

NVIDIA GRID VIRTUAL PC AND VIRTUAL APPS GPU-ACCELERATED PERFORMANCE FOR THE VIRTUAL ENTERPRISE

Desktop virtualization has been around for many years, but some organizations still struggle to deliver a user experience that stands up to what workers have enjoyed on physical PCs. While IT has traditionally settled for a “good enough” user experience, today’s workforce is more tech savvy and increasingly made up of digital natives who expect a dynamic, multimedia-rich experience.

NVIDIA GRID® Virtual PC (GRID vPC) and GRID® Virtual Apps (GRID vApps) improve virtual desktops and applications for every user, with proven performance built on NVIDIA® Tesla® GPUs for exceptional productivity, security, and IT manageability. The virtualization software divides Tesla GPU resources, so the GPU can be shared across multiple virtual machines running any application.

Here are three powerful reasons to deploy GRID vPC and GRID vApps in your data center.

Reason 1: Every App Is a Graphics App.

Even simple productivity applications found in Microsoft Windows 10 (Win10), Office 2016, web browsers, and streaming video can benefit from GPU acceleration. A recent study showed that the number of applications that use graphics acceleration has doubled since 2012. Today, over 60 percent of enterprise users work with at least one of these applications.¹



Traditional desktop and laptop PCs boost application performance with embedded or integrated GPUs. However, when making the transition from physical to virtual, IT has traditionally left the computer graphics burden—such as from DirectX and OpenGL workloads and video streaming—to be handled by a server CPU.

¹ Data from Lakeside Software’s SysTrack Community, 2017.

This has caused a performance bottleneck that limits scale and compromises the user experience. Now, IT departments can take advantage of GPU technology to deliver a virtual desktop experience that's as good as a physical device.

Reason 2: Be Prepared for the “New Normal.”



Information workers, designers, and engineers can realize up to a **42%** increase in productivity through the use of multiple displays.²

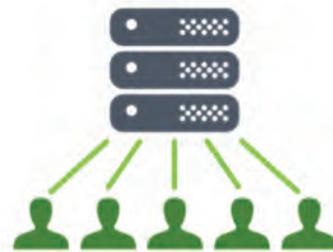
Knowledge workers have changed the way they work. In today's fast-paced, digital world, technology has changed information accessibility. Knowledge workers have multiple devices and expect access to the applications they need, anytime, with seamless experience between those devices. The “new normal” includes multiple devices, multiple 4K resolution monitors, and consistent performance. No longer used only by professionals in the financial services or healthcare industries, multi-monitor setups have quickly become an affordable and effective way to boost productivity.

At the same time, software developers and electronic design automation (EDA) engineers and designers that require Linux-based development environments can also increase productivity by utilizing the like-native experience that GRID vPC provides.

Reason 3: Improve Density while Increasing User Satisfaction.

Today's virtual desktops and applications require graphics acceleration for a native-PC experience. With NVIDIA GRID, you can deliver GPU-acceleration to every virtualized user and support up to 32 concurrent users on a GPU card.³ Because work that was typically done by the CPU has been offloaded to the GPU, the user has a much better experience and more users can be supported.

VMware, Dell, and NVIDIA conducted a side-by-side user testing on a system with and without a GPU, and 99 percent of users rated the GPU-accelerated virtual desktop “like a native PC.”⁴ Despite being loaded with 30 percent more users than the CPU-only virtual desktop, the GPU-accelerated virtual desktop environment provided an incomparably better experience.



30% more users can be supported with a better Win10 VDI user experience with a GPU.⁴

NVIDIA GRID enables organizations to extend the reach of their virtual desktop infrastructure (VDI) investment and give every user a native-PC experience that meets the needs of today's graphic-intensive applications.

² John Peddie Research. October 16, 2017. Jon Peddie Research: Multiple Displays Can Increase Productivity By 42%. Retrieved from <https://www.jonpeddie.com/press-releases/jon-peddie-research-multiple-displays-can-increase-productivity-by-42/>.

³ This assumes the NVIDIA® Tesla® M10 GPU with 1 GB frame buffer per user.

⁴ Results based on side-by-side testing from 136 respondents at VMworld 2017, US. Testing done on WebGL, Google Earth, and YouTube.

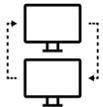
NVIDIA GRID Delivers Productivity Without Boundaries.

Native-PC Experience



Utilize NVIDIA GRID software with Tesla GPUs to deliver responsive, virtual workspaces that perform like native PCs to every user. Now, mobile professionals and knowledge workers can enjoy virtual desktops that are optimized for Windows 10 and office applications.

Multiple Monitor



Support multiple HD displays and up to two 4K-resolution displays to expand your workspace, enhance multi-tasking, and increase productivity.

Video Performance



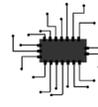
Optimize video performance and scalability through hardware encode/decode via NVIDIA NVENC. Support the latest video-compression standards, including HEVC/H.265 codec.

Improve Server Density



Bypass the challenges of modern applications and operating systems such as Windows 10 that are more graphics intensive and result in lower server density of VDI deployments. Increase user density by 30 percent and deliver a better user experience with NVIDIA virtual GPU solutions.⁵

Linux OS Applications



Deliver the benefits of like-native performance of VDI environments to software developers and EDA designers that rely on a Linux-based development environment.

Investment Protection



Support increasing graphical requirements of Windows 10 and modern productivity applications, with continuous innovation delivered through every vGPU software release.

NVIDIA GRID Feature List

Configuration and Deployment	GRID vPC	GRID vApps
Desktop Virtualization	✓	
Remote Desktop Session Host (RDSH) App Hosting	✓	✓
RDSH Desktop Hosting	✓	✓
Windows OS Support	✓	✓
Linux OS Support	✓ ⁶	
GPU Pass-Through Support ⁷		✓
Bare Metal Support ⁸		✓
NVIDIA Graphics Driver	✓	✓
Guaranteed Quality-of-Service Scheduling ⁹	✓	✓

Display	GRID vPC	GRID vApps
Maximum Hardware Rendered Display	Four HD, Two 4K ⁶	One ¹⁰
Maximum Resolution	4096 x 2160 ⁶	1280 x 1024

⁵ Lakeside Software, Inc. "Elevating User Experience Through GPU Acceleration: A Windows 10 versus Windows 7 Analysis." Lakeside Software White Paper. 2017.

⁶ Support starts with the NVIDIA virtual GPU software March 2018 release (version 6.0).

⁷ GPU pass-through is only supported on 1:1 profiles.

⁸ Only NVIDIA® Tesla® M6 hardware is supported as a primary display device.

⁹ Scheduling options include fixed share, equal share, and best effort/time slicing.

¹⁰ GRID vApps supports one 1280x1024 display from the GPU card. However, Citrix XenApp renders to an offscreen buffer, so it can support multiple software-rendered displays at higher resolutions.

Data Center Management	GRID vPC	GRID vApps
Host-, Guest-, and Application-Level Monitoring ¹¹	✓	✓
Live Migration ⁶	✓	✓

Support	GRID vPC	GRID vApps
NVIDIA Direct Enterprise-Level Technical Support	✓	✓
Maintenance Releases, Defect Resolutions, and Security Patches for up to Three Years ¹²	✓	✓

Advanced Professional Features	GRID vPC	GRID vApps
CUDA/OpenCL		✓ ¹³

Graphics Features and APIs	GRID vPC	GRID vApps
NVENC	✓	
OpenGL Extensions, Including WebGL	✓	✓
DirectX	✓	✓

Profiles ¹⁴	GRID vPC	GRID vApps
Max Frame Buffer Supported	2 GB ⁶	24 GB
Available Profiles	0B, 1B, 2B ⁶	24A, 16A, 12A, 8A, 6A, 4A, 3A, 2A, 1A

NVIDIA Virtual GPU Hardware

GPUs	GRID vPC	GRID vApps
Best Density and Performance for the Knowledge Worker	M10	M10
Best for Blade Form Factor	M6, P6	M6, P6
All Supported GPUs	V100, P100, P40, P4, P6, M60, M10, M6	V100, P100, P40, P4, P6, M60, M10, M6

¹¹ Application-level monitoring is only available starting with the NVIDIA virtual GPU August 2017 release (version 5.0).

¹² Listed support is available with an active Support, Update, and Maintenance (SUMs) contract.

¹³ CUDA/OpenCL is only supported for NVIDIA Maxwell™ 8A profile on NVIDIA GRID 4.x and earlier releases.

¹⁴ Profiles supported have dependency on GPU selected. For more information, read the [NVIDIA virtual GPU software user guide](#).