NVIDIA GRID™ K1 AND K2
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.

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 3D-intensive

Access their most important apps from anywhere, on any device

---

The NVIDIA GRID K1 and K2 boards provide:

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 high-performance 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.

Specifications

	GRID K1	GRID K2
Number of GPUs	4 Kepler GPUs	2 high-end Kepler GPUs
Total NVIDIA® CUDA® cores	768	3,072
Total memory size	16 GB DDR3	8 GB GDDR5
Max power	130 W	225 W
Board length	10.5"	10.5"
Board height	4.4"	4.4"
Board width	Dual slot	Dual slot
Display IO	None	None
Aux power	6-pin connector	8-pin connector
PCIe	x16	x16
PCIe generation	Gen3 (Gen2 compatible)	Gen3 (Gen2 compatible)
Cooling solution	Passive	Passive

Software Partners

	NVIDIA DRIVER	REMOTE WORKSTATION CERTIFICATION	API	GRID K1	GRID K2
VIRTUALIZED APPLICATIONS					
Citrix XenApp 6.5 with OpenGL 4.3	✔		NVIDIA CUDA DirectX 9,10,11 OpenGL 4.3	✔	✔
VIRTUAL DESKTOPS					
Citrix XenDesktop 5.6 FP1 with NVIDIA GRID Software[1]	✔	✔	NVIDIA CUDA DirectX 9,10,11 OpenGL 4.3	✔	✔
Microsoft RemoteFX in Windows Server 2012			DirectX 9,10,11 OpenGL 1.1	✔	✔
VMware View 5.2 with vSGA[2]			DirectX 9 OpenGL 2.1	✔	✔
VIRTUAL REMOTE WORKSTATIONS					
Citrix XenDesktop 5.6 FP1 with HDX 3D Pro	✔	✔	NVIDIA CUDA DirectX 9, 10, 11 OpenGL 4.3	4 Users	2 Power Users

OEM Systems Partners

For more information or to purchase available systems, visit **www.nvidia.com/vdi**