



Certified OEM Platforms for EOL Products

Application Note

Document History

DA-09018-001_v08

Version	Date	Authors	Description of Change
01	May 17, 2018	VK, SM	Initial Release
02	July 23, 2019	VK, SM	<ul style="list-style-type: none">• Added Tesla M6 and Tesla P100 products• Updated application note to meet current NVIDIA standards
03	May 18, 2020	VK, SM	Added Tesla M60 and Tesla P4 products
04	May 26, 2020	VK, SM	Updated "Supported Servers" section to include Tesla M60 and Tesla P4
05	June 30, 2020	VK, SM	Updated heading in Table 5
06	July 27, 2020	VK, SM	Updated Dell listing in Table 4
07	January 12, 2021	VK, SM	Added Tesla P6 and Tesla P40 products
08	May 4, 2021	VK, SM	Added Tesla V100 product

Table of Contents

Overview	1
NVIDIA GRID K1	2
NVIDIA GRID K1 Hardware Specifications	2
NVIDIA GRID K2	3
NVIDIA GRID K2 Hardware Specifications	3
Tesla M6	4
Tesla M6 Hardware Specifications.....	4
Tesla P100	5
Tesla P100 PCIe Specifications	5
Tesla P100 SXM2 Specifications.....	5
Tesla M60	6
Tesla M60 PCIe Specifications	6
Tesla P4	7
Tesla P4 PCIe Specifications	7
Tesla P6	8
Tesla P6 Hardware Specifications	8
Tesla P40	9
Tesla P40 Hardware Specifications	9
Tesla V100	10
Tesla V100 Hardware Specifications.....	10
Supported Servers	11

List of Tables

Table 1.	NVIDIA GRID K1 and NVIDIA GRID K2 Supported Servers	11
Table 2.	Tesla M6 Supported Servers.....	15
Table 3.	Tesla P100 Supported Servers	15
Table 4.	Tesla M60 Supported Servers.....	18
Table 5.	Tesla P4 Supported Servers	22
Table 6.	Tesla P6 Supported Servers	24
Table 7.	Tesla P40 Supported Servers	25
Table 8.	Tesla V100 Supported Servers.....	29

Overview

This application note lists the OEM server platforms which are currently supported for the following NVIDIA GRID® and NVIDIA® Tesla® products.

- ▶ NVIDIA GRID K1
- ▶ NVIDIA GRID K2
- ▶ Tesla M6
- ▶ Tesla P100
- ▶ Tesla M60
- ▶ Tesla P4
- ▶ Tesla P6
- ▶ Tesla P40
- ▶ Tesla V100

NVIDIA GRID K1

The NVIDIA GRID K1 is a dual-slot 10.5 inch PCI Express Gen3 graphics board with four NVIDIA Kepler™ graphics processing units (GPUs). The NVIDIA GRID K1 has 16 GB of DDR3 memory (4 GB per GPU), and a 130 W maximum power limit. The NVIDIA GRID K1 graphics board uses a passive heat sink that requires system airflow to properly operate the card within thermal limits. It is designed to accelerate graphics in virtual desktop environments, making it the ideal graphics processor for Microsoft RemoteFX and VMware vSGA.

NVIDIA GRID K1 Hardware Specifications

The following list provides the hardware specifications for NVIDIA GRID K1.

- ▶ Four GK107 GPUs
- ▶ PCI Express 3.0 × 16 system interface
- ▶ Physical dimensions: 10.5 inches × 4.4 inches (dual-slot)
- ▶ Board power: 130 W (maximum)
- ▶ One 6-pin PCI Express power connector

NVIDIA GRID K2

The NVIDIA GRID K2 is a dual-slot 10.5 inch PCI Express Gen3 graphics card with two high-end NVIDIA Kepler GPUs. The NVIDIA GRID K2 has 8 GB of GDDR5 memory (4 GB per GPU), and a 225 W maximum power limit. The NVIDIA GRID K2 graphics board uses a passive heat sink that requires system airflow to properly operate the card within thermal limits. It is designed to accelerate graphics in virtual remote workstation and virtual desktop environments.

NVIDIA GRID K2 Hardware Specifications

The following list provides the hardware specifications for NVIDIA GRID K2.

- ▶ Two GK104 GPUs
- ▶ PCI Express Gen3 × 16 system interface
- ▶ Physical dimensions: 10.5 inches × 4.4 inches (dual-slot)
- ▶ Board power: 225 W (maximum)

Tesla M6

The NVIDIA Tesla M6 is an MXM 3.1 Type B card with a single NVIDIA Maxwell™ GM204 GPU. It has 8 GB GDDR5 on-board memory and a 100 W maximum power limit.

Tesla M6 is specifically designed to fit into constrained space available in blade servers. NVIDIA does not ship it with a cooling solution attached. However, it provides thermal specifications that OEMs can use to design their custom heat sinks.

Tesla M6 Hardware Specifications

The following list provides the hardware specifications for Tesla M6.

- ▶ One GM204 GPU
- ▶ MXM Form Factor
- ▶ Physical dimensions: 3.2 inches × 4.1 inches
- ▶ Board power: 100 W (maximum)

Tesla P100

The NVIDIA Tesla P100 GPU Accelerator for PCIe is a dual-slot 10.5 inch PCI Express Gen3 card with a single NVIDIA Pascal™ GP100 GPU. It uses a passive heat sink for cooling, which requires system air flow to properly operate the card within its thermal limits. The Tesla P100 PCIe supports double precision (FP64), single precision (FP32) and half precision (FP16) compute tasks, unified virtual memory and page migration engine.

Tesla P100 GPU Accelerator is available in three different configurations.

- ▶ Tesla P100 PCIe 12GB
- ▶ Tesla P100 PCIe 16GB
- ▶ Tesla P100 SXM2

Tesla P100 for PCIe is available in two memory configurations

- ▶ Tesla P100 PCIe with 16GB HBM2
- ▶ Tesla P100 PCIe with 12GB HBM2

Tesla P100 PCIe Specifications

The following list provides the PCIe specifications for Tesla P100 SXM2 GPU Accelerator.

- ▶ One GP100 GPU
- ▶ PCI Express Gen3 × 16 system interface
- ▶ Physical dimensions: 10.5 inches × 4.4 inches (dual-slot)
- ▶ Board power: 250 W (maximum)

Tesla P100 SXM2 Specifications

The following list provides the hardware specifications for Tesla P100 SXM2 GPU Accelerator.

- ▶ One GP100 GPU
- ▶ SXM2 Form factor
- ▶ Physical dimensions: 5.5inches x 3.1inches x .5inches
- ▶ Board power: 300 W (maximum)

Tesla M60

The NVIDIA Tesla M60 is a dual-slot 10.5 inch PCI Express Gen3 graphics card with two high-end NVIDIA Maxwell™ GPUs. The Tesla M60 has 16 GB GDDR5 memory (8 GB per GPU) and a 300 W maximum power limit. The board is offered in a 300 W passively cooled variant that requires system airflow to properly operate the card within its thermal limits or in a 240 W actively cooled version. It is designed for single precision GPU compute tasks as well as to accelerate graphics in virtual remote workstation and virtual desktop environments.

A main feature of the Tesla M60 board is the support of the NVIDIA GRID software which includes NVIDIA GRID vGPU™. This technology enables the virtualization of physical GPUs into full-featured virtual GPUs providing maximum performance and scalability. For performance optimization this board utilizes NVIDIA GPU Boost™. By adjusting the GPU clock dynamically, maximum performance is achieved within the power cap limit (300 W or 240 W).

Tesla M60 PCIe Specifications

The following list provides the PCIe specifications for Tesla M60 PCIe GPU Accelerator.

- ▶ Two GM204 GPU
- ▶ PCI Express Gen3 × 16 system interface
- ▶ Physical dimensions: 10.5 inches × 4.4 inches (dual-slot)
- ▶ Board power: 300 W (maximum)

The Tesla M60 board is available in four variants. Each version supports a single unidirectional airflow.

- ▶ PG402 SKU 40 supports passive cooling with left-to-right airflow
- ▶ PG402 SKU 60 supports passive cooling with right-to-left airflow
- ▶ PG402 SKU 80 supports active cooling with straight extender
- ▶ PG402 SKU 80 supports active cooling with long offset extender

Tesla P4

The NVIDIA Tesla P4 is a single-slot, low profile, 6.6-inch PCI Express Gen3 GPU Accelerator with an NVIDIA Pascal GPU. The Tesla P4 has 8 GB GDDR5 memory and a 75 W maximum power limit. The Tesla P4 is offered as a 75 W passively cooled board that requires system air flow to properly operate the card within thermal limits.

The NVIDIA Tesla P4 features optimized INT8 instructions aimed at deep learning inference computations. As a result, the NVIDIA Tesla P4 delivers 21 TOPs (Tera Operations per second) of inference performance, enabling smart responsive artificial intelligence (AI)-based services. Tesla P4 is also supported with Virtual GPU solutions

Tesla P4 PCIe Specifications

The following list provides the PCIe specifications for Tesla P4

- ▶ One GP104
- ▶ PCI Express Gen3 × 16 system interface
- ▶ Physical dimensions: 6.6 inches × 2.7 inches (single-slot)
- ▶ Board power: 75 W (maximum)

Tesla P6

The NVIDIA Tesla P6 is an MXM 3.1 Type B card with a single NVIDIA Pascal GP104 GPU. It has 16 GB GDDR5 video memory and a 90 W maximum power limit. This board is intended for accelerated graphics in virtual remote workstation and virtual desktop environments, as well as for single precision GPU compute tasks.

Tesla P6 is specifically designed to fit into constrained space available in blade servers.

NVIDIA does not ship it with a cooling solution attached. However, it provides thermal specifications that OEMs can use to design their custom heat sinks.

Tesla P6 Hardware Specifications

The following list provides the hardware specifications for Tesla P6.

- ▶ One GP104 GPU
- ▶ MXM Form Factor
- ▶ Physical dimensions: 3.2 inches × 4.1 inches
- ▶ Board power: 90 W (maximum)

Tesla P40

The NVIDIA Tesla P40 GPU Accelerator is a dual-slot 10.5-inch PCI Express Gen3 graphics card based on a high-end NVIDIA Pascal GPU. The Tesla P40 GPU Accelerator has 24 GB GDDR5 memory and a 250 W maximum power limit.

Tesla P40 Hardware Specifications

The following list provides the hardware specifications for Tesla P40.

- ▶ One GP102 GPU
- ▶ PCI Express Gen3 × 16 system interface
- ▶ Physical dimensions: 10.5 inches × 4.4 inches (dual-slot)
- ▶ Board power: 250 W (maximum)

Tesla V100

The NVIDIA Tesla V100 GPU Accelerator for PCIe is a dual-slot 10.5 inch PCI Express Gen3 card with a single NVIDIA Volta GV100 GPU. It uses a passive heat sink for cooling, which requires system air flow to properly operate the card within its thermal limits. The Tesla V100 GPU Accelerator has 16 GB HBM2 memory and a 250 W maximum power limit

Tesla V100 Hardware Specifications

- ▶ One GV100 GPU
- ▶ PCI Express Gen3 × 16 system interface
- ▶ Physical dimensions: 10.5 inches × 4.4 inches (dual-slot)
- ▶ Board power: 250 W (maximum)

Supported Servers

The following tables contain the supported servers and models.

- ▶ NVIDIA GRID K1 and NVIDIA GRID K2 (Table 1)
- ▶ Tesla M6 (Table 2)
- ▶ Tesla P100 (Table 3)
- ▶ Tesla M60 (Table 4)
- ▶ Tesla P4 (Table 5)
- ▶ Tesla P6 (Table 6)
- ▶ Tesla P40 (Table 7)
- ▶ Tesla V100 (Table 8)

Table 1. NVIDIA GRID K1 and NVIDIA GRID K2 Supported Servers

Manufacturer	Model	Rack Units	Node per Chassis	NVIDIA GRID K1	NVIDIA GRID K2
ASRock Rack	2U2N-F/4GC612	2	2	-	2
ASUS	ESC4000 G2	2	1	2	4
ASUS	ESC4000 G3	2	1	2	4
ASUS	ESC4000 G3S	2	1	2	4
ASUS	ESC4000/FDR G2	2	1	2	4
ASUS	ESC8000 G3	3	1	2	4
ASUS	RS920-E7/RS8	2	1	2	2
ASUS	RS926-E7/RS8	2	1	2	2
Bull	Bullx R421 E3	1	2	2	3
Cisco	UCS C240 M3	2	1	2	2
Cisco	UCS C240 M4	2	1	2	2
Cisco	UCS C460 M4	4	1	2	2
Cubix	RPS NVGrid K2	8	1	-	4
Cubix	SPS Grid K1 JagFast	4	1	2	-
Dell	PowerEdge C4130	1	1	3	4

Manufacturer	Model	Rack Units	Node per Chassis	NVIDIA GRID K1	NVIDIA GRID K2
Dell	PowerEdge C8220X	4	4	-	2
Dell	PowerEdge R720	2	1	2	2
Dell	PowerEdge R730	2	1	2	2
Dell	XC730-16G	2	1	2	2
Dell	PowerEdge T620	2	1	-	4
Dell	PowerEdge T630	5	1	-	4
Dell	PowerEdge VRTX	2	1	-	1
Dell	Precision Appliance for Wyse	2	1	2	2
Dell	Precision R7610	2	2	-	2
Dell	Precision R7910	2	1	-	2
Exxact	Quantum TXR113-1000R	1	1	2	4
Exxact	Quantum TXR130-1000R	1	1	2	3
Exxact	Quantum TXR230-0512R	2	1	2	4
Exxact	Quantum TXR211-1000R	2	1	2	4
Exxact	Quantum TXR110-2000R	1	1	2	4
Exxact	Quantum TXR231-1000R	2	1	2	4
FUJITSU	CELSIUS C620	1	1	-	1
FUJITSU	CELSIUS C740	1	1	-	1
FUJITSU	CELSIUS M740	5	1	2	2
FUJITSU	CELSIUS R940	5	1	2	3
FUJITSU	PRIMERGY CX2570 M1	2	2	1	1
FUJITSU	PRIMERGY RX2540 M1	2	1	2	2
FUJITSU	PRIMERGY RX350 S8	4	1	2	2
FUJITSU	PRIMERGY TX300 S8	5	1	2	2
Gigabyte	G250-S88	2	1	2	4
Gigabyte	R280-G20	2	1	2	3
Hitachi	Compute Blade 520H	6	8	-	1
Hitachi	HA8000/RS220 AN2,BN2	2	1	1	1
HPE	DL380z Gen9 Virtual Workstation	2	1	2	2
HPE	Apollo XL250a Gen9	5	5	2	-
HPE	ProLiant DL380p Gen8	2	1	2	2
HPE	ProLiant DL380 Gen9	2	1	2	2
HPE	ProLiant SL250s Gen8	4	1	-	3
HPE	ProLiant SL270s Gen8 SE	4	1	-	4

Manufacturer	Model	Rack Units	Node per Chassis	NVIDIA GRID K1	NVIDIA GRID K2
HPE	ProLiant WS460c Gen8	10	16	1**	1**
HPE	ProLiant WS460c Gen9	10	16	1**	1**
HPE	Hyper Converged 380	2	1	2	2
Huawei	Tecal CH221 V2	12	8	1	1
Huawei	Tecal RH2288H V2	2	1	1	1
Huawei	FusionServer RH2288H V3	2	1	2	2
Huawei	FusionServer RH5885H V3	4	1	2	1
Huawei	FusionServer XH622 V3	4	4	2	2
Inspur	NF5288	2	1	2	4
Inspur	NF5588M3	4	1	2	4
Leadtek	WinFast GS2000	2	1	2	4
Leadtek	WinFast GS2020	2	1	2	4
Lenovo	Flex System x240 M5	10	7	1	1
Lenovo	NeXtScale nx360 M4	6	6	2	2
Lenovo	NeXtScale nx360 M5	6	4	2	2
Lenovo	System x iDataPlex dx360 M4	2	1	2	2
Lenovo	System x3650 M4	2	1	-	2
Lenovo	System x3650 M5	2	1	2	2
Lenovo	System x3850 X6	4	1	2	2
Lenovo	System x3950 X6	8	1	2	2
NEC	Express 5800/R120e-2M	2	1	1	1
NEC	Express 5800/R120f-2M	2	1	1	1
NEC	Express 5800/R120g-2M	2	1	1	1
Nutanix	NX-3155G-G4	2	1	-	3
Nutanix	NX-3155G-G5	2	1	-	3
Nutanix	NX-3175-G4	1	1	1	1
Nutanix	NX-3175-G5	1	1	1	1
Nutanix	NX-7110	2	1	2	3
One Stop Systems	1U Expansion Chassis	1	1	-	4
One Stop Systems	2U Expansion Chassis	1	1	-	8
Pivot3	vSTAC R2S P Cubed	2	1	1	1
QCT	QuantaGrid D51BV-2U	2	1	2	2
QCT	STRATOS S210-X2A2J	2	1	2	4
SGI	Rackable C1104G-RP5	1	1	2	3
SGI	Rackable C2108-GP5	2	1	2	4

Manufacturer	Model	Rack Units	Node per Chassis	NVIDIA GRID K1	NVIDIA GRID K2
Sugon	I620-G15	2	1	2	2
Sugon	W760-G10	2	1	2	2
Sugon	W580I-G10	4	1	2	4
Supermicro	SYS-1017R-WR	1	1	-	1
Supermicro	SYS-1027GR	1	1	2	3
Supermicro	SYS-1028GQ	1	1	2	4
Supermicro	SYS-1028GR	1	1	2	3
Supermicro	SYS-1028U / 6018U	1	1	1	1
Supermicro	SYS-2027GR	1	-	2	4
Supermicro	SYS-2028GR	2	1	2	4
Supermicro	SYS-2028TP-DC1FR	2	2	1	1
Supermicro	SYS-2028U / 6028U	2	1	-	3
Supermicro	SYS-7047GR-TRF	4	1	2	4
Supermicro	SYS-7048GR-TR	4	1	2	4
Supermicro	SYS-F627G	4	4	2	3
Tyan	FT77AB7059	4	1	-	4
Tyan	GA80-B7061	2	1	-	2
Tyan	TA77-B7061	2	1	-	4
VDI-Appliance	IO-100 G3	1	1	2	4
VDI-Appliance	IO-150 G3	1	1	2	4
VDI-Appliance	IO-250 G3	1	1	2	4
VDI-Appliance	IO-275 G3	1	1	2	4
VDI-Appliance	IO-285 G3	1	1	2	4

Note: **With expansion chassis

Table 2. Tesla M6 Supported Servers

Manufacturer	Model	Rack Units	Node per Chassis	Tesla M6
Cisco	UCS B200 M4	6	8	1
HPE	ProLiant WS460c Gen9	10	16	1 or 4**
HPE	ProLiant WS460c Gen8	10	16	1 or 4**
HPE	Synergy 480 Gen10	10	12	1 or 7**
HPE	Synergy 480 Gen9	10	12	1 or 7**
Amulet Hotkey	CoreStation VM630	10	16	1
Amulet Hotkey	Corestation VM640	10	16	1

Note: **With expansion chassis

Table 3. Tesla P100 Supported Servers

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P100
Cisco	UCS C240 M5	2	1	2
Cisco	UCS C240 M4	2	1	2
Cisco	UCS C480 M5	4	1	6
Dell	PowerEdge C4140	1	1	4
Dell	PowerEdge C4130	1	1	4
Dell	PowerEdge R740	2	1	3
Dell	PowerEdge R740xd	2	1	3
Dell	PowerEdge R730	2	1	2
Dell	XC740xd	2	1	3
Dell	R940xa	4	2	4
Dell	R840	2	2	2
Dell	PowerEdge T640	5	1	4
Dell	PowerEdge T630	5	1	4
HPE	Apollo pc40 Server	1	1	4
HPE	Apollo sx40 Server	1	1	4***
HPE	Apollo XL190r Gen10	2	2	2
HPE	Apollo XL190r Gen9	2	2	2
HPE	ProLiant DL380 Gen10	2	1	3
HPE	ProLiant DL380 Gen9	2	1	2
HPE	Apollo XL270d Gen10	4	2	8***

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P100
HPE	Apollo XL270d Gen9	4	2	8
HPE	ProLiant DL580 Gen10	4	1	4
HPE	Superdome Flex	6	2	4
Lenovo	ThinkSystem SR650	2	1	2
Lenovo	ThinkSystem SD530/D2	2	2	4
Lenovo	ThinkSystem SR670	2	1	4
Lenovo	System x3650 M5	2	1	2
Lenovo	ThinkAgile HX3520-G	2	1	2
Lenovo	ThinkAgile HX3521-G	2	1	2
Lenovo	NeXtScale nx360 M5	6	4	2
Supermicro	SYS-1018GR / 5018GR	1	1	2
Supermicro	SYS-1028GQ	1	1	4***
Supermicro	SYS-1029GQ	1	1	4
Supermicro	SYS-2028GR	2	1	6
Supermicro	SYS-2029U / 6029U	2	1	2
Supermicro	SYS-4028GR	4	1	8***
Supermicro	SYS-7048GR-TR	4	1	4
Supermicro	SYS-7049GP-TRT	4	1	4
Supermicro	SYS-F628G	4	4	3
ASRock Rack	3U8G-C612	3	1	8
ASUS	ESC4000 G3	2	1	4
ASUS	ESC4000 G3S	2	1	4
ASUS	ESC8000 G3	3	1	8
ASUS	ESC4000 G4	2	1	4
FUJITSU	PRIMERGY CX2570 M4	2	2	2
FUJITSU	PRIMERGY RX2540 M4	2	1	2
FUJITSU	CX400M1	2	1	4
FUJITSU	PRIMERGY CX400 M4	2	1	4
Gigabyte	G190-H44	1	1	4
Gigabyte	T180-G23	1	1	4
Gigabyte	T180-G24	1	1	4
Gigabyte	G180-G00	1	1	2
Gigabyte	G250-G50	2	2	8
Gigabyte	G250-G51	2	1	8
Gigabyte	G250-G52	2	1	8

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P100
Gigabyte	G250-S88	2	1	8
Gigabyte	R280-G20	2	1	3
Gigabyte	G481-HA0	4	1	10
Huawei	G560 V5	4	1	8
Huawei	FusionServer RH2288H V3	2	1	2
Huawei	FusionServer RH2288H V5	2	1	2
Huawei	FusionServer XH622 V3	4	4	2
Inspur	NF5280M5	2	1	4
Inspur	NF5280M4	2	1	2
Inspur	NF5288M4	4	1	4
Inspur	NF5568M4	4	1	4
Inventec	K800G3	2	1	2
Inventec	TB800G4	2	1	1
Penguin Computing	Relion 1903GT	1	1	3
Penguin Computing	Relion 2908GT	2	1	8
PNY	PNYSER14 Series	1	1	4
PNY	PNYSER28 Series	2	1	8
PNY	PNYSER48 Series	4	1	8
PNY	PNYSRA48 Series	4	1	8
QCT	QuantaGrid D51BV-2U	2	1	2
QCT	QuantaGrid D52BV-2U	2	1	4
QCT	QuantaGrid Q72D-2U	2	1	2
QCT	QuantaGrid D52G-4U	4	1	8
Sugon	W740-G20	2	1	4
Sugon	W580-G20	4	1	4
Sugon	W780	4	1	8
Tyan	GA88-B5631	1	1	4
Tyan	TN76-B7102	2	1	2
Tyan	TA80-B7071	2	1	4
Tyan	FA77-B7119	4	1	10
Tyan	FT77D-B7109	4	1	8
Tyan	FT76-B7922	4	1	4
Tyan	FT77C-B7079	4	1	8
VDI-Appliance	IO-110 G3	1	1	6
VDI-Appliance	IO-150 G3	1	1	6

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P100
VDI-Appliance	IO-250 G3	1	1	6
VDI-Appliance	IO-275 G3	1	1	6
VDI-Appliance	IO-285 G3	1	1	6

Notes:
 **With expansion chassis
 ***SXM2

Table 4. Tesla M60 Supported Servers

Manufacturer	Model	Rack Units	Node per Chassis	Tesla M60
Cisco	UCS C240 M5	2	1	2
Cisco	UCS C240 M4	2	1	2
Cisco	HyperFlex HX240c M5	2	1	2
Cisco	HyperFlex HX240c M4	2	1	1
Cisco	UCS C480 M5	4	1	6
Cisco	UCS C460 M4	4	1	2
Dell	PowerEdge C4130	1	1	4
Dell	PowerEdge R740	2	1	3
Dell	PowerEdge R740xd	2	1	3
Dell	PowerEdge R730	2	1	2
Dell	XC740xd	2	1	3
Dell	XC730-16G	2	1	2
Dell	Dell EMC VxRail V470	2	1	2
Dell	Dell EMC VxRail V470F	2	1	2
Dell	Dell EMC VxRail V570	2	1	3
Dell	Dell EMC VxRail V570F	2	1	3
Dell	Dell EMC VxRail V570	2	1	3
Dell	Precision Appliance for Wyse	2	1	2
Dell	PowerEdge T640	5	1	4
Dell	PowerEdge T630	5	1	4
HPE	Apollo XL190r Gen9	2	2	2
HPE	ProLiant DL380 Gen10	2	1	3
HPE	ProLiant DL380 Gen9	2	1	2
HPE	Hyper Converged 380	2	1	2

Manufacturer	Model	Rack Units	Node per Chassis	Tesla M60
HPE	ProLiant DL580 Gen9	4	1	4
HPE	Apollo XL250a Gen9	5	5	2
HPE	ProLiant WS460c Gen9	10	16	1**
HPE	Synergy 480 Gen10	10	12	2**
HPE	Synergy 480 Gen9	10	12	1
Lenovo	ThinkSystem SR650	2	1	2
Lenovo	ThinkSystem SD530/D2	2	2	2
Lenovo	ThinkSystem SR860	4	1	2
Lenovo	System x3650 M5	2	1	2
Lenovo	Converged HX3510-G	2	1	2
Lenovo	ThinkAgile HX3520-G	2	1	2
Lenovo	ThinkAgile HX3521-G	2	1	2
Lenovo	System x3850 X6	4	1	2
Lenovo	ThinkStation P710	4	1	1
Lenovo	System x3500 M5	5	1	2
Lenovo	ThinkStation P910	5	1	1
Lenovo	NeXtScale nx360 M5	6	4	4
Lenovo	System x3950 X6	8	1	4
Supermicro	SYS-1018R-WR	1	1	1
Supermicro	SYS-1018GR / 5018GR	1	1	2
Supermicro	SYS-1028GQ	1	1	4
Supermicro	SYS-1028GR	1	1	3
Supermicro	SYS-1028U / 6018U	1	1	1
Supermicro	SYS-1029U / 6019U	1	1	1
Supermicro	SYS-2028GR	2	1	4
Supermicro	SYS-2028TP-DC1FR	2	2	1
Supermicro	SYS-2028U / 6028U	2	1	2
Supermicro	SYS-F628R3	4	4	3
Acer	Acer Altos R480 F4	2	1	8*
Advantech	AGS-913I	1	1	2
Advantech	AGS-923	2	1	2
Advantech	HPC-7400-S813	4	1	2
ASRock Rack	2U2N-F/4GC612	2	2	2
ASRock Rack	3U8G-C612	3	1	4
ASUS	ESC4000 G3	2	1	4

Manufacturer	Model	Rack Units	Node per Chassis	Tesla M60
ASUS	ESC4000 G3S	2	1	4
ASUS	ESC8000 G3	3	1	4
Cubix	SPS Tesla M60 JagFastS	4	1	2
Cubix	RPS Tesla M60 JagFastR	8	1	8*
Exxact	Quantum TXR113-1000R	1	1	4
Exxact	Quantum TXR130-1000R	1	1	3
Exxact	Quantum TXR110-2000R	1	1	4
Exxact	Quantum TXR230-0512R	2	1	4
Exxact	Quantum TXR211-1000R	2	1	4
Exxact	Quantum TXR231-1000R	2	1	4
FUJITSU	CELSIUS C740	1	1	1
FUJITSU	PRIMERGY CX2570 M2	2	2	2
FUJITSU	PRIMERGY CX2570 M4	2	2	2
FUJITSU	PRIMERGY RX2540 M1	2	1	2
FUJITSU	PRIMERGY RX2540 M2	2	1	2
FUJITSU	PRIMERGY RX2540 M4	2	1	2
FUJITSU	PRIMERGY TX2560 M2	4	1	2
Gigabyte	G190-H44	1	1	4
Gigabyte	T180-G23	1	1	4
Gigabyte	T180-G24	1	1	4
Gigabyte	G180-G00	1	1	2
Gigabyte	G250-G50	2	2	8*
Gigabyte	G250-G51	2	1	8*
Gigabyte	G250-G52	2	1	8*
Gigabyte	G250-S88	2	1	8*
Gigabyte	R280-G20	2	1	3
Gigabyte	G291-280	2	1	8*
Hitachi	HA8000/RS220 AN2,BN2	2	1	1
Hitachi Vantara	Hitachi Advanced Server DS225	2	1	4
H3C	UIS R390X G2	2	1	3
Huawei	FusionServer RH2288H V3	2	1	2
Huawei	FusionServer RH2288H V5	2	1	2
Huawei	FusionServer CH220 V3	12	16	2
Inspur	Yitian NF 5288M4	2	1	4

Manufacturer	Model	Rack Units	Node per Chassis	Tesla M60
Inspur	NF5280M5	2	1	4
Inventec	K888G3	2	1	4
Leadtek	WinFast GS1020	1	1	4
Leadtek	WinFast GS1020S	1	1	2
Leadtek	WinFast GS1020V	1	1	3
Magma	ExpressBox 3600-AB	4	1	8
NEC	Express 5800/R120g-2M	2	1	1
NEC	Express 5800/T120g	5	1	1
Nutanix	NX-3175-G4	1	1	1
Nutanix	NX-3175-G5	1	1	1
Nutanix	NX-3155G-G5	2	1	2
Penguin Computing	Relion 1904GT	1	1	4
PNY	PNYSER14 Series	1	1	4
PNY	PNYSER28 Series	2	1	4
QCT	QuantaGrid D51BV-2U	2	1	2
QCT	QuantaGrid D52BV-2U	2	1	4
Simplivity	OmniStack Integrated Solution with Huawei FusionServer	2	1	2
Simplivity	OmniStack Integrated Solution with Lenovo System x	2	1	2
Sugon	I620-G20	2	1	2
Sugon	W740-G20	2	1	4
Sugon	W580-G20	4	1	4
Themis	RESXR5-2U	2	1	1
Tyan	TA80-B7071	2	1	4
VDI-Appliance	IO-110 G3	1	1	4
VDI-Appliance	IO-150 G3	1	1	4
VDI-Appliance	IO-250 G3	1	1	4
VDI-Appliance	IO-275 G3	1	1	4
VDI-Appliance	IO-285 G3	1	1	4

Notes:

*The maximum number of supported Tesla M60 cards per system when using NVIDIA GRID vGPU is four

**With expansion chassis

Table 5. Tesla P4 Supported Servers

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P4
Cisco	UCS C220 M5	1	1	2
Cisco	UCS C240 M5	2	1	6
Cisco	HyperFlex HX240c M5	2	1	6
Dell	PowerEdge R740	2	1	6
Dell	PowerEdge R740xd	2	1	6
Dell	PowerEdge R7425	2	1	6
Dell	PowerEdge R730	2	1	4
Dell	XC740xd	2	1	6
HPE	ProLiant DL360 Gen10	1	1	2
HPE	ProLiant DL360 Gen9	1	1	2
HPE	ProLiant DL380 Gen10	2	1	5
HPE	ProLiant DL380 Gen9	2	1	3
HPE	Edgeline EL1000	1	1	2
HPE	Edgeline EL4000	1	1	4
HPE	Apollo XL270d Gen9	4	2	8
Lenovo	ThinkSystem SR630	1	1	2
Lenovo	ThinkSystem HR630X	1	1	1
Lenovo	ThinkSystem SR650	2	1	2
Lenovo	ThinkAgile HX3520-G	2	1	2
Lenovo	ThinkAgile HX3521-G	2	1	2
Lenovo	NeXtScale nx360 M5	6	4	2
Supermicro	SYS-1018GR / 5018GR	1	1	2
Supermicro	SYS-1028GQ	1	1	4
Supermicro	SYS-1029P-W	1	1	1
Supermicro	SYS-1028U / 6018U	1	1	1
Supermicro	SYS-1029U / 6019U	1	1	2
Supermicro	SYS-2028GR	2	1	6
Supermicro	SYS-2029U / 6029U	2	1	4
Supermicro	SYS-2029GP-TR	2	1	6
Supermicro	SYS-4028GR	4	1	8
Supermicro	SYS-7048GR-TR	4	1	4
Gigabyte	G190-H44	1	1	4
Gigabyte	G191-H44	1	1	4

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P4
Gigabyte	T180-G23	1	1	4
Gigabyte	T180-G24	1	1	4
Gigabyte	G180-G00	1	1	2
Gigabyte	G250-G50	2	2	8
Gigabyte	G250-G51	2	1	8
Gigabyte	G250-G52	2	1	8
Gigabyte	G250-S88	2	1	8
Gigabyte	G291-280	2	1	8
Gigabyte	G291-2G0	2	1	16
Gigabyte	R281-G30	2	1	3
Gigabyte	R281-3C2	2	1	3
H3C	R4900 G3	2	1	4
H3C	R4700 G3	1	1	2
Huawei	G5500/G530 V5/GP308	4	1	6
Huawei	G5500/G530 V5/GP316	4	1	16
Huawei	G5500/G560 V5/GP608	4	1	8
Huawei	G560 V5	4	1	8
Huawei	FusionServer RH2288H V3	2	1	2
Huawei	FusionServer RH2288H V5	2	1	7
Huawei	FusionServer CH220 V3	12	16	6
Inspur	NF5280M5	2	1	4
Inspur	NF5280M4	2	1	2
Inspur	NF5288M5	2	1	8
Inspur	NF5288M4	2	1	4
Inspur	NF5468M5	4	1	16
Inspur	NF5568M4	4	1	4
Inventec	K800G3	2	1	1
NEC	Express 5800/R120g-2M	2	1	2
Penguin Computing	Relion 1904GT	1	1	4
Penguin Computing	Relion 2908GT	2	1	8
Penguin Computing	Relion XE2112GT	2	1	3
Penguin Computing	Relion XE1114GT	1	1	4
Penguin Computing	Relion XE2118GT	2	1	8
QCT	QuantaGrid D52BV-2U	2	1	4
QCT	QuantaGrid Q72D-2U	2	1	2

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P4
QCT	QuantaGrid D52G-4U	4	1	16
Sugon	W760-G30	2	1	6
Sugon	W580-G20	4	1	4
Sugon	X785-G30	4	1	20
Tyan	FT77C-B7079	4	1	8
VDI-Appliance	IO-110 G3	1	1	6
VDI-Appliance	IO-150 G3	1	1	6
VDI-Appliance	IO-250 G3	1	1	6
VDI-Appliance	IO-275 G3	1	1	6
VDI-Appliance	IO-285 G3	1	1	6

Table 6. Tesla P6 Supported Servers

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P6
Cisco	UCS B200 M5	6	8	2
Cisco	UCS B480 M5	6	4	4
HPE	ProLiant BL460c Gen10	10	16	1 or 4*
HPE	Synergy 480 Gen10	10	12	1 or 7*
Amulet Hotkey	CoreStation VM630	10	16	1
Amulet Hotkey	CoreStatuion VM640	10	16	1
Amulet Hotkey	Corestation VFC640	2	4	4

Notes: Mixed GPU configurations within a server are not supported.
*With expansion chassis.

Table 7. Tesla P40 Supported Servers

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P40
Cisco	UCS C240 M5	2	1	2
Cisco	HyperFlex HX240c M5	2	1	2
Cisco	UCS C480 M5	4	1	6
Dell	PowerEdge C4140	1	1	4
Dell	PowerEdge C4130	1	1	4
Dell	PowerEdge R740	2	1	3
Dell	PowerEdge R740xd	2	1	3
Dell	PowerEdge R7425	2	1	3
Dell	PowerEdge R730	2	1	2
Dell	XC740xd	2	1	3
Dell	R940xa	4	2	4
Dell	R840	2	2	2
Dell	Dell EMC VxRail V570	2	1	3
Dell	Dell EMC VxRail V570F	2	1	3
Dell	PowerEdge T630	5	1	4
HPE	Apollo pc40 Server	1	1	4
HPE	Apollo XL190r Gen10	2	2	2
HPE	Apollo XL190r Gen9	2	2	2
HPE	ProLiant DL380 Gen10	2	1	3
HPE	ProLiant DL385 Gen10	2	1	3
HPE	ProLiant DL380 Gen9	2	1	2
HPE	SimpliVity 380	2	1	1
HPE	SimpliVity 380G	2	1	2
HPE	SimpliVity 2600	2	2	1
HPE	Apollo XL270d Gen10	4	2	8
HPE	Apollo XL270d Gen9	4	2	8
HPE	ProLiant ML350 Gen10	4	1	4
HPE	ProLiant DL580 Gen10	4	1	4
HPE	ProLiant BL460c Gen10	10	16	1
HPE	Synergy 480 Gen10	10	12	2**
HPE	Synergy 480 Gen9	10	12	2**
Lenovo	ThinkSystem SR650	2	1	2
Lenovo	ThinkSystem SD530/D2	2	2	4

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P40
Lenovo	ThinkSystem SR670	2	1	4
Lenovo	System x3650 M5	2	1	2
Lenovo	ThinkAgile HX3520-G	2	1	2
Lenovo	ThinkAgile HX3521-G	2	1	2
Lenovo	NeXtScale nx360	8	1	16
Lenovo	NeXtScale nx360 M5	6	4	2
Supermicro	SYS-1018GR / 5018GR	1	1	2
Supermicro	SYS-1028GQ	1	1	4
Supermicro	SYS-1029GQ	1	1	4
Supermicro	SYS-1028GR	1	1	3
Supermicro	SYS-1028U / 6018U	1	1	1
Supermicro	SYS-1029U / 6019U	1	1	1
Supermicro	SYS-2028GR	2	1	6
Supermicro	SYS-2028U / 6028U	2	1	2
Supermicro	SYS-2029U / 6029U	2	1	2
Supermicro	SYS-2029GP-TR	2	1	6
Supermicro	SYS-4028GR	4	1	8
Supermicro	SYS-7048GR-TR	4	1	4
Supermicro	SYS-7049GP-TRT	4	1	4
Supermicro	SYS-F628G	4	4	4
Advantech	AGS-913I	1	1	2
Advantech	AGS-923	2	1	4
Advantech	HPC-7483-S923	4	1	4
ASRock Rack	2U2N-F/4GC612	2	2	4
ASRock Rack	2U2G	2	1	2
ASUS	ESC4000 G3	2	1	4
ASUS	ESC4000 G3S	2	1	4
ASUS	ESC8000 G3	3	1	8
ASUS	ESC4000 G4	2	1	4
ASUS	RS720-E9-RS8-G	2	1	2
FUJITSU	PRIMERGY CX2570 M4	2	2	2
FUJITSU	PRIMERGY RX2540 M4	2	1	2
FUJITSU	PRIMERGY CX400 M4	2	1	4
Gigabyte	G190-H44	1	1	4
Gigabyte	G191-H44	1	1	4

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P40
Gigabyte	T180-G23	1	1	4
Gigabyte	T180-G24	1	1	4
Gigabyte	T181-Z70	1	1	4
Gigabyte	G180-G00	1	1	2
Gigabyte	G250-G50	2	2	8
Gigabyte	G250-G51	2	1	8
Gigabyte	G250-G52	2	1	8
Gigabyte	G250-S88	2	1	8
Gigabyte	R280-G20	2	1	3
Gigabyte	G291-280	2	1	8
Gigabyte	G221-Z30	2	1	2
Gigabyte	G291-Z20	2	1	8
Gigabyte	R281-G30	2	1	3
Gigabyte	R281-3C2	2	1	3
Gigabyte	G481-HA0	4	1	10
Hitachi	HA8000V	2	1	3
Hitachi Vantara	Hitachi Advanced Server DS225	2	1	4
H3C	R4900 G3	2	1	3
Huawei	G5500/G530 V5/GP308	4	1	4
Huawei	G5500/G560 V5/GP608	4	1	8
Huawei	FusionServer RH2288H V5	2	1	2
Inspur	Yitian NF 5288M4	2	1	4
Inspur	NF5280M5	2	1	4
Inspur	NF5280M4	2	1	2
Inspur	NF5288M5	2	1	8
Inspur	NF5288M4	2	1	4
Inspur	NF5468M5	4	1	8
Inspur	NF5568M4	4	1	4
Nutanix	NX-3175-G5	1	1	1
Nutanix	NX-3170-G6	1	1	1
Nutanix	NX-3155G-G5	2	1	2
Nutanix	NX-3155G-G6	2	1	2
Penguin Computing	Relion 1904GT	1	1	4
Penguin Computing	Relion 2904GT	2	1	4

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P40
Penguin Computing	Relion 2908GT	2	1	8
Penguin Computing	Relion XE1114GT	1	1	4
Penguin Computing	Relion XE2114GT	2	1	4
Penguin Computing	Relion XE4118GT	4	1	10*
PNY	PNYSER14 Series	1	1	4
PNY	PNYSER28 Series	2	1	8
PNY	PNYSER48 Series	4	1	8
PNY	PNYSRA48 Series	4	1	8
QCT	QuantaGrid D51BV-2U	2	1	2
QCT	QuantaGrid D52BV-2U	2	1	4
QCT	QuantaGrid Q72D-2U	2	1	2
QCT	QuantaGrid D52G-4U	4	1	8
Sugon	W760-G20	2	1	2
Sugon	W740-G20	2	1	4
Sugon	W580-G20	4	1	4
Sugon	W780	4	1	4
Tyan		2	1	4
Tyan	FA77-B7119	4	1	10
Tyan	FT77D-B7109	4	1	8
Tyan	FT76-B7922	4	1	4
Tyan	FT77C-B7079	4	1	8
VDI-Appliance	IO-110 G3	1	1	6
VDI-Appliance	IO-150 G3	1	1	6
VDI-Appliance	IO-250 G3	1	1	6
VDI-Appliance	IO-275 G3	1	1	6
VDI-Appliance	IO-285 G3	1	1	6

Notes:

*Mixed GPU configurations within a server are not supported.

**With expansion chassis.

Table 8. Tesla V100 Supported Servers

Manufacturer	Model	Rack Units	Node per Chassis	Tesla V100
Cisco	UCS C240 M5	2	1	2
Cisco	HyperFlex HX240c M5	2	1	2
Cisco	UCS C480 M5	4	1	6
Dell	PowerEdge C4140	1	1	4**
Dell	PowerEdge C4130	1	1	4**
Dell	PowerEdge R740	2	1	3
Dell	PowerEdge R740xd	2	1	3
Dell	PowerEdge R7425	2	1	3
Dell	R940xa	4	2	4
Dell	DSS 8440	4	1	10
Dell	R840	2	2	2
Dell	Dell EMC VxRail V570	2	1	3
Dell	Dell EMC VxRail V570F	2	1	3
Dell	PowerEdge T640	5	1	4
Dell	PowerEdge T630	5	1	4
HPE	Apollo pc40 Server	1	1	4
HPE	Apollo sx40 Server	1	1	4**
HPE	SGI 8600 (XA780i Gen 10)	1	1	4**
HPE	Apollo XL190r Gen10	2	2	2
HPE	ProLiant DL380 Gen10	2	1	3
HPE	ProLiant DL385 Gen10	2	1	3
HPE	ProLiant DL380 Gen9	2	1	2
HPE	Apollo XL270d Gen10	4	2	8**
HPE	Apollo XL270d Gen9	4	2	8
HPE	ProLiant DL580 Gen10	4	1	4
HPE	Superdome Flex	6	2	4
Lenovo	ThinkSystem SR650	2	1	2
Lenovo	ThinkSystem SR665	2	1	3
Lenovo	ThinkSystem SR655	2	1	3
Lenovo	ThinkSystem HR650X	2	1	2
Lenovo	ThinkSystem SD530/D2	2	2	4
Lenovo	ThinkSystem SR670	2	1	4
Lenovo	ThinkSystem SR860	4	1	2

Manufacturer	Model	Rack Units	Node per Chassis	Tesla V100
Lenovo	ThinkAgile HX3520-G	2	1	2
Lenovo	ThinkAgile HX3521-G	2	1	2
Supermicro	SYS-1028GQ	1	1	4
Supermicro	SYS-1029GQ	1	1	4**
Supermicro	SYS-1028GR	1	1	2
Supermicro	SYS-1029U / 6019U	1	1	1
Supermicro	SYS-2028GR	2	1	4
Supermicro	SYS-2029U / 6029U	2	1	2
Supermicro	SYS-2029GP-TR	2	1	4
Supermicro	SYS-4028GR	4	1	8**
Supermicro	SYS-4029GP- TVRT	4	1	8**
Supermicro	SYS-7048GR-TR	4	1	4
Supermicro	SYS-7049GP-TRT	4	1	4
Supermicro	SYS-6049GP	4	1	8
ASUS	ESC4000 G3	2	1	4
ASUS	ESC4000 G3S	2	1	4
ASUS	ESC8000 G3	3	1	8
ASUS	ESC8000 G4	4	1	8
ASUS	ESC4000 G4	2	1	4
ATOS	Sequana X1125 blade	1	1	4**
FUJITSU	PRIMERGY CX2570 M4	2	2	4
FUJITSU	PRIMERGY RX2540 M4	2	1	2
FUJITSU	PRIMERGY CX400 M4	2	1	4**
Gigabyte	G191-H44	1	1	4
Gigabyte	T180-G23	1	1	4
Gigabyte	T180-G24	1	1	4
Gigabyte	T181-Z70	1	1	4
Gigabyte	G180-G00	1	1	2
Gigabyte	G250-G51	2	1	8
Gigabyte	G250-G52	2	1	8
Gigabyte	G291-280	2	1	8
Gigabyte	G291-281	2	1	8
Gigabyte	G221-Z30	2	1	2
Gigabyte	G291-Z20	2	1	8
Gigabyte	R281-G30	2	1	3

Manufacturer	Model	Rack Units	Node per Chassis	Tesla V100
Gigabyte	R281-3C2	2	1	3
Gigabyte	G481-HA0	4	1	10
Gigabyte	G482-Z50	4	1	10
Hitachi	HA8000V	2	1	3
Hitachi Vantara	Hitachi Advanced Server DS225	2	1	4
H3C	R4900 G3	2	1	3
H3C	R5200 G3	4	1	10
Huawei	G5500/G530 V5/GP308	4	1	4
Huawei	G5500/G560 V5/GP608	4	1	8
Huawei	G560 V5	4	1	8**
Huawei	FusionServer RH2288H V5	2	1	2
Inspur	NF5280M5	2	1	4
Inspur	NF5280M4	2	1	2
Inspur	NF5288M5	2	1	8**
Inspur	NF5468M5	4	1	8**
Nutanix	NX-3155G-G7	2	1	2
Nutanix	NX-3155G-G6	2	1	2
Penguin Computing	Altus X01114GT	1	1	4
Penguin Computing	Relion 1903GT	1	1	3
Penguin Computing	Relion 2904GT	2	1	4
Penguin Computing	Relion 2908GT	2	1	8
Penguin Computing	Relion XE2112GT	2	1	3
Penguin Computing	Relion XE1114GT	1	1	4
Penguin Computing	Relion XE2114GT	2	1	4
Penguin Computing	Relion XE2118GT	2	1	4
Penguin Computing	Relion XE4118GT	4	1	10*
Penguin Computing	Relion X01114GT	1	1	4
PNY	PNYSER28 Series	2	1	8
PNY	PNYSER48 Series	4	1	8
PNY	PNYSRA48 Series	4	1	8
QCT	QuantaGrid D52BV-2U	2	1	4
QCT	QuantaGrid D52G-4U	4	1	8
Sugon	X795-G30	4	1	8
Tyan	GA88-B8021	1	1	4

Manufacturer	Model	Rack Units	Node per Chassis	Tesla V100
Tyan	GA88-B5631	1	1	4
Tyan	TN76-B7102	2	1	2
Tyan	FA77-B7119	4	1	10
Tyan	FT77D-B7109	4	1	8
Tyan	FT48T-B7105	4	1	5
Tyan	FT48B-B7100	4	1	4
Tyan	FT77C-B7079	4	1	8
VDI-Appliance	IO-110 G3	1	1	4
VDI-Appliance	IO-150 G3	1	1	4
VDI-Appliance	IO-250 G3	1	1	4
VDI-Appliance	IO-275 G3	1	1	4
VDI-Appliance	IO-285 G3	1	1	4

Notes:

*Mixed GPU configurations within a server are not supported.

**With expansion chassis.

Notice

The information provided in this specification is believed to be accurate and reliable as of the date provided. However, NVIDIA Corporation ("NVIDIA") does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information. 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 publication supersedes and replaces all other specifications for the product that may have been previously supplied.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and other changes to this specification, at any time and/or to discontinue any product or service without notice. Customer should obtain the latest relevant specification 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. NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this specification.

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 these specifications will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do 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 specification. NVIDIA does not accept any 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 specification, 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 specification. 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 specification is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

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. 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 NVIDIA terms and conditions of sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, NVIDIA GPU Boost, NVIDIA GRID, NVIDIA GRID vGPU, NVIDIA Kepler, NVIDIA Maxwell, NVIDIA Pascal, and Tesla are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2018, 2019, 2020, 2021 NVIDIA Corporation. All rights reserved.