



# 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 - Data Center Servers.....	5
Chapter 5. List of NVIDIA-Certified Systems - Workstations.....	13
Chapter 6. List of NVIDIA-Certified Systems - Mobile Workstations.....	15
Chapter 7. List of NVIDIA-Certified Systems - High-Density VDI Servers.....	16
Chapter 8. List of NVIDIA-Certified Systems - Edge Systems.....	18
Chapter 9. NVIDIA-Certified Systems Supported Software.....	21

---

# 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 - 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 <sup>1</sup>	NVIDIA AI Enterprise Compatible <sup>2</sup>
Advantech	<a href="#">SKY-640V2</a>	A100, A30	Bare Metal
Altos Computing	<a href="#">BrainSphere R685 F5</a>	RTX A6000, A40	Bare Metal
AMAX	<a href="#">AceleMax DGS- 140W</a>	A100, A30, A40	Bare Metal
AMAX	<a href="#">AceleMax DGS- 214A</a>	A100, A30, A40	Bare Metal
AMAX	<a href="#">AceleMax DGS- 224AS</a>	HGX A100 - 4 GPU	Bare Metal
AMAX	<a href="#">AceleMax DGS- 260W</a>	A100, A30, A40	Bare Metal
AMAX	<a href="#">AceleMax DGS- 410W</a>	A100, A30, A40	Bare Metal
AMAX	<a href="#">AceleMax DGS- 428A</a>	A100, A30, A40	Bare Metal
AMAX	<a href="#">AceleMax DGS- 428AS</a>	HGX A100 - 8 GPU	Bare Metal
AMAX	<a href="#">AceleMax DGS- 428WS</a>	HGX A100 - 8 GPU	Bare Metal
AMAX	<a href="#">AceleMax DL- E440W</a>	A100, A30, A40	Bare Metal
ASRock Rack	<a href="#">1U4G-ROME</a>	A100, A30	Bare Metal
ASRock Rack	<a href="#">2U-Open19-m3.xlarge.x86</a>	A100	Bare Metal
ASRock Rack	<a href="#">2U-Open19-n3.xlarge.x86</a>	A100	Bare Metal
ASRock Rack	<a href="#">2U4G-ROME/2T</a>	A100, A30	Bare Metal
ASRock Rack	<a href="#">4U8G-ROME2/2T</a>	A100	Bare Metal
ASRock Rack	<a href="#">4U10G-ICX2/2T</a>	A100	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs <sup>1</sup>	NVIDIA AI Enterprise Compatible <sup>2</sup>
ASRock Rack	<a href="#">4U10G-ROME2/2T</a>	A100	Bare Metal
ASUS	<a href="#">ESC4000-E10 / ESC4000-E10S</a>	A100	Bare Metal, vSphere <sup>3</sup>
ASUS	<a href="#">ESC4000A-E10</a>	A100, A30, A40, RTX A6000	Bare Metal, vSphere <sup>3</sup>
ASUS	<a href="#">ESC8000A-E11</a>	A100	Bare Metal
ASUS	<a href="#">ESC8000A-E11</a>	A40, A30	Bare Metal, vSphere <sup>3</sup>
ASUS	<a href="#">ESC8000A-E12</a>	H100	Bare Metal
Atipa Technologies	<a href="#">Altezza SX227-32G3</a>	A100, A40	Bare Metal
Atipa Technologies	<a href="#">Altezza SX427-32G10</a>	A100, A40	Bare Metal
Atipa Technologies	<a href="#">Altezza SX427-32G8</a>	A100, A40	Bare Metal
Atipa Technologies	<a href="#">Altezza SX427-32G8SXM4</a>	HGX A100 - 8 GPU	Bare Metal
Atipa Technologies	<a href="#">Procyon SE218-8G4</a>	A100, T4	Bare Metal
Atipa Technologies	<a href="#">Procyon SE218-8G8</a>	A100, A30	Bare Metal
Atipa Technologies	<a href="#">Procyon SE228-16G4SXM4</a>	HGX A100 - 4 GPU	Bare Metal
Atipa Technologies	<a href="#">Procyon SE228-16G8</a>	A100, A40, A30	Bare Metal
Atipa Technologies	<a href="#">Procyon SE228-32G3</a>	A100	Bare Metal
Atipa Technologies	<a href="#">Procyon SE428-32G10</a>	A100, A30	Bare Metal
Atipa Technologies	<a href="#">Procyon SE428-32G8</a>	A100, A30, A40	Bare Metal
Atipa Technologies	<a href="#">Procyon SE428-32G8SXM4</a>	HGX A100 - 8 GPU	Bare Metal
ATOS	<a href="#">BullSequana SA20G</a>	A100	Bare Metal, vSphere <sup>3</sup>
ATOS	<a href="#">BullSequana X410-A5</a>	A100	Bare Metal
Boston	<a href="#">ANNA Ampere L1</a>	HGX A100 - 4 GPU	Bare Metal
Boston	<a href="#">ANNA Ampere M1</a>	A100, A40, A30	Bare Metal
Boston	<a href="#">ANNA Ampere S2</a>	A100, A40	Bare Metal
Boston	<a href="#">ANNA Ampere XL1</a>	A100, A40, A30	Bare Metal
BOXX	<a href="#">RAXX P4G</a>	A40	Bare Metal
Cisco	<a href="#">UCS C240 M6 Rack Server</a>	A100, A30	Bare Metal, vSphere <sup>3</sup>
Colfax	<a href="#">Colfax CX2460s-EK8 2U Rackmount Server</a>	HGX A100 - 4 GPU	Bare Metal
Colfax	<a href="#">Colfax CX41060s-XK8 4U Rackmount Server</a>	A100, A30, A40	Bare Metal



Partner	Data Center Server	Supported NVIDIA GPUs <sup>1</sup>	NVIDIA AI Enterprise Compatible <sup>2</sup>
Colfax	<a href="#">Colfax CX4860s-EK8 4U Rackmount Server</a>	A100, A30, A40	Bare Metal
Colfax	<a href="#">Colfax CX4880s-EK8 4U Rackmount Server</a>	HGX A100 - 8 GPU	Bare Metal
Colfax	<a href="#">Colfax CX4880s-XK8 4U Rackmount Server</a>	HGX A100 - 8 GPU	Bare Metal
Colfax	<a href="#">Colfax ProEdge SXP9000 Workstation</a>	A100, A30, A40	Bare Metal
Dell	<a href="#">DSS8440</a>	A100, A40, T4	Bare Metal, vSphere <sup>3</sup>
Dell	<a href="#">PowerEdge R650</a>	T4	Bare Metal
Dell	<a href="#">PowerEdge R7515</a>	A30	Bare Metal
Dell	<a href="#">PowerEdge R7525</a>	H100, A100, A30, A40, A10, T4	Bare Metal, vSphere <sup>3</sup>
Dell	<a href="#">PowerEdge R740/ R740xd</a>	A100, A30, A40, A10, T4	Bare Metal, vSphere <sup>3</sup>
Dell	<a href="#">PowerEdge R750</a>	H100, A100, A30, A40, A10, T4	Bare Metal, vSphere <sup>3</sup>
Dell	<a href="#">PowerEdge R750xa</a>	H100, A100, A30, A40, A10, T4	Bare Metal, vSphere <sup>3</sup>
Dell	<a href="#">PowerEdge XE8545</a>	HGX A100 - 4 GPU	Bare Metal, vSphere <sup>3</sup>
Dell	<a href="#">PowerEdge XR11</a>	T4	Bare Metal
Dell	<a href="#">PowerEdge XR12</a>	A100, A40, A30, A10, T4	Bare Metal, vSphere <sup>3</sup>
Dell	<a href="#">VxRail V570 / V570F</a>	A100, A40, A30, T4	Bare Metal, vSphere <sup>3</sup>
Dell	<a href="#">VxRail V670F</a>	A100, A40, A30, T4	Bare Metal, vSphere <sup>3</sup>
Exxact	<a href="#">TensorEX TS2-185671979</a>	A100, A30, A40	Bare Metal
Exxact	<a href="#">TensorEX TS2-169843133</a>	A100	Bare Metal
Exxact	<a href="#">TensorEX TS2-150341732</a>	A100, A30	Bare Metal
Exxact	<a href="#">TensorEX TS4-194492555</a>	A100, A30, A40	Bare Metal
Exxact	<a href="#">TensorEX TS2-184920243</a>	A100	Bare Metal
Exxact	<a href="#">TensorEX TS2-171138844</a>	HGX A100 - 4 GPU	Bare Metal

<b>Partner</b>	<b>Data Center Server</b>	<b>Supported NVIDIA GPUs <sup>1</sup></b>	<b>NVIDIA AI Enterprise Compatible <sup>2</sup></b>
Exxact	<a href="#">TensorEX TS4-195183185</a>	HGX A100 - 8 GPU	Bare Metal
Exxact	<a href="#">TensorEX TS4-173535991</a>	A100, A30, A40	Bare Metal
Exxact	<a href="#">TensorEX TS1-147097335</a>	A100, A30, A40	Bare Metal
Exxact	<a href="#">TensorEX TS2-144629457</a>	A100, A30, A40	Bare Metal
Exxact	<a href="#">TensorEX TS4-168747704</a>	HGX A100 - 8 GPU	Bare Metal
Exxact	<a href="#">TensorEX TS4-133524070</a>	A100, A30, A40	Bare Metal
Exxact	<a href="#">TensorEX TWS-115999024</a>	A100, A30, A40	Bare Metal
Fujitsu	<a href="#">PRIMERGY GX2460 M1</a>	A100, A30	Bare Metal, vSphere <sup>3</sup>
Fujitsu	<a href="#">PRIMERGY GX2570 M6</a>	HGX A100 - 8 GPU	Bare Metal
Fujitsu	<a href="#">PRIMERGY RX2540 M5</a>	A100, T4	Bare Metal, vSphere <sup>3</sup>
Fujitsu	<a href="#">PRIMERGY RX2540 M6</a>	A100, A100X, A40, A30, A30X, RTX A6000	Bare Metal, vSphere <sup>3</sup>
GIGABYTE	<a href="#">E152-ZEO</a>	A100	Bare Metal
GIGABYTE	<a href="#">E162-Z20</a>	A100, A40	Bare Metal
GIGABYTE	<a href="#">G242-Z11</a>	A100, T4	Bare Metal
GIGABYTE	<a href="#">G242-Z12</a>	A100, A40	Bare Metal
GIGABYTE	<a href="#">G262-ZR0</a>	HGX A100 - 4 GPU	Bare Metal
GIGABYTE	<a href="#">G292-Z20</a>	NVIDIA A100 for PCIe	Bare Metal
GIGABYTE	<a href="#">G292-Z24</a>	A100, A30	Bare Metal
GIGABYTE	<a href="#">G292-Z40</a>	A100, A40, A30	Bare Metal
GIGABYTE	<a href="#">G292-Z44</a>	A100, A40, A10	Bare Metal
GIGABYTE	<a href="#">G292-280</a>	A100, A40	Bare Metal
GIGABYTE	<a href="#">G482-Z54</a>	A100, A30, A40	Bare Metal
GIGABYTE	<a href="#">G492-HA0</a>	A100, A40	Bare Metal
GIGABYTE	<a href="#">G492-H80</a>	A100, A40	Bare Metal
GIGABYTE	<a href="#">G492-ID0</a>	HGX A100 - 8 GPU	Bare Metal
GIGABYTE	<a href="#">G492-ZD2</a>	HGX A100 - 8 GPU	Bare Metal
GIGABYTE	<a href="#">G492-Z51</a>	A100, A30	Bare Metal
GIGABYTE	<a href="#">R262-ZA2</a>	A100, RTX A6000	Bare Metal
GIGABYTE	<a href="#">R281-G30</a>	A100, T4	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs <sup>1</sup>	NVIDIA AI Enterprise Compatible <sup>2</sup>
GIGABYTE	<a href="#">R282-G30</a>	A100, A40	Bare Metal
GIGABYTE	<a href="#">R282-Z93</a>	A100, A40	Bare Metal, vSphere <sup>3</sup>
GIGABYTE	<a href="#">R282-Z96</a>	A100, A10	Bare Metal
H3C	<a href="#">UniServer R4900 G5</a>	A100, A30, A10	Bare Metal
H3C	<a href="#">UniServer R5300 G5</a>	A100, A40, A30	Bare Metal, vSphere <sup>3</sup>
H3C	<a href="#">UniServer R5500 G5 AMD</a>	HGX A100 - 8 GPU	Bare Metal
H3C	<a href="#">UniServer R5500 G5 Intel</a>	HGX A100 - 8 GPU	Bare Metal
Hitachi	<a href="#">HA8000V/DL360 Gen10</a>	T4	Bare Metal
Hitachi	<a href="#">HA8000V/DL360 Gen10 Plus</a>	A100, T4	Bare Metal
Hitachi	<a href="#">HA8000V/DL380 Gen10 Plus</a>	A100, A40, A30, A10, T4	Bare Metal, vSphere <sup>3</sup>
Hitachi Vantara	<a href="#">Hitachi Advanced Server DS120 G2</a>	T4	Bare Metal
Hitachi Vantara	<a href="#">Hitachi Advanced Server DS220 G2</a>	A100, T4	Bare Metal, vSphere <sup>3</sup>
HPE	<a href="#">ProLiant DL360 Gen10 Plus</a>	T4	Bare Metal
HPE	<a href="#">ProLiant DL380 Gen10</a>	A100, T4	Bare Metal, vSphere <sup>3</sup>
HPE	<a href="#">ProLiant DL380 Gen10 Plus</a>	A100, A40, A30, A10, T4	Bare Metal, vSphere <sup>3</sup>
HPE	<a href="#">ProLiant DL385 Gen10 Plus V2</a>	A100, A40, T4	Bare Metal
HPE	<a href="#">Apollo 2000 (XL290 Gen10 Plus)</a>	A100, T4	Bare Metal
HPE	<a href="#">Apollo 6500 (XL270d Gen10)</a>	A100	Bare Metal, vSphere <sup>3</sup>
HPE	<a href="#">Apollo 6500 (XL645d Gen10 Plus)</a>	A100, A40, A10, HGX A100 - 4 GPU	Bare Metal
HPE	<a href="#">Apollo 6500 (XL675d Gen10 Plus)</a>	A100, A40, A10, HGX A100 - 8 GPU	Bare Metal, vSphere <sup>3</sup>
Inspur	<a href="#">NF5280M6</a>	A100, A30, A10, T4	Bare Metal, vSphere <sup>3</sup>
Inspur	<a href="#">NF5468A5</a>	A100, A40, A30, T4	Bare Metal, vSphere <sup>3</sup>
Inspur	<a href="#">NF5468M5</a>	A100, A40, A30, T4	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs <sup>1</sup>	NVIDIA AI Enterprise Compatible <sup>2</sup>
Inspur	<a href="#">NF5468M6</a>	A100, A40, A30, T4	Bare Metal, vSphere <sup>3</sup>
Inspur	<a href="#">NF5488A5</a>	HGX A100 - 8 GPU	Bare Metal
Inspur	<a href="#">NF5688M6</a>	HGX A100 - 8 GPU	Bare Metal
Koi Computers	<a href="#">XG-EA21-2UDPE10</a>	A100, A40	Bare Metal
Koi Computers	<a href="#">XG-EG21-2UDP893</a>	A100, A40	Bare Metal
Koi Computers	<a href="#">XG-EG21-2UDP896</a>	A100	Bare Metal
Koi Computers	<a href="#">XG-EG21-2UDP940</a>	A100, A30	Bare Metal
Koi Computers	<a href="#">XG-EG21-2UDP944</a>	A100	Bare Metal
Koi Computers	<a href="#">XG-EG21-2UUP411</a>	A100	Bare Metal
Koi Computers	<a href="#">XG-EG21-2UUP920</a>	A100	Bare Metal
Koi Computers	<a href="#">XG-EG21-2UUP924</a>	A100, A30	Bare Metal
Koi Computers	<a href="#">XG-EG21-4UDP854</a>	A100, A30	Bare Metal
Koi Computers	<a href="#">XG-EG21-4UDP951</a>	A100, A30	Bare Metal
Koi Computers	<a href="#">XG-XA21-2UDPE10 / XG-XA21-2UDP10S</a>	A100	Bare Metal
Koi Computers	<a href="#">XG-XG21-2UDP830</a>	A100, A40	Bare Metal
Koi Computers	<a href="#">XG-XG21-4UDP980</a>	A100	Bare Metal
Koi Computers	<a href="#">XG-XG21-4UDP9A0</a>	A40	Bare Metal
Lambda	<a href="#">Hyperplane 8-A100</a>	HGX A100 - 8 GPU	Bare Metal
Lambda	<a href="#">Hyperplane 4-A100</a>	HGX A100 - 4 GPU	Bare Metal
Lambda	<a href="#">Scalar 8 AMD g</a>	A100, A30	Bare Metal
Lambda	<a href="#">Scalar 2U</a>	A100	Bare Metal
Lambda	<a href="#">Scalar 8 AMD</a>	A100, A30, A40	Bare Metal
Lanner Electronics	<a href="#">FX-3420</a>	A30	Bare Metal
Leadtek	<a href="#">WinFast GS2040T</a>	A100, RTX A6000	Bare Metal
Leadtek	<a href="#">WinFast GS2045T</a>	A100	Bare Metal
Leadtek	<a href="#">WinFast GS4840</a>	A100	Bare Metal
Leadtek	<a href="#">WinFast GS4845</a>	A100, RTX A6000	Bare Metal
Leadtek	<a href="#">WinFast WS2040</a>	A100, RTX A6000	Bare Metal
Lenovo	<a href="#">ThinkSystem SR650 V2 Rack Server</a>	A100, A40, A30 , RTX A6000	Bare Metal
Lenovo	<a href="#">ThinkSystem SR665</a>	A100, A40, A30	Bare Metal, vSphere <sup>3</sup>
Lenovo	<a href="#">ThinkSystem SR670</a>	A100, A40	Bare Metal
Lenovo	<a href="#">ThinkSystem SR670 V2</a>	A100, A40, A30, T4	Bare Metal, vSphere <sup>3</sup>

Partner	Data Center Server	Supported NVIDIA GPUs <sup>1</sup>	NVIDIA AI Enterprise Compatible <sup>2</sup>
Lenovo	<a href="#">Lenovo ThinkSystem ST650 V2 Server</a>	RTX A6000	Bare Metal
NEC	<a href="#">Express5800/R120h-2M</a>	A100, T4	Bare Metal, vSphere <sup>3</sup>
NEC	<a href="#">Express5800/R120i-2M</a>	A100, A40	Bare Metal, vSphere <sup>3</sup>
NEC	<a href="#">NX7700x/A50102E-2 v2</a>	A100, T4	Bare Metal, vSphere <sup>3</sup>
NEC	<a href="#">NX7700x/A6010E-2</a>	A100, A40	Bare Metal, vSphere <sup>3</sup>
Nettrix	<a href="#">X620 G40</a>	A100, A30	Bare Metal
Nettrix	<a href="#">X640 G40</a>	A100, A40, A30	Bare Metal
Nettrix	<a href="#">X660 G45</a>	HGX A100 - 8 GPU	Bare Metal
One Stop Systems	<a href="#">EOS 4a / EB4400/ SDS 4a / EB4400</a>	A100	Bare Metal
QCT	<a href="#">QuantaGrid D43N-3U</a>	A100, A30, A10, RTX A6000	Bare Metal, vSphere <sup>3</sup>
QCT	<a href="#">QuantaGrid D53XQ-2U</a>	A100, A30, T4	Bare Metal
QCT	<a href="#">QuantaPlex T43Z-2U</a>	T4	Bare Metal
Supermicro	<a href="#">A+ Server 2114GT-DNR</a>	A100, A40, A10	Bare Metal
Supermicro	<a href="#">A+ Server AS-2124GQ-NART</a>	HGX A100 - 4 GPU	Bare Metal, vSphere <sup>3</sup>
Supermicro	<a href="#">A+ Server AS-4124GO-NART</a>	HGX A100 - 8 GPU	Bare Metal
Supermicro	<a href="#">A+ Server AS-4124GS-TNR</a>	A100, A30, A40	Bare Metal, vSphere <sup>3</sup>
Supermicro	<a href="#">CloudDC SuperServer SYS-620C-TN12R</a>	RTX A6000	Bare Metal
Supermicro	<a href="#">GPU Blade SBA-4119SG</a>	A100, A40, A10, T4	Bare Metal
Supermicro	<a href="#">GPU SuperServer SYS-120GQ-TNRT</a>	A100, A40, A30, A10, T4	Bare Metal
Supermicro	<a href="#">GPU SuperServer SYS-210GP-DNR</a>	A100, A40, A30	Bare Metal
Supermicro	<a href="#">GPU SuperServer SYS-220GP-TNR</a>	A100, A40, A30, A10, T4	Bare Metal, vSphere <sup>3</sup>
Supermicro	<a href="#">GPU SuperServer SYS-220GQ-TNAR+</a>	HGX A100 - 4 GPU	Bare Metal, vSphere <sup>3</sup>
Supermicro	<a href="#">GPU SuperServer SYS-420GP-TNAR</a>	HGX A100 - 8 GPU	Bare Metal
Supermicro	<a href="#">GPU SuperServer SYS-420GP-TNR</a>	A100, A40, A30, A10, T4	Bare Metal

Partner	Data Center Server	Supported NVIDIA GPUs <sup>1</sup>	NVIDIA AI Enterprise Compatible <sup>2</sup>
Supermicro	<a href="#">GPU SuperWorkstation SYS-740GP-TNRT</a>	A100, A40, A30, A10, T4	Bare Metal
Supermicro	<a href="#">Ultra SuperServer SYS-120U-TNR</a>	A100, A40, A30, A10, T4	Bare Metal, vSphere <sup>3</sup>
Supermicro	<a href="#">Ultra SuperServer SYS-220U-TNR</a>	A100, A40, A30, A10, T4	Bare Metal, vSphere <sup>3</sup>
xFusion	<a href="#">FusionServer 2288H V6</a>	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 NVIDIA GPUs that are supported by vSphere on each partner system are listed on the [VMWare Compatibility Guide](#).

# Chapter 5. 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 <sup>1</sup>
ASUS	<a href="#">ExpertCenter E500 G9</a>	RTX A2000
BOXX	<a href="#">APEXX A3.02</a>	RTX A6000, RTX A5000
BOXX	<a href="#">APEXX A3.03</a>	RTX A4000, RTX A2000
BOXX	<a href="#">APEXX E3.04</a>	RTX A4000, RTX A2000
BOXX	<a href="#">APEXX S3.06</a>	RTX A6000
BOXX	<a href="#">APEXX S3.07</a>	RTX A6000, RTX A5000, RTX A4500, RTX A4000
BOXX	<a href="#">APEXX S3.08</a>	RTX A6000
BOXX	<a href="#">APEXX T3.05</a>	RTX A6000, RTX A5000
BOXX	<a href="#">APEXX T3.07</a>	RTX A6000
BOXX	<a href="#">APEXX T4L</a>	RTX A6000
BOXX	<a href="#">APEXX W3L</a>	RTX A6000, RTX A5000
BOXX	<a href="#">APEXX W4L</a>	RTX A6000, RTX A5000
BOXX	<a href="#">FLEXX S2G</a>	RTX A6000, RTX A4000
BOXX	<a href="#">RAXX P3</a>	RTX A6000, RTX A5000
BOXX	<a href="#">RAXX T3</a>	RTX A6000
Dell	<a href="#">Precision 3650 Tower Workstation</a>	RTX A5000
Dell	<a href="#">Precision 5820 Tower</a>	RTX A5000, RTX A4000
HP	<a href="#">Z2 G8 Tower Workstation</a>	RTX A5000, RTX A4000

Partner	Workstation	Supported NVIDIA GPUs <sup>1</sup>
HP	<a href="#">Z4 G4 Workstation</a>	RTX A6000, RTX A5000, RTX A4000
HP	<a href="#">ZCentral 4R</a>	RTX A6000, RTX A5000, RTX A4000
HP	<a href="#">Z6 G4 Workstation</a>	RTX A6000, RTX A4000
HP	<a href="#">Z8 G4 Workstation</a>	RTX A6000, RTX A5000, RTX A4000
Leadtek	<a href="#">WinFast WS750</a>	RTX A5500, RTX A4500
Leadtek	<a href="#">WinFast WS940</a>	RTX A6000, RTX A5500, RTX A5000, RTX A4500, RTX A4000
Leadtek	<a href="#">WinFast WS945</a>	RTX A6000, RTX A5500, RTX A5000, RTX A4500, RTX A4000
Leadtek	<a href="#">WinFast WS1040</a>	RTX A6000, RTX A5000, RTX A4000
Leadtek	<a href="#">WinFast WS2040</a>	RTX A6000, RTX A5000, RTX A4000
Lenovo	<a href="#">ThinkPad P16</a>	RTX A5500
Lenovo	<a href="#">ThinkStation P350 Tower</a>	RTX A5000, RTX A4000
Lenovo	<a href="#">ThinkStation P520 Workstation</a>	RTX A6000, RTX A5500, RTX A5000, RTX A4000
Lenovo	<a href="#">ThinkStation P520c Workstation</a>	RTX A5000, RTX A4000
Lenovo	<a href="#">ThinkStation P620 Tower Workstation</a>	RTX A6000, RTX A5500, RTX A5000, RTX A4000
Lenovo	<a href="#">ThinkStation P720 Tower Workstation</a>	RTX A6000, RTX A5000, RTX A4000
Lenovo	<a href="#">ThinkStation P920 Tower Workstation</a>	RTX A6000, RTX A5500, RTX A5000
Microway	<a href="#">Whisperstation</a>	RTX A6000
Supermicro	<a href="#">A+ SuperWorkstation 5014A-TT</a>	RTX A6000, RTX A5500, RTX A4500
Supermicro	<a href="#">A+ SuperWorkstation 5039A-i</a>	RTX A6000
Supermicro	<a href="#">GPU SuperServer SYS-740GP-TNRT</a>	RTX A6000, RTX A5000
Supermicro	<a href="#">UP Workstation SYS-531A-IL</a>	RTX A6000, RTX A5500, RTX A4500, RTX A2000
Supermicro	<a href="#">UP Workstation SYS-540A-TR</a>	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 6. 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 <sup>1</sup>
BOXX	<a href="#">GoBOXX SLM 17</a>	RTX A5500
Dell	<a href="#">Precision 7560 Workstation</a>	RTX A5000, RTX A4000
Dell	<a href="#">Precision 7760 Workstation</a>	RTX A5000, RTX A4000
Lenovo	<a href="#">ThinkPad P1 Gen 4 Mobile Workstation</a>	RTX A5000, RTX A4000
Lenovo	<a href="#">ThinkPad P15 Gen 2 Mobile Workstation</a>	RTX A5000, RTX A4000
Lenovo	<a href="#">ThinkPad P17 Gen 2 Mobile Workstation</a>	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 7. 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 <sup>1</sup>
ATOS	<a href="#">BullSequana SA20G</a>	A16
DataON	<a href="#">AZS-6208G</a>	A16
Dell	<a href="#">PowerEdge R740/R740xd</a>	A16
Dell	<a href="#">PowerEdge R750</a>	A16
Dell	<a href="#">PowerEdge R750xa</a>	A16
Dell	<a href="#">PowerEdge R7515</a>	A16
Dell	<a href="#">PowerEdge R7525</a>	A16
GIGABYTE	<a href="#">E152-ZE0</a>	A16
GIGABYTE	<a href="#">E162-Z20</a>	A16
GIGABYTE	<a href="#">G242-Z12</a>	A16
GIGABYTE	<a href="#">G292-280</a>	A16
GIGABYTE	<a href="#">G292-Z24</a>	A16
GIGABYTE	<a href="#">G292-Z44</a>	A16
GIGABYTE	<a href="#">G292-Z45</a>	A16
GIGABYTE	<a href="#">G492-H80</a>	A16
GIGABYTE	<a href="#">G492-HA0</a>	A16
GIGABYTE	<a href="#">R282-Z93</a>	A16
HPE	<a href="#">ProLiant DL380 Gen10</a>	A16
HPE	<a href="#">ProLiant DL380 Gen10 Plus</a>	A16
HPE	<a href="#">ProLiant DL385 Gen10 Plus</a>	A16
QCT	<a href="#">QuantaGrid D43N-3U</a>	A16
QCT	<a href="#">QuantaGrid D53XQ-2U</a>	A16

Partner	High-Density VDI Servers	Supported NVIDIA GPUs <sup>1</sup>
Supermicro	<a href="#">SYS-740GP-TNRT</a>	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 8. List of NVIDIA-Certified Systems - Edge Systems

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

Partner	System	Supported NVIDIA GPUs <sup>1</sup>	Enterprise Edge <sup>2</sup>	Industrial Edge <sup>3</sup>	NVIDIA AI Enterprise Compatible <sup>4</sup>
Aaeon	<a href="#">BOXER-8332</a>	A2		Yes	Bare Metal
Advantech	<a href="#">ACP-2020G</a>	A2		Yes	Bare Metal
Advantech	<a href="#">EPC-B3522</a>	RTX A4500		Yes	
Advantech	<a href="#">EPC-B3588</a>	RTX A4500		Yes	
Advantech	<a href="#">SKY-640V2</a>	A30	Yes		Bare Metal
Advantech	<a href="#">USM-501</a>	RTX A6000		Yes	Bare Metal
Aetina	<a href="#">AIP-FQ47-B1</a>	RTX A4500		Yes	
Aetina	<a href="#">AIS-D422-A1</a>	A2, T4		Yes	Bare Metal
Axiomtek	<a href="#">IPC-962-525</a>	A2	Yes	Yes	Bare Metal
DataON	<a href="#">AZS-6208G</a>	A2	Yes	Yes	Bare Metal
Dell	<a href="#">PowerEdge R640</a>	A2	Yes		Bare Metal, vSphere <sup>5</sup>
Dell	<a href="#">PowerEdge R6515</a>	A2	Yes		Bare Metal, vSphere <sup>5</sup>
Dell	<a href="#">PowerEdge R650</a>	A2, T4	Yes		Bare Metal, vSphere <sup>5</sup>
Dell	<a href="#">PowerEdge R740/R740xd</a>	A100, A40, A30, A2, T4	Yes		Bare Metal, vSphere <sup>5</sup>
Dell	<a href="#">PowerEdge R750</a>	A100, A40, A30, A2, T4	Yes		Bare Metal, vSphere <sup>5</sup>
Dell	<a href="#">PowerEdge R750xa</a>	A2	Yes		Bare Metal, vSphere <sup>5</sup>
Dell	<a href="#">PowerEdge R7515</a>	A30, A2	Yes		Bare Metal, vSphere <sup>5</sup>

Partner	System	Supported NVIDIA GPUs <sup>1</sup>	Enterprise Edge <sup>2</sup>	Industrial Edge <sup>3</sup>	NVIDIA AI Enterprise Compatible <sup>4</sup>
Dell	<a href="#">PowerEdge R7525</a>	A100, A40, A30, A2, T4	Yes		Bare Metal, vSphere <sup>5</sup>
Dell	<a href="#">PowerEdge XR11</a>	A2, T4	Yes		Bare Metal, vSphere <sup>5</sup>
Dell	<a href="#">PowerEdge XR12</a>	A100, A40, A30, A2, T4	Yes		Bare Metal, vSphere <sup>5</sup>
Fujitsu	<a href="#">PRIMERGY RX2540 M6</a>	A100, A40, A30, A2, T4	Yes		Bare Metal, vSphere <sup>5</sup>
Fujitsu	<a href="#">PRIMERGY TX1330 M5</a>	A2	Yes		Bare Metal
GIGABYTE	<a href="#">E152-ZE0</a>	A2	Yes		Bare Metal
GIGABYTE	<a href="#">E162-220</a>	A2	Yes		Bare Metal
GIGABYTE	<a href="#">G242-P33</a>	A100	Yes		Bare Metal
HPE	<a href="#">ProLiant DL110 Gen10 Plus</a>	A2		Yes	Bare Metal
HPE	<a href="#">ProLiant DL360 Gen10</a>	A2	Yes		Bare Metal
HPE	<a href="#">ProLiant DL360 Gen10 Plus</a>	A2, T4	Yes		Bare Metal
HPE	<a href="#">ProLiant DL380 Gen10</a>	A100, A2, T4	Yes		Bare Metal, vSphere <sup>5</sup>
HPE	<a href="#">ProLiant DL380 Gen10 Plus</a>	A100, A40, A30, T4	Yes		Bare Metal, vSphere <sup>5</sup>
HPE	<a href="#">ProLiant DL385 Gen10 Plus V2</a>	A100, A40, A2, T4	Yes		Bare Metal, vSphere <sup>5</sup>
Inspur	<a href="#">NF5280M6</a>	A2	Yes		Bare Metal
Lanner Electronics	<a href="#">FX-3420</a>	A30	Yes		Bare Metal
Lanner Electronics	<a href="#">LEC-2290E</a>	A2		Yes	Bare Metal
Lenovo	<a href="#">ThinkSystem SE350 Edge Server</a>	A2	Yes	Yes	Bare Metal
Lenovo	<a href="#">ThinkEdge SE450 Edge Server</a>	A30, A2	Yes	Yes	Bare Metal
Lenovo	<a href="#">ThinkSystem SR650 V2 Rack Server</a>	A100, A40, A30, RTX A6000	Yes		Bare Metal
Lenovo	<a href="#">ThinkSystem SR665</a>	A100, A40, A30	Yes		Bare Metal, vSphere <sup>5</sup>

Partner	System	Supported NVIDIA GPUs <sup>1</sup>	Enterprise Edge <sup>2</sup>	Industrial Edge <sup>3</sup>	NVIDIA AI Enterprise Compatible <sup>4</sup>
Mercury	<a href="#">RESX07-2U22R</a>	A100		Yes	Bare Metal
Neousys	<a href="#">Nuvo-7166GC</a>	A2		Yes	Bare Metal
One Stop Systems	<a href="#">Rigel Edge Supercomputer</a>	HGX A100 - 4 GPU		Yes	Bare Metal
QCT	<a href="#">QuantaGrid D53X-1U</a>	A2	Yes		Bare Metal
Siemens AG	<a href="#">IPC647E</a>	A2, RTX A6000		Yes	Bare Metal
Supermicro	<a href="#">CloudDC SuperServer SYS-120C-TN10R</a>	A2	Yes		Bare Metal
Supermicro	<a href="#">IoT SuperServer SYS-220HE-FTNR</a>	A40	Yes		Bare Metal
Supermicro	<a href="#">IoT SuperServer SYS-E403-12P-FN2T</a>	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 systems are tested for remote systems management using the industry standard Redfish API.
3. Industrial Edge systems are tested at elevated ambient temperatures to demonstrate their ability to function effectively outside a typical datacenter environment.
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 9. 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 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, 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.

