

NVIDIA GRID™ OVERVIEW

Imagine if responsive Windows and rich multimedia experiences were available via virtual desktop infrastructure, even those with intensive graphics needs.

NVIDIA makes this possible with NVIDIA GRID, a virtualized GPU that delivers a faster, interactive graphical experience to a variety of users across all devices.

WHY GPUS FOR VIRTUAL DESKTOPS AND APPLICATIONS?

Your smartphone has a GPU, your Tablet has a GPU, your PC has a GPU and your workstation has a high powered GPU. There's a reason for this. The users expect a certain level of visual experience that you can't get from a CPU alone. GRID enables VDI for everyone.

WHAT IS A VIRTUAL DESKTOP?

Virtual Desktop Infrastructure (VDI) is the practice of hosting a desktop OS within a virtual machine (VM) running on a hosted, centralized or remote server. It allows IT departments to offer a secure "access anywhere" desktop experience to end users.


WHAT IS A VIRTUAL APPLICATION?

A virtualized application, more commonly referred to as a virtual user session allows multiple users to run applications in a single OS (ie: Windows Server).

WHY NVIDIA?

Continuing our tradition of pioneering technology, NVIDIA's Kepler Architecture is the world's first virtualization enabled GPU. NVIDIA GRID K1 & K2 boards bring this technology to the widest breadth of partner products and solutions.


SERVER POWERED BY NVIDIA GRID



GPUS ACCELERATE VIRTUAL DESKTOPS AND APPLICATIONS

GRID OVERVIEW

WHO NEEDS NVIDIA GRID™?




DESIGNER/ENGINEER: Traditionally an NVIDIA Quadro® user, this customer creates and works with complicated datasets using graphics intensive applications (3D design, medical diagnostics, etc.).

Related industries: Oil & Gas, Manufacturing, Media & Entertainment, Medical Imaging, Public Sector

Application examples: Dassault Systèmes (CATIA, SolidWorks), Enovia, Siemens NX, Autodesk (3ds Max, Inventor, Maya, etc.)*

*Plus all the same applications a Power User and Knowledge Worker can use.



POWER USER: A user of visual data (such as 3D images, and 2D graphs and line charts), this customer represents a new business opportunity. Customer may also be someone who tried Desktop and Application Virtualization without GPU acceleration and was not satisfied.

Related industries: Healthcare (nurses' stations, doctors' offices, doctors' tablets), Educational Institutions (target engineering and design schools), Government (simulation and training, Geospatial research), Manufacturing (Product Data Management, Product Lifecycle Management, Manufacturing Floor/Job Site, Support).

Application examples: SolidWorks View, PTC Creo View and View Express, Autodesk Inventor View, Siemens NX Viewer, Team Center Visualization. Manufacturing Product Lifecycle applications, Adobe Imaging applications, medical imaging applications, Microsoft Office PowerPoint Authoring, Video editing applications.**

**Plus all the same applications a Knowledge Worker can use.



KNOWLEDGE WORKER: This is typically an office worker who uses Office applications, email, video conferencing and rich media internet applications.

Related industries: Financial Services (Retail, Commercial and Investment Banking, Insurance), Manufacturing, Life Sciences, Oil & Gas, Media & Entertainment, Telecommunications, Government, Education, Technology.

Application examples: Microsoft Office (PowerPoint graphics), Photoshop, etc.

EXAMPLES OF INDUSTRIES WHERE GRID CAN BE USED




FINANCIAL INSTITUTIONS



GOVERNMENT



HEALTHCARE PROVIDERS



MEDIA & ENTERTAINMENT




MANUFACTURERS




EDUCATIONAL INSTITUTIONS

NVIDIA GRID FOR VDI AND APPLICATION VIRTUALIZATION

NVIDIA GRID PLATFORMS FOR VDI




NVIDIA GRID PLATFORMS FOR APPLICATIONS




GLOSSARY OF TERMS

TERM	DESCRIPTION
Citrix Receiver	A lightweight software application that runs on the user's Windows, Mac, Linux, iOS, and Android devices and connects to a XenDesktop virtual machine or XenApp session in the data center.
Citrix XenApp	Citrix's virtualization product that provides remote session capabilities, delivering remote applications or virtual sessions.
Citrix XenDesktop	Citrix's desktop virtualization product that provides remote desktop capabilities.
Citrix XenServer	The commercial hypervisor from Citrix that allows multiple operating systems to run on a single server node.
Cloud Computing	The use of computing resources (hardware and software) that are delivered as a service over a network. Clouds can be public (delivered over the internet) or private (contained within a company's intranet).
Dedicated GPU	A case where the full GPU is used by a virtual machine, not sharing it with other VMs.
GPU Pass-Through	The remoting technology that connects a virtual machine (VM) to a dedicated CPU. The technology to do this from NVIDIA is known as NMOS (NVIDIA Multi-OS), from the capability to have multiple OS's running on a system, each with their own GPU.
GRID	NVIDIA technology that is a combination of both hardware and software to deliver the ultimate virtualized experience.
GRID K1	NVIDIA's GRID K1 is a graphics card designed specifically for virtualization use cases. It carries four Kepler GPUs and a total of 16GB of video memory, making it ideal for multi-user configurations.
GRID K2	NVIDIA's GRID K2 is a graphics card designed specifically for workstation virtualization. It offers two high-end Kepler GPUs, which can support a total of two high-end designers and engineers. The GRID K2 can also be used for power users and knowledge workers, especially where the IT department wants to deploy a single resource type across the entire enterprise.
Host Machine	A computer on which a hypervisor is running one or more virtual machines is a host machine. Each of those virtual machines is called a guest machine. The hypervisor presents to the guest operating systems a virtual operating platform and manages the execution of the guest operating systems.
Hypervisor	In computing, a hypervisor or virtual machine manager (VMM) is a piece of computer software, firmware or hardware that creates and runs virtual machines.
Kepler	The world's first cloud enabled GPU, and the latest architecture from NVIDIA. NVIDIA products built on the NVIDIA® Kepler™ architecture have names such as NVIDIA Quadro® K5000, NVIDIA Tesla® K20, and GRID K2 (K for Kepler).
Microsoft RemoteFX	Microsoft's product that enables multiple users to have access to the acceleration of a single hardware GPU. It is available as a feature of Microsoft Server 2008 R2 and Server 2012.
NVIDIA Windows Server Driver	The NVIDIA graphics driver that loads on a server OS. This driver works with RemoteFX to deliver graphics acceleration.

GPU pass-through was first introduced in 2009 to deliver multiple hardware accelerated operating systems on workstations.



Kepler, the world's first cloud enabled GPU, has up to 7 billion transistors.



NVIDIA has been supporting, shipping and certifying a Windows Server graphics driver since 2003.

GLOSSARY OF TERMS CONTINUED

TERM	DESCRIPTION
NVIDIA ESX Driver	This NVIDIA driver runs on the ESX (otherwise known as vSphere) hypervisor, which allows VMs running on top of ESX to get access to a shared NVIDIA GPU for hardware graphics acceleration when using vSGA (see definition below).
NVIDIA Driver	NVIDIA's certified driver that runs inside the virtual machine, giving graphics acceleration directly to the applications.
Remote Workstation	A workstation instance that is running from a data center, and delivered via the network to a client device. "Remote" can be on premise to the user's office, or can be off premise to a partner site, travel location, or home user.
Virtual Machine	The instance of an OS that runs on top of a hypervisor, using the abstracted hardware image presented by the hypervisor.
Virtualization	The practice of abstracting the virtual machine (VM) from the physical hardware that it runs on. In practice, virtualization is used primarily to run multiple VMs on a single piece of physical hardware.
> Virtual Desktop Infrastructure (VDI)	The practice of hosting a desktop operating system within a virtual machine (VM) running on a hosted, centralized or remote server.
> Hardware Virtualization	The creation of a virtual machine that acts like the real hardware running underneath the hypervisor, or a subset of that hardware. The software executed on these virtual machines recognizes the underlying physical hardware resources (i.e. the OS can load a native driver for that hardware and interact directly with it).
> Virtual User Session	The practice of hosting applications in a single server OS and delivering those applications to multiple remote users.
> Hardware Virtualized GPU	NVIDIA Kepler-based K1 or K2 board that allows multiple users to share a single GPU and allowing each user direct access to the GPU hardware. This improves user density while providing true performance and compatibility.
> Software Virtualization	Software virtualization acts as an interface between the hardware and the VM, creating the ability to adapt to different underlying hardware configurations. In this case, the software acts like the hardware resource, passing commands to the hypervisor which can execute them either on real hardware or emulated hardware.
> NVIDIA GRID Virtual GPU (NVIDIA GRID vGPU)	Key technology used to enable the hardware virtualization of the GPU which allows multiple virtual machines to interact directly with a GPU. The GRID Virtual GPU manages GPU resources that enable multiple users to share common hardware while improving user density, and creates a true PC experience in the cloud.
VMware ESX	The commercial hypervisor from VMware that allows multiple operating systems to run on a single server node.
VMware View	VMware's desktop virtualization product that provides remote desktop capabilities.
VMware vDGA	virtual Dedicated Graphics Acceleration. VMware's productization of GPU pass-through in ESX. Each GPU in the system is dedicated to a single VM, and the NVIDIA native graphics driver is loaded in each of those VM's operating systems.
VMware vSGA	virtual Shared Graphics Acceleration. The software virtualization of a GPU in ESX, allowing multiple users to access a single GPU for graphics acceleration. In this case, the NVIDIA ESX driver is loaded in the hypervisor.

NVIDIA GRID enables a true PC experience for the hundreds of millions of power users who increasingly want to bring their own devices to work.



The GRID Virtual GPU technology is integrated into a commercially available system level hypervisor, like Citrix XenServer.

NVIDIA GRID™ BOARD SPECIFICATIONS

NVIDIA GRID boards are designed for higher GPU densities, with 2 or 4 GPUs per card. For example, with a GRID K1 card, there are 4 GPUs per card. A server with two x16 PCIe slots can hold 2 of these cards, for a total of 8 fully accelerated VMs.



Feature	GRID K1	GRID K2
Number of GPUs	4 entry Kepler GPUs	2 high-end Kepler GPUs
Total NVIDIA CUDA cores	768 cores	3072 cores
Total GPU memory size	16 GB DDR3 (4 GB per GPU)	8 GB GDDR5 (4 GB per GPU)
Max power	130 W	225 W
Board length	10.5"	
Board height	4.4"	
Board width	Dual slot	
Display IO	None	
Aux power	6-pin connector	8-pin connector
PCIe	x16	
PCIe generation	Gen3 (Gen2 compatible)	
Cooling solution	Requires GRID certified server, passive cooling	
Graphics APIs	Open GL 4.3, DirectX 9, 10, 11, CUDA	