RTX 6000 Ada
ASUS	ExpertCenter E500 G9	RTX A2000
BOXX	APEXX A3.02	RTX A6000, RTX A5000
BOXX	APEXX A3.03	RTX A4000, RTX A2000
BOXX	APEXX E3.04	RTX A4000, RTX A2000
BOXX	APEXX S3.06	RTX A6000
BOXX	APEXX S3.07	RTX A6000, RTX A5000, RTX A4500, RTX A4000
BOXX	APEXX S3.08	RTX 6000 ADA, RTX A6000
BOXX	APEXX T3.05	RTX A6000, RTX A5000
BOXX	APEXX T3.07	RTX 6000 ADA, RTX A6000, RTX A2000
BOXX	APEXX T4L	RTX A6000
BOXX	APEXX W3L	RTX A6000, RTX A5000
BOXX	APEXX W4L	RTX A6000, RTX A5000
BOXX	FLEXX S1G	RTX 6000 ADA
BOXX	FLEXX S2G	RTX A6000, RTX A4000
BOXX	RAXX P3	RTX A6000, RTX A5000
BOXX	RAXX T3	RTX A6000
Dell Technologies	Precision 3650 Tower Workstation	RTX A5000

Partner	Workstation	Supported NVIDIA GPUs ¹
Dell Technologies	Precision 5820 Tower	RTX A5000, RTX A4000
Dell Technologies	Precision 5860 Tower	RTX 6000 Ada
HP	Z2 G8 Tower Workstation	RTX A5000, RTX A4000
HP	Z4 G4 Workstation	RTX A6000, RTX A5000, RTX A4000
HP	ZCentral 4R	RTX A6000, RTX A5000, RTX A4000
HP	Z6 G4 Workstation	RTX A6000, RTX A4000
HP	Z8 G4 Workstation	RTX A6000, RTX A5000, RTX A4000
Leadtek	WinFast WS750	RTX A5500, RTX A4500
Leadtek	WinFast WS850	RTX 6000 Ada
Leadtek	WinFast WS940	RTX A6000, RTX A5500, RTX A5000, RTX A4500, RTX A4000
Leadtek	WinFast WS945	RTX 6000 Ada, RTX A6000, RTX A5500, RTX A5000, RTX A4500, RTX A4000
Leadtek	WinFast WS1040	RTX 6000 Ada, RTX A6000, RTX A5500, RTX A5000, RTX A4500, RTX A4000
Leadtek	WinFast WS1050	RTX 6000 Ada
Leadtek	WinFast WS2040	RTX A6000, RTX A5500, RTX A5000, RTX A4500, RTX A4000
Leadtek	WinFast WS2050	RTX 6000 Ada, RTX 5000 Ada
Lenovo	ThinkPad P16	RTX A5500
Lenovo	ThinkStation P350 Tower	RTX A5000, RTX A4000
Lenovo	ThinkStation P520 Workstation	RTX A6000, RTX A5500, RTX A5000, RTX A4000
Lenovo	ThinkStation P520c Workstation	RTX A5000, RTX A4000
Lenovo	ThinkStation P620 Tower Workstation	RTX A6000, RTX A5500, RTX A5000, RTX A4000
Lenovo	ThinkStation P720 Tower Workstation	RTX A6000, RTX A5000, RTX A4000
Lenovo	ThinkStation P920 Tower Workstation	RTX A6000, RTX A5500, RTX A5000
Lenovo	ThinkSystem SR650 V3 Rack Server	RTX A4500
Lenovo	ThinkSystem SR655 V3 Rack Server	RTX A4500
Microway	Whisperstation	RTX A6000
Siemens	Simatic IPC847E	RTX A4000

Partner	Workstation	Supported NVIDIA GPUs ¹
Supermicro	A+ SuperWorkstation 5014A-TT	RTX A6000, RTX A5500, RTX A4500
Supermicro	A+ SuperWorkstation 5039A-i	RTX A6000
Supermicro	GPU SuperServer SYS-740GP-TNRT	RTX A6000, RTX A5000
Supermicro	UP Workstation SYS-531A-IL	RTX A6000, RTX A5500, RTX A4500, RTX A2000
Supermicro	UP Workstation SYS-540A-TR	RTX A6000, RTX A5500, RTX A4500, RTX A2000

**Note:**

1. NVIDIA-Certified systems are tested with each supported NVIDIA GPU to ensure that they can work together to provide the highest level of performance and reliability.

Chapter 7. List of NVIDIA-Certified Systems - Mobile Workstations

The following systems have been validated by NVIDIA as NVIDIA-Certified Systems - Mobile Workstations.

NVIDIA-Certified workstations provide the hardware foundation of [Data Science Workstations](#) that combine optimized hardware with a data science software stack built on NVIDIA CUDA-X AI. They utilize enterprise-class GPUs to accelerate data science workloads, and are validated for optimal performance, reliability, and compatibility with NVIDIA software and services.

Partner	Workstation	Supported NVIDIA GPUs ¹
BOXX	GoBOXX SLM 17	RTX A5500
Dell Technologies	Precision 7560 Workstation	RTX A5000, RTX A4000
Dell Technologies	Precision 7760 Workstation	RTX A5000, RTX A4000
Dell Technologies	Precision 7680	RTX 5000 Ada, RTX 4000 Ada,
Lenovo	ThinkPad P1 Gen 4 Mobile Workstation	RTX A5000, RTX A4000
Lenovo	ThinkPad P15 Gen 2 Mobile Workstation	RTX A5000, RTX A4000
Lenovo	ThinkPad P17 Gen 2 Mobile Workstation	RTX A5000, RTX A4000



Note:

1. NVIDIA-Certified systems are tested with each supported NVIDIA GPU to ensure that they can work together to provide the highest level of performance and reliability.

Chapter 8. List of NVIDIA-Certified Systems - High-Density VDI Servers

The following systems have been validated by NVIDIA as NVIDIA-Certified Systems - High-Density VDI Servers.

Partner	High-Density VDI Servers	Supported NVIDIA GPUs ¹
ATOS	BullSequana SA20G	A16
DataON	AZS-6208G	A16
Dell Technologies	PowerEdge R740/R740xd	A16
Dell Technologies	PowerEdge R750	A16
Dell Technologies	PowerEdge R750xa	A16
Dell Technologies	PowerEdge R760xa Rack Server	A16
Dell Technologies	PowerEdge R7615 Rack Server	A16
Dell Technologies	PowerEdge R7625 Rack Server	A16 (8 GPUs)
Dell Technologies	PowerEdge R7515	A16
Dell Technologies	PowerEdge R7525	A16
Dell Technologies	PowerEdge R760	A16
Dell Technologies	PowerEdge R960 Server Node	A16
GIGABYTE	E152-ZE0	A16
GIGABYTE	E162-220	A16
GIGABYTE	G242-Z12	A16
GIGABYTE	G292-280	A16
GIGABYTE	G292-Z24	A16
GIGABYTE	G292-Z44	A16
GIGABYTE	G292-Z45	A16
GIGABYTE	G492-H80	A16
GIGABYTE	G492-HA0	A16
GIGABYTE	R282-Z93	A16

Partner	High-Density VDI Servers	Supported NVIDIA GPUs ¹
Hitachi Vantara	Hitachi Advanced Server DS220 G2	A16
HPE	ProLiant DL380 Gen10	A16
HPE	ProLiant DL380 Gen10 Plus	A16
HPE	ProLiant DL385 Gen10 Plus	A16
QCT	QuantaGrid D43N-3U	A16
QCT	QuantaGrid D53XQ-2U	A16
Supermicro	SYS-740GP-TNRT	A16



Note:

1. NVIDIA-Certified systems are tested with each supported NVIDIA GPU to ensure that they can work together to provide the highest level of performance and reliability.

Chapter 9. List of NVIDIA-Certified Systems - Edge Systems

The following servers and workstations have been validated by NVIDIA as NVIDIA-Certified Systems - Edge Systems.

Partner	Server / Workstation	Supported NVIDIA GPUs ¹	Enterprise Edge Server ²	Industrial Edge Workstation ³	NVIDIA AI Enterprise Compatible ⁴
Aaeon	BOXER-8332	A2		Yes	Bare Metal
Advantech	ACP-2020G	A2		Yes	Bare Metal
Advantech	EPC-B3522	RTX A4500		Yes	
Advantech	EPC-B3588	RTX A4500		Yes	
Advantech	EPC-B5587	RTX A6000	Yes		Bare Metal
Advantech	MIC-770V2	RTX A2000		Yes	Bare Metal
Advantech	MIC-770V3	L4		Yes	Bare Metal
Advantech	SKY-640V2	A30	Yes		Bare Metal
Advantech	USM-501	RTX A6000		Yes	Bare Metal
Aetina	AIP-FQ47-B1	RTX A4500		Yes	
Aetina	AIP-FR68-A1	RTX A6000		Yes	Bare Metal
Aetina	AIS-D422-A1	L4		Yes	Bare Metal
Aewin	SCB-1932c	A30		Yes	Bare Metal
Axiomtek	IPC-962-525	A2	Yes	Yes	Bare Metal
DataON	AZS-6208G	A40, A30, A2	Yes	Yes	Bare Metal
Dell Technologies	PowerEdge C6220 Server Node	A2	Yes		Bare Metal
Dell Technologies	PowerEdge R640	A2	Yes		Bare Metal, vSphere ⁵
Dell Technologies	PowerEdge R6515	A2	Yes		Bare Metal, vSphere ⁵

Partner	Server / Workstation	Supported NVIDIA GPUs ¹	Enterprise Edge Server ²	Industrial Edge Workstation ³	NVIDIA AI Enterprise Compatible ⁴
Dell Technologies	PowerEdge R650	L4, A2, T4	Yes		Bare Metal, vSphere ⁵
Dell Technologies	PowerEdge R660 Rack Server	L4, A2	Yes		Bare Metal
Dell Technologies	PowerEdge R6615 Rack Server	L4, A2	Yes		Bare Metal
Dell Technologies	PowerEdge R6625 Rack Server	L4	Yes		Bare Metal
Dell Technologies	PowerEdge R760xa Rack Server	L40, L4, A100, A40, A2	Yes		Bare Metal
Dell Technologies	PowerEdge R740/R740xd	A100, A40, A30, A2, T4	Yes		Bare Metal, vSphere ⁵
Dell Technologies	PowerEdge R750 Rack Server	L40, L4, A100, A40, A30, A2, T4	Yes		Bare Metal, vSphere ⁵
Dell Technologies	PowerEdge R750XA Rack Server	L40, L4, A2	Yes		Bare Metal, vSphere ⁵
Dell Technologies	PowerEdge R7515 Rack Server	L4	Yes		Bare Metal
Dell Technologies	PowerEdge R7525 Rack Server	L40, L4, A100, A40, A30, A2, T4	Yes		Bare Metal, vSphere ⁵
Dell Technologies	PowerEdge R760 Rack Server	L40, L4, H100, A100, A40, A30, A2	Yes		Bare Metal
Dell Technologies	PowerEdge R760xd2 Rack Server	L4	Yes		Bare Metal
Dell Technologies	PowerEdge R760xs Rack Server	A2	Yes		Bare Metal
Dell Technologies	PowerEdge R7615 Rack Server	L40, L4, A100, A40, A30, A2	Yes		Bare Metal

Partner	Server / Workstation	Supported NVIDIA GPUs ¹	Enterprise Edge Server ²	Industrial Edge Workstation ³	NVIDIA AI Enterprise Compatible ⁴
Dell Technologies	PowerEdge R7625 Rack Server	L40, L4, H100, A100, A40, A30	Yes		Bare Metal
Dell Technologies	PowerEdge R960 Rack Server	A100	Yes		Bare Metal
Dell Technologies	PowerEdge XR11 Rack Server	L4	Yes		Bare Metal
Dell Technologies	PowerEdge XR12 Rack Server	L4	Yes		Bare Metal
Dell Technologies	PowerEdge XR4520C	L4, A30, A2	Yes	Yes	Bare Metal
Dell Technologies	PowerEdge XR5610 Rack Server	L4, A2	Yes		Bare Metal
Dell Technologies	PowerEdge XR7620 Rack Server	L4, A100, A30, A2	Yes		Bare Metal
Dell Technologies	PowerEdge XR8620t Compute Sled	L4	Yes		Bare Metal
Fujitsu	PRIMERGY RX2540 M6	L4, A100, A40, A30, A2, T4	Yes		Bare Metal, vSphere ⁵
Fujitsu	PRIMERGY TX1330 M5	L4, A2	Yes		Bare Metal
GIGABYTE	E152-ZEO	A2	Yes		Bare Metal
GIGABYTE	E162-220	A2	Yes		Bare Metal
GIGABYTE	G242-P33	A100	Yes		Bare Metal
Hitachi Vantara	Hitachi Advanced Server DS220 G2	A2	Yes		Bare Metal
HPE	ProLiant DL110 Gen10 Plus	A2		Yes	Bare Metal
HPE	ProLiant DL360 Gen10	A2	Yes		Bare Metal
HPE	ProLiant DL360 Gen10 Plus	A2, T4	Yes		Bare Metal

Partner	Server / Workstation	Supported NVIDIA GPUs ¹	Enterprise Edge Server ²	Industrial Edge Workstation ³	NVIDIA AI Enterprise Compatible ⁴
HPE	ProLiant DL380 Gen10	A100, A2, T4	Yes		Bare Metal, vSphere ⁵
HPE	ProLiant DL380 Gen10 Plus	A100, A40, A30, T4	Yes		Bare Metal, vSphere ⁵
HPE	ProLiant DL385 Gen10 Plus V2	A100, A40, A2, T4	Yes		Bare Metal, vSphere ⁵
Inspur Electronic Information	NF5280M6	A2	Yes		Bare Metal
Kontron	CG2400	A30		Yes	Bare Metal
Kontron	ME1310	L4		Yes	Bare Metal
Lanner Electronics	ECA-5540	L4		Yes	
Lanner Electronics	FX-3420	A30	Yes		Bare Metal
Lanner Electronics	LEC-2290E	A2		Yes	Bare Metal
Lenovo	ThinkSystem SE350 Edge Server	A2	Yes	Yes	Bare Metal
Lenovo	ThinkEdge SE450 Edge Server	L40, L4, A30, A2	Yes	Yes	Bare Metal
Lenovo	ThinkEdge SE360 V2	A2		Yes	Bare Metal
Lenovo	ThinkSystem SR630 V3 Rack Server	A2	Yes		
Lenovo	ThinkSystem SR650 V2 Rack Server	A100, A40, A30, RTX A6000	Yes		Bare Metal
Lenovo	ThinkSystem SR650 V3 Rack Server	A2	Yes		
Lenovo	ThinkSystem SR665	A100, A40, A30	Yes		Bare Metal, vSphere ⁵
Mercury	RESX07-2U22R	A100		Yes	Bare Metal
Neousys	Nuvo-7166GC	A2		Yes	Bare Metal
Nutanix/Dell	XC7525-24	A100	Yes		Bare Metal, Nutanix

Partner	Server / Workstation	Supported NVIDIA GPUs ¹	Enterprise Edge Server ²	Industrial Edge Workstation ³	NVIDIA AI Enterprise Compatible ⁴
Nutanix/Dell	XC750-24	A100	Yes		Bare Metal, Nutanix
Nutanix/HPE	ProLiant DX380 Gen10 Plus	A100	Yes		Bare Metal, Nutanix
Nutanix/HPE	DX385 Gen10 Plus v2	A100	Yes		Bare Metal, Nutanix
One Stop Systems	Rigel Edge Supercomputer	HGX A100 - 4 GPU		Yes	Bare Metal
QCT	QuantaGrid D53X-1U	A2	Yes		Bare Metal
Siemens AG	IPC647E	A2, RTX A6000		Yes	Bare Metal
Supermicro	CloudDC SuperServer SYS-120C-TN10R	A2	Yes		Bare Metal
Supermicro	IoT SuperServer SYS-220HE-FTNR	A40	Yes		Bare Metal
Supermicro	IoT SuperServer SYS-221HE-FTNRD	H100	Yes		Bare Metal
Supermicro	IoT SuperServer SYS-E403-12P-FN2T	A30	Yes		Bare Metal


Note:

1. NVIDIA-Certified systems are tested with each supported NVIDIA GPU to ensure that they can work together to provide the highest level of performance and reliability.
2. Enterprise Edge servers are tested for remote systems management using the industry standard Redfish API.
3. Industrial Edge workstations are tested at elevated ambient temperatures to demonstrate their ability to function effectively outside a typical datacenter environment. Please refer to your hardware provider for the specific NVIDIA GPUs that are supported are for this use case.
4. Only systems equipped with eligible NVIDIA GPUs are NVIDIA AI Enterprise Compatible. See the [NVIDIA AI Enterprise technical documentation](#) for the list of eligible GPUs and more information.
5. The NVIDIA GPUs that are supported by vSphere on each partner system are listed on the [VMWare Compatibility Guide](#).

Chapter 10. NVIDIA-Certified Systems Supported Software

NVIDIA-Certified Systems Testing Software Environment

NVIDIA-Certified systems are tested in a standardized software environment that provides the highest levels of stability and performance. The current software test environment for NVIDIA-Certified systems uses production versions of the following:

- ▶ Ubuntu 20.04
- ▶ NVIDIA drivers
- ▶ MLNX_OFED network adaptor drivers
- ▶ NVIDIA Cloud Native Core with NVIDIA GPU Operator and Kubernetes

NVIDIA also supports its AI software running on Red Hat Enterprise Linux and CentOS. More details about OS support can be found at <https://docs.nvidia.com/deeplearning/frameworks/support-matrix/index.html>.

List of NVIDIA-Certified Systems Supported Software

NVIDIA supports the following software running on NVIDIA-Certified systems:

- ▶ Containers in general availability published by NVIDIA on the [NGC Catalog](#) including:
 - ▶ TensorFlow
 - ▶ PyTorch
 - ▶ TensorRT
 - ▶ Triton Inference Server
 - ▶ CUDA
 - ▶ RAPIDS
 - ▶ Clara Train
 - ▶ NVIDIA TAO
 - ▶ NVIDIA Riva
 - ▶ NVIDIA Omniverse

Free online support for NVIDIA's NGC containers is available on the [NVIDIA Developer Forums](#).

Enterprise support subscriptions for NVIDIA-Certified 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

© 2023 NVIDIA CORPORATION & AFFILIATES. All rights reserved.

