

VMWARE HORIZON AND NVIDIA GRID vGPU

Q. What is NVIDIA GRID vGPU?

A. GRID vGPU is a graphics acceleration technology from NVIDIA that enables a single GPU (graphics processing unit) to be shared among multiple virtual desktops. When NVIDIA GRID cards (installed in an x86 host) are used in a desktop and app virtualization solution running on VMware vSphere® 6.x, application graphics can be rendered with superior performance compared to non-hardware-accelerated environments. This capability is useful for graphics-intensive use cases such as designers in a manufacturing setting, architects, engineering labs, higher education, oil and gas exploration, clinicians in a healthcare setting, as well as for power users and knowledge workers who need access to rich 2D and 3D graphical interfaces.

Q. What are the components of the NVIDIA GRID solution?

A. To implement the NVIDIA GRID solution, you need a supported hypervisor (a minimum version of vSphere 6.0 for vGPU deployments, and a minimum version of vSphere 5.5 for vDGA or vSGA deployments), a virtualization product such as Horizon, NVIDIA GRID software, licensing and support, and certified hardware which includes your NVIDIA data center GPUs.

Q. What NVIDIA GRID software editions are available and what use cases does each edition support?

A. NVIDIA GRID offers three software editions, each perfectly suited to a particular use case:

- NVIDIA GRID Virtual Applications (vApp) – For use with VMware Horizon Hosted Apps and other solution such as Citrix XenApp.
- NVIDIA GRID Virtual PC (vPC) – For virtual desktop delivering standard PC applications, browser, and multimedia.
- NVIDIA GRID Virtual Workstation (vWS) – Ideal for mainstream and high-end designers who use powerful 3D content creation applications like Dassault CATIA, SOLIDWORKS, and 3DEXcite, Siemens NX, PTC Creo, Schlumberger Petrel, or Autodesk Maya. Only vWS includes the NVIDIA Quadro® driver.

NVIDIA GRID software editions are available as a convenient annual subscription or a perpetual license in combination with support.

1. Assumes cost of subscription, NVIDIA GRID software, and hardware, with 3-year amortization of hardware of 2 Tesla M10 cards supporting 87 vApps users 64 vPC users

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

Q. Does NVIDIA offer a software edition targeted for knowledge workers?

A. For knowledge workers, NVIDIA recommends the vPC or vApp editions. Whether you want a virtual PC or virtual desktop, you can run your applications with “PC-like” and consistent performance on any device with cost effective scaling. These software editions offer GPU-accelerated performance for as little as \$2/user/month for NVIDIA GRID vApps and \$6/user/month for NVIDIA GRID vPC.¹

Q. Does VDI implementation for knowledge workers need graphics acceleration?

A. Windows 10 is the most graphically intensive operating system, requiring 32% more graphics consumption than Windows 7. Updated versions of applications used in Windows 10 (i.e. Microsoft Office, Skype) also drive higher graphics consumption than their counterparts in Windows 7. As such, without a GPU to offload additional processing requirements from the CPU, organizations could experience poor acceptance of VDI due to slow performance, reduced feature sets and even applications refusing to launch.

Furthermore, the addition of compute and graphics resources to offload the CPU in a GPU accelerated VDI environment also means more users can be supported on a server. Windows 10 VDI architects can increase server density (in CPU bound cases) by 30% through the inclusion of vGPU.²

Q. For high-end designers, what benefits does the NVIDIA GRID vWS software edition provide?

A. The NVIDIA GRID vWS includes a certified Quadro driver to ensure that users get the same features expected of a physical workstation, including anti-aliasing, realistic models, enhanced application performance, and application certification.

Q. What are the benefits of VMware Horizon with NVIDIA GRID vGPU?

A. VMware Horizon delivers an enhanced user experience with Blast Performance, built on NVIDIA GRID vGPU technology, offering the following benefits:

- 100 percent compatibility with any graphics-enabled application. By enabling access to the NVIDIA graphics drivers that come with the vPC, vApps or vWorkstation license, a VMware virtual machine with GRID now delivers the same user experience as expected from a local machine.

- Cloud-based, PC, and workstation-class graphics without the need to be tethered to a workstation. Deliver wider access to 3D applications across multiple devices and locations.
- Support for real-time collaboration by centralizing large graphics datasets across end users.
- Enhanced workspace productivity for geographically dispersed teams (including designers, clinicians, researchers, and knowledge workers).
- High-end graphics performance with native NVIDIA driver support ensuring that applications behave the same way they did on their previous desktop systems, all within vSphere-based desktop virtualization environments.
- Better cost-effectiveness and scale with sharing of GPU hardware across multiple VMs and users.
- Access to immersive, hardware-accelerated 3D graphics on any device with VMware Horizon.
- Greater protection against data loss from workstations or laptops, with high-value graphics datasets centrally hosted and secured.
- Certification by NVIDIA and VMware, for a growing portfolio of leading ISV applications like ESRI and Siemens provided with the Quadro driver that comes with the vWorkstation license.

Q. How is GRID vGPU different from other graphics-acceleration approaches?

A. GRID vGPU vs. vSGA – While vSGA (Virtual Shared Graphics Acceleration) offers the benefit of sharing NVIDIA hardware across a larger ratio of end users, GRID vGPU offers superior graphics performance through better application compatibility using native NVIDIA drivers, as well as support for newer releases of OpenGL and DirectX. GRID vGPU in vSphere allow for simple management with a single-pane-of-glass management console. This allows enterprise IT departments the ability to deliver a customized graphics to meet individual end-user requirements. VMware recommends GRID vGPU-based implementations of VMware Horizon where the economics and scalability of GPU-sharing are required, while delivering superior performance. vSGA deployment requires an NVIDIA GRID vPC license.

GRID vGPU vs. vDGA – While vDGA (Virtual Dedicated Graphics Acceleration) delivers superior performance with native NVIDIA driver support demanded by designers using dedicated workstations, it is constrained in offering cost-effective scalability. Only one user can be mapped to an individual NVIDIA GPU with vDGA, so

this approach is best applied in high-end use cases where the economics of GPU-sharing is secondary to performance. When dedicated and shared GPUs are needed, GRID vGPU can manage both needs from a single management interface, with better portability across hosts. vDGA deployment requires an NVIDIA GRID vWorkstation license.

Q. What if I've already deployed vSGA or vDGA in my environment?

A. Customers who have deployed vSGA or vDGA with NVIDIA GRID K1 and K2 hardware with NVIDIA Tesla datacenter GPUs are able to migrate seamlessly to GRID vGPU with vSphere 6 Enterprise Plus. Since the NVIDIA Tesla GPUs are the same (for vSGA, vDGA, and vGPU), customers can upgrade to vSphere 6, change their virtual machine settings, and install the NVIDIA graphics driver to migrate to GRID vGPU.

Q. What are some typical use cases for vDGA?

A. Designers can benefit from dedicated graphics acceleration using applications such as 3D design or computer-aided drawing and engineering. Clinicians can access detailed CT scans and MRIs from their mobile stations and devices. vDGA is ideally suited to those environments where the economics of GPU-sharing are not a requirement, and where the use case dictates a level of performance that is equivalent to a dedicated workstation. When dedicated and shared GPUs are needed, GRID vGPU can manage both needs from a single management interface, with better portability across hosts.

Q. Since GRID vGPU is a capability enabled by vSphere 6.x, can I use it to support my Citrix environment instead of VMware Horizon?

A. While existing Citrix XenDesktop environments deployed on XenServer can now migrate to vSphere and reap the benefits of GRID vGPU, there are significant incremental benefits from additionally migrating XenDesktop environments to VMware Horizon. VMware offers a single platform based on an all-VMware solution that spans the device to the data center, with unified access to VDI desktops, packaged apps, RDS-hosted apps and desktops, as well as SaaS-based apps. This all-VMware stack offers end-to-end management and automation that minimizes support and management costs, reducing customer OpEx and TCO compared with fragmented multi-vendor desktop virtualization solution stacks. Learn how VMware offers a Safe Passage for Citrix Customers.

Q. What is the maximum number of monitors supported? How many users can be supported on a single NVIDIA card? How much video memory can be allocated to users?

A. Refer to the table below for these specifications:

| NVIDIA GRI Graphics Board | Virtual GPU Profile | Application Certifications | Graphics Memory | MAX Displays Per User | MAX Resolution Per Display |
|---------------------------|---------------------|----------------------------|-----------------|-----------------------|----------------------------|
| Tesla M60 | M60-8Q | | 8G | 4 | 4096x2160 |
| | M60-4Q | | 4G | 4 | 4096x2160 |
| | M60-2Q | | 2G | 4 | 4096x2160 |
| | M60-1Q | | 1G | 2 | 4096x2160 |
| | M60-0Q | | 512MB | 2 | 2560x1600 |
| | M60-1B | | 1G | 4 | 2560x1600 |
| | M60-0B | | 512MB | 2 | 2560x1600 |
| | M60-8A | | 8G | 1 | 1280x1024 |
| | M60-4A | | 4G | 1 | 1280x1024 |
| | M60-2A | | 2G | 1 | 1280x1024 |
| | M60-1A | | 1G | 1 | 1280x1024 |
| Tesla M10 | M10-8Q | | 8G | 4 | 4096x2160 |
| | M10-4Q | | 4G | 4 | 4096x2160 |
| | M10-2Q | | 2G | 4 | 4096x2160 |
| | M10-1Q | | 1G | 2 | 4096x2160 |
| | M10-0Q | | 512MB | 2 | 2560x1600 |
| | M10-1B | | 1G | 4 | 2560x1600 |
| | M10-0B | | 512MB | 2 | 2560x1600 |
| | M10-8A | | 8G | 1 | 1280x1024 |
| | M10-4A | | 4G | 1 | 1280x1024 |
| | M10-2A | | 2G | 1 | 1280x1024 |
| | M10-1A | | 1G | 1 | 1280x1024 |
| Tesla M6 | M6-8Q | | 8G | 4 | 4096x2160 |
| | M6-4Q | | 4G | 4 | 4096x2160 |
| | M6-2Q | | 2G | 4 | 4096x2160 |
| | M6-1Q | | 1G | 2 | 4096x2160 |
| | M6-0Q | | 512MB | 2 | 2560x1600 |
| | M6-1B | | 1G | 4 | 2560x1600 |
| | M6-0B | | 512MB | 2 | 2560x1600 |
| | M6-8A | | 8G | 1 | 1280x1024 |
| | M6-4A | | 4G | 1 | 1280x1024 |
| | M6-2A | | 2G | 1 | 1280x1024 |
| | M6-1A | | 1G | 1 | 1280x1024 |

Q. What applications are compatible with GRID vGPU?

A. Since GRID vGPU includes the power of the NVIDIA GPU and the long-proven NVIDIA graphics driver, the vast majority of professional applications ranging from engineering, design, media, and entertainment, to oil and gas, and office productivity are supported with NVIDIA graphics.

Q. Where can I find the Hardware Compatibility List for GRID vGPU?

A. NVIDIA GRID vGPU is a capability of the NVIDIA GRID solution which includes the NVIDIA GRID license and a Tesla GPU. The server compatibility list for NVIDIA GRID can be found on the [NVIDIA Website](#).

Q. Where can I get more information?

A. For more information on the VMware Horizon with NVIDIA GRID vGPU solution, contact your VMware or partner account executive, or visit <https://www.vmware.com/products/horizon.html>.

For more information on NVIDIA GRID vGPU with VMware Horizon, visit <http://www.nvidia.com/object/vmware.html>.

For more information on entitlement, packaging and licensing for NVIDIA GRID visit <http://images.nvidia.com/content/grid/pdf/161207-GRID-Packaging-and-Licensing-Guide.pdf>.

