



Virtual GPUs: The new engine of the modern data center

Unlocking productivity for Education and Government institutions

Contents

Delivering a modern education experience.....	2
Run mission-critical applications with complete security.....	3
Accelerate virtual desktops in the data center	3
NVIDIA and HPE: Making it easy to get started with vGPUs	4
Conclusion.....	4

Digital transformation is having a profound effect on virtually every aspect of business and industry. As companies increasingly adopt advanced digital technologies, trends like cloud computing, employee mobility, and big data analytics are forcing dramatically different ways of doing business and disrupting traditional IT delivery models. Today's IT departments are having to fundamentally rethink how they deliver desktops, applications, and graphics to their workforce in order to capitalize on the massive competitive advantages offered by these trends. To survive and stay agile in this evolving landscape, organizations must provide their workers with access to corporate data, applications, and collaboration tools on their preferred devices, while still taking steps to properly protect their most valuable asset – their data.

As cloud computing takes hold, organizations are increasingly transitioning to a hybrid delivery model that merges on-premises IT with cloud to deliver many advantages. The on-demand nature of the cloud offers tremendous opportunities for organizations to rapidly provision infrastructure and apps, increase scalability, improve efficiency, and control costs, while on-premises IT provides enhanced security and control. IT departments are now leveraging virtualization, a software architecture that creates a virtual version of a resource like a server, desktop, operating system, storage, or network, to enable cloud computing. Virtualization allows organizations to more efficiently use IT resources, enhance the flexibility of existing hardware, reduce costs, and move toward a more agile IT model that provides computing resources on-demand.

Also serving to further complicate the business IT landscape is the proliferation of personal devices. Many modern organizations are struggling to manage an end-user computing environment that suddenly consists of a wide variety of both traditional endpoints and personal devices. Today's users expect the same seamless access to data and technology that they have become accustomed to in their personal lives, so organizations must quickly adapt their IT approaches to deliver the on-demand availability that their users expect. Providing this rapid access to data for users, regardless of their location, can result in tremendous productivity gains for both organizations and employees – but only if IT can figure out a simple, secure, and cost-efficient way to do it.

Virtual desktop infrastructure (VDI) can be an effective way to deliver desktops and apps to workers; however, it can fall short when it comes to users that rely on complex graphics or video rendering. For users who regularly work with graphics-intensive applications, graphics processing unit (GPU) acceleration is a must. And as IDC points out, many of the modern applications commonly used by organizations today (Microsoft Office 2016/365, various web browsers, video/unified communication) increasingly require high-end graphics capabilities, which is driving the need for GPU utilization across the board. Virtual GPUs (vGPUs) render graphics on a host server rather than on a physical endpoint device in order to meet the need for a highly immersive graphics experience while still providing today's workers with the ability to be mobile. With vGPUs, organizations can deliver the performance and power needed to accelerate their workers' most demanding workloads in a virtual environment that IT has complete control over.

In this paper, we will examine two vertical markets – higher education and government – that are rapidly adopting vGPUs to streamline IT management and enhance the user experience.

Delivering a modern education experience

Today's student body is highly mobile, social, and connected, and they carry with them more personal technology than ever. Recent research indicates that the average college student brings anywhere from three to seven internet-connected devices with them to campus, including smartphones, tablets, and laptops. Students not only expect to be able to use their preferred devices exactly like they do at home, but they also increasingly want to use their personal technology for learning. For example, engineering students may want to access sophisticated computer-aided design (CAD) programs on a laptop in their dorm room. Art or photography students may need to use image-editing software like Photoshop on a tablet in order to complete their assignments. And regardless of concentration, it's a fact that most online students will use a mobile device at some point during their education to access and complete coursework.

As students ramp-up their mobile learning practices, many university IT departments are struggling to manage the rapid influx of connected students wielding a variety of different devices. They are facing challenges like increasing security across a growing quantity of endpoints, meeting expectations for wireless connectivity across campus, and supporting an expanding number of distance learning programs. Many schools are fighting to upgrade aging labs, or build additional lab space to accommodate an expanding curriculum and alleviate the scheduling demands on existing classrooms. On top of that, every computer on campus must be kept up-to-date with proper application licenses, patches, maintenance, and upgrades.

By leveraging vGPUs, universities can dramatically simplify the management of their IT resources, reduce IT costs, and discover more flexible ways to use lab areas and classrooms. For students, vGPUs deliver an experience that feels native and “just works,” enabling them to realize power and performance for their educational workloads regardless of whether they are sitting in a coffee shop or in a lab on campus. This anywhere, anytime access to coursework and applications untethers students from classrooms and allows them to take advantage of borderless

learning. With this increased flexibility, students are able to complete their work on their own schedules and on their preferred devices, which broadens the availability of higher education to a much wider range of students.

Stevens Institute of Technology in Hoboken, New Jersey recently leveraged NVIDIA vGPU software to fundamentally change how they deliver technology to students. In the past, the university would issue workstation-class laptops to each of its new 800-900 freshmen, each loaded with the applications they would need for their courses. However, this practice quickly became costly and inefficient, and the laptops became increasingly harder to manage over time. To alleviate these problems, the school began to architect a scalable and flexible virtualized infrastructure that would provide fast, reliable access to the environments and applications needed by students and faculty from any desktop, laptop, or mobile device.

“This major change to how we deliver technology to our students will easily save us more than a million dollars per year. With this new system, there are few limits to how we can innovate education and delivery.”

– VP of IT and CIO David Dodd, Stevens Institute of Technology

The university's new graphics-accelerated virtualized environment consisted of HP blade servers equipped with NVIDIA GRID™ K2 cards for users running graphics-intensive applications. Four additional blade servers equipped with NVIDIA GRID™ K1 cards allowed access to Adobe Creative Suite, while other servers provided non-accelerated access to Windows desktops with standard office applications. Today, the school's technologists estimate that the new environment has saved them over a million dollars per year, and freed them from the daily management of a fleet of laptops so they can think more strategically about how to deliver innovative educational solutions to both local and remote students.

Run mission-critical applications with complete security

Every day in the government sector, agencies run mission-critical workloads that have a completely different set of requirements than enterprise applications. Data security is the primary concern for this vertical market, which consistently operates in an environment where a breach in cybersecurity could threaten national security. Whether a Cabinet-level department or a Humvee conducting a top-secret mission in a remote area, the government sector constantly faces the proposition of how to secure all of their confidential data. Historically, IT departments in this sector would physically remove hard drives from workstations and manually secure them at the end of each day, but with virtualization, they now have the ability to move that data back into the data center to reduce the risk of moving it out to endpoint devices.

Many users in the government sector are considered power users, because they regularly use highly graphical applications in order to do their jobs. These users rely on high power, performance, and availability from their computing applications. Historically, it was simply not an option to virtualize graphics-intensive applications such as VISSIM or Geographic Information System (GIS) software. While the government sector has been moving toward virtualization for years for the security benefits, the challenge has always been that they still need the power and performance to run these highly visual applications. The Windows 10 mandate is another key factor driving virtualization adoption in the government sector. In February 2016, the U.S. Department of Defense (DoD) directed all agencies currently on legacy operating systems to standardize on Windows 10 by spring of 2017. This effort will help a number of DoD agencies abandon the outdated technology that they have relied upon for years, and enable them to transition to more efficient operating systems – ones that will be positioned to support virtualization capabilities and hybrid IT models.

Legacy virtualization solutions simply cannot address the demands of mission-critical applications. vGPU solutions are helping government agencies use graphics-accelerated virtualization to provide users a fast, powerful, and secure way to streamline workflows. Moving workstation and desktop capabilities into the data center provides the security needed by users in this sector, because all processing and visualization takes place on the server with only encrypted pixels delivered to the analysts and monitors. Graphics virtualization also helps to lower costs, allowing large, bulky workstations to be replaced with inexpensive thin clients and enabling IT administrators to remotely manage to save on maintenance costs.

Accelerate virtual desktops in the data center

NVIDIA virtual GPU solutions bring graphics acceleration to the data center, enabling IT to extend the reach of virtualization to every user while delivering a user experience that rivals physical PCs. Housed in the data center, NVIDIA vGPU products run powerful graphics-intensive applications that reside in the cloud rather than on the endpoint device itself. This new “graphics-accelerated data center” is the realization of a fully-virtualized end user computing environment, where NVIDIA GPU technology powers every desktop, every app, and every workstation,

making virtualized workspaces accessible to an entire student body or government agency. NVIDIA vGPU solutions are designed to meet the increasing graphics demands of today's modern applications and Windows 10 operating systems, and eliminate constrained workflows that limit user productivity.

NVIDIA vGPU solutions address many of the challenges faced by IT departments as digital technologies rapidly advance and the preferences of their workforce fundamentally change.

Increasing business agility

Survival in a highly competitive atmosphere demands greater agility to stay ahead of the competition, and in many cases, it's up to IT to deliver. NVIDIA vGPU solutions help liberate users from the constraints of PCs, workstations, offices, and classrooms, transforming workflows by allowing teams to access the data they need from any location and on any device of their choosing. NVIDIA helps organizations eliminate serialized workflows that inhibit agility, and collaborate in real-time without borders or limits.

Supporting a diverse workforce

Today's IT organizations must support a broad and diverse workforce that includes remote, distributed offices, as well as mobile users. In addition to raising the bar on productivity and user experience, NVIDIA vGPU solutions allow IT teams to centralize the management of data and applications in the data center to deliver virtual workspaces with simplified manageability.

Strengthening data security

As the cybersecurity attack surface expands and threats become increasingly sophisticated, organizations must ensure their most valued assets are protected from theft or loss without constraining productivity. NVIDIA vGPU solutions improve overall security while ensuring data and intellectual property are centrally stored in the data center. Data is continuously protected from breaches, loss, or damage, enabling organizations to securely collaborate without risk.

NVIDIA and HPE: Making it easy to get started with vGPUs

NVIDIA's industry-leading GPU technologies are the platform of choice among the most demanding computer users in the world, from scientists, to designers, artists, and gamers. Since NVIDIA's invention of the GPU in 1999, GPU computing has moved beyond the PC gaming market and evolved into a disruptive technique that is pioneering a new, supercharged form of computing. Organizations are increasingly turning to NVIDIA vGPU solutions to build powerful, reliable, and professional virtual environments that provide a user experience that is equivalent to a desktop or workstation.

NVIDIA is making it easier than ever for organizations to [get started with vGPUs](#). NVIDIA currently offers three vGPU software editions that are each perfectly suited to a specific use case:

- [NVIDIA GRID™ Virtual Apps](#) (GRID vApps): Ideal for PC level applications and server-based desktops.
- [NVIDIA GRID™ Virtual PC](#) (GRID vPC): For users who want a virtual desktop but need great user experience leveraging Windows applications, browsers, and high definition video.
- [NVIDIA® Quadro® Virtual Data Center Workstation](#) (Quadro vDWS): For users who want to be able to use remote professional graphics applications with full performance on any device, anywhere.

Organizations can also choose the licensing and support option that is best for them, either as a convenient annual subscription or a perpetual license in combination with [software, updates, and maintenance support \(SUMS\)](#). NVIDIA virtual GPU software runs on [NVIDIA® Tesla® GPUs](#) based on the NVIDIA® Pascal® and Maxwell® GPU architectures and is supported on all Pascal GPUs with the NVIDIA® Tesla® P6, and NVIDIA® Tesla® P40 recommended for enterprise deployment.

[Hewlett Packard Enterprise \(HPE\)](#) is a certified NVIDIA virtualization partner offering a selection of enterprise-grade platforms for organizations that want to deliver high-performance virtual workspaces. HPE's best-selling [ProLiant DL380 Gen10 servers](#) now support NVIDIA vGPU software with NVIDIA® Tesla® M10 GPUs for graphics virtualization, and accommodate up to three double wide or five single wide GPUs for workload acceleration. HPE's industry-leading [high performance computing \(HPC\) platforms](#) deliver better user densities than competitive server solutions, which helps to dramatically cut the costs of deploying virtual machines. All HPE HPC platforms are also capable of supporting the fifth version of NVIDIA GRID, the latest software upgrade which was announced in September 2017.

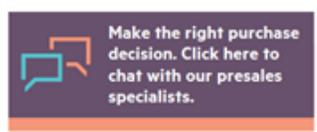
Conclusion

Today's IT departments are tasked with supporting a highly mobile and diverse workforce, and this is especially true in the higher education and government sectors. vGPUs are addressing this challenge by providing seamless access to data, enhancing user productivity, and driving dramatically higher levels of speed, scalability, and security. NVIDIA vGPU solutions bring the power of GPU computing to virtual desktops, apps, and workstations, allowing organizations to deliver an immersive, high-quality user experience for every user, regardless of location. When

combined with HPC server platforms from partners like HPE, organizations can leverage graphics acceleration for a variety of modern workloads. With these best-in-class solutions, organizations can deliver cost-effective performance that scales, transform workflows, and liberate users from the confines of PCs and offices – allowing them to collaborate in real time, from anywhere, and on any device.

Learn more at

<http://www.nvidia.com/virtual-gpu.html>



Sign up for updates



© Copyright 2017 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

The NVIDIA logo, GRID, Quadro, Tesla, Pascal, and Maxwell are trademarks of NVIDIA Corporation in the U.S. and other countries. All other third-party trademark(s) is/are the property of their respective owner(s).

a00036205enw, November 2017