PART NUMBER:
VCGRIDK1M-PB

NVIDIA GRID K1

GRAPHICS-ACCELERATED VIRTUAL DESKTOPS AND APPLICATIONS

NVIDIA GRID™ technology offers the ability to offload graphics processing from the CPU to the GPU in virtualized environments. This gives the data center manager the freedom to deliver true PC graphics-rich experiences to more virtual users for the first time.

The NVIDIA GRID K1 by PNY board provides:

GPU Virtualization[1]

GRID boards feature the NVIDIA® Kepler™ architecture that, for the first time, allows hardware virtualization of the GPU. This means multiple users can share a single GPU, improving user density while providing true PC performance and compatibility.

Low-Latency Remote Display

NVIDIA's patented low-latency remote display technology greatly improves the user experience by reducing the lag that users feel when interacting with their virtual machine. With this technology, the,virtual desktop screen is pushed,directly to the remoting protocol.

H.264 Encoding[2]

The Kepler GPU includes a highperformance H.264 engine capable of encoding simultaneous streams with superior quality. This provides a giant leap forward in cloud server efficiency by offloading the CPU from encoding functions and allowing these functions to scale with the number of GPUs in a server.

Power Efficiency

GRID GPUs are designed to provide data center-class power efficiency, including the revolutionary new,streaming multiprocessor, called «SMX». The result is an innovative, proven solution that delivers revolutionary performance per-watt for the enterprise data center.

Maximum User Density

NVIDIA GRID boards have an optimized multi-GPU design that helps to maximize user density. GRID K1 boards, which include four Kepler-based GPUs and 16 GB of memory, are designed to host the maximum number of concurrent users. GRID K2 boards, which include two higher-end Kepler GPUs and 8 GB of memory, deliver maximum density for users of graphics-intensive applications.

24/7 Reliability

GRID boards are designed, built, and tested by NVIDIA for 24/7 operation. Working closely with leading server vendors such as Cisco, Dell, HP, IBM, and SuperMicro ensures that GRID cards perform optimally and reliably for the life of the system.

Widest Range of Virtualization Solutions

GRID boards enable GPU-capable virtualization solutions from Citrix, Microsoft, and VMware, delivering the flexibility to choose from a wide range of proven solutions.

GRID K1 - PRODUCT SPECIFICATION

NUMBER OF GPUS	4 x Entry Kepler Based GPUs
TOTAL GPU CORES	768
TOTAL MEMORY SIZE	16 GB DDR3
MAX POWER	130 W
AUX POWER	6-Pin Connector
BOARD LENGHT	265 mm
BOARD HIGHT	110 mm
BOARD WIDTH	Dual Slot
DISPLAY IO	None
PCIE	x16
PCI GENERATION	GEN3 (GEN2 Compatible)
COOLING SOLUTION	passive

[1] 1. Available for Citrix Xenserver | 2. Feature needs application support

AIRFLOW

VCGRIDK1M-PB
Airflow from right to left

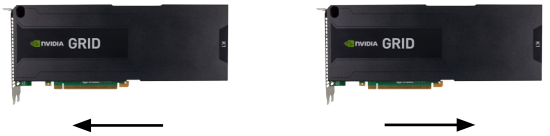

Airflow from left to right

IT MANAGERS CAN NOW

- Leverage industry-leading virtualization solutions, including Citrix, Microsoft and VMware
- Add the most graphics-intensive users to virtual solutions
- Improve the productivity of all users

USERS CAN NOW

- Explore highly responsive windows and rich multimedia experiences
- Access all critical applications including the most 3-D intensive
- Access their most important apps from anywhere, on any device

SOFTWARE PARTNERS

The NVIDIA compatibility guarantee ensures that virtualized users experience the same state-of-the-art graphics they have at their desk. NVIDIA works with over 100 leading companies to ensure this experience meets their stringent application certification standards. A list of these solutions can be found at www.nvidia.com/gridcertifications.

VIRTUAL GPU TECHNOLOGY

The GRID vGPU manager allows for simple management of user profiles with a single-pane-of-glass management console. IT managers can easily assign the optimal amount of graphics memory and deliver a customized graphics profile to meet the specific needs of each user. Every virtual desktop has dedicated graphics memory, just like they would at their desk, so they always have the resources they need to launch and run their applications.
GRID vGPU enables up to eight users to share each physical GPU, assigning the graphics resources of the available GPUs to virtual machines in a balanced approach. Each NVIDIA GRID K1 card has four GPUs, allowing 32 users to share a single card.

	NVIDIA COMPATIBILITY GUARANTEE	APPLICATION CERTIFICATIONS	GRAPHICS APIs SUPPORTED	GRID K1	GRID K2
VIRTUALIZED APPLICATIONS					
Citrix XenApp	✓		DirectX 9,10,11 OpenGL 4.4	✓	✓
VIRTUAL DESKTOPS					
Citrix XenDesktop with HDX 3D Pro and NVIDIA GRID vGPU[1]	✓	✓	DirectX 9,10,11 OpenGL 4.4	Up to 32 Users	Up to 16 Users
VMware Horizon View with NVIDIA GRID vGPU[1]					
VIRTUAL REMOTE WORKSTATIONS[2]					
Citrix XenDesktop with HDX 3D Pro	✓	✓	NVIDIA CUDA DirectX 9, 10, 11 OpenGL 4.4	4 Users	2 High-end Users
VMware Horizon View with vDGA					

[1]NVIDIA GRID™ vGPU™ is supported on compatible versions of Citrix XenServer and VMware vSphere. Consult Citrix and VMware for compatibility.
[2] Path-through, one GPU per user

VGPU PROFILES MEAN CUSTOMIZED, DEDICATED GRAPHICS MEMORY

NVIDIA GRID™ vGPU™ brings the full benefit of NVIDIA hardware-accelerated graphics to virtualized solutions. This technology provides exceptional graphics performance for virtual desktops equivalent to local PCs when sharing a GPU among multiple users. GRID vGPU is the industry's most advanced technology for sharing true GPU hardware acceleration between multiple virtual desktops — without compromising the graphics experience. Application features and compatibility are exactly the same as they would be at the user's desk.
With GRID vGPU technology, the graphics commands of each virtual machine are passed directly to the GPU, without translation by the hypervisor. This allows the GPU hardware to be time-sliced to deliver the ultimate in shared virtualized graphics performance.

Virtual GPU Profile[3]	Application Certifications	Graphics Memory	Max Displays per User	Max Resolution per Display	Max Users Per Graphics Board	Recomended Use Case
K180Q	yes	4 GB	4	2560 x 1600	4	Entry Designer
K160Q	yes	2 GB	4	2560 x 1600	8	Power User
K140Q	yes	1 GB	2	2560 x 1600	16	Power User
K120Q	yes	512 MB	2	2560 x 1600	32	Power User

[3] Graphics cards profiles offered in virtual system. The GRID K1 graphics chip is virtualized to e.g. 2 x K180Q, 4 x K160Q, 1 x K180Q and 2 x K160Q, etc. The virutal graphics card profile is chosen during setup of the virtual system.

NVIDIA | PNY

PNY Technologies EUROPE
sales@pny.eu | T: +33 (0) 556 13 7575