Certified OEM Platforms for EOL Products

Application Note

DA-09018-001_v02 | July 2019

Document History

DA-09018-001_v02

Version	Date	Authors	Description of Change
01	May 17, 2018	VK, SM	Initial Release
02	July 23, 2019	VK, SM	• Added Tesla M6 and Tesla P100 products • Updated application note to meet current NVIDIA standards

Table of Contents

List of Tables

Overview

This application note lists the OEM server platforms which are currently supported for the NVIDIA GRID® K1, NVIDIA GRID K2, NVIDIA® Tesla® M6, and Tesla P100 products.

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 graphics processing unit (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 × .5inches
- Board power: 300 W (maximum)

Supported Servers

The following tables contain the supported servers and models using NVIDIA GRID K1 and NVIDIA GRID K2 (Table 1), Tesla M6 (Table 2), and Tesla P100 (Table 3).

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
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

Manufacturer	Model	Rack Units	Node per Chassis	NVIDIA GRID K1	NVIDIA GRID K2
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-G2O	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
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

Manufacturer	Model	Rack Units	Node per Chassis	NVIDIA GRID K1	NVIDIA GRID K2
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
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

Manufacturer	Model	Rack Units	Node per Chassis	NVIDIA GRID K1	NVIDIA GRID K2
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***
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

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P100
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
Gigabyte	G250-S88	2	1	8
Gigabyte	R280-G2O	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

Manufacturer	Model	Rack Units	Node per Chassis	Tesla P100
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
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

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 GRID, NVIDIA Kepler, 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

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