



Modernizing healthcare workloads with virtual GPU technology

Accelerate mobility, productivity, and security

Contents

Adding the power of GPUs to virtualization environments	2
Drive mobility while simplifying IT management.....	2
A consistently exceptional graphics experience for every user	3
Complementary technologies that deliver the best virtual experience.....	3
Conclusion	4

The rapid upsurge of smart personal devices has forced organizations in every industry to fundamentally rethink the way they deliver IT applications and services to a progressively mobile workforce. In the healthcare sector, clinicians are increasingly using mobile devices such as smartphones, tablets, and personal laptops to access patient records, analyze medical images, and share information that allows them to better diagnose and treat patients. This has left hospitals to seek new IT delivery models that support this productivity paradigm shift, and offer greater mobility, security, and IT manageability. Many have turned to virtualization to quickly and cost-effectively provision many high performance, secure virtual workstations using the compute power of a single physical hardware resource.

Rapidly expanding security requirements around patient data and the proven benefits of digitizing hospital operations are the chief drivers of virtualization adoption in the healthcare sector. Hospitals are striving to adhere to federal regulations that mandate meaningful use of electronic health records (EHRs), but simultaneously must be mindful of data security in order to meet Health Insurance Portability and Accountability Act of 1996 (HIPAA) and Health Information Technology for Economic and Clinical Health (HITECH) Act regulations that govern the use and protection of medical data. At the same time, today's hospitals are expected to deliver more enhanced, patient-centric services while operating under the same strict capital budgets that have persisted for years. Virtualization is proven to not only transform the bottom line by lowering the costs of IT delivery to an ever-expanding array of complex endpoints, but also drive profitability by allowing doctors to see and treat even more patients in the same amount of time.

Until recently, clinicians had to log in and out of a dedicated physical workstation in order to view the most current and accurate medical images. This was not only often inconvenient from a location perspective, but quickly added up to hours of lost productivity as they stopped to authenticate themselves into a shared workstation countless times per shift. Virtualization allows doctors and nurses to quickly access critical patient data and detailed medical images in reference quality from wherever they are, and on a variety of different devices ranging from traditional endpoints (desktops and laptops) to personal devices (smartphones and tablets).

Adding the power of GPUs to virtualization environments

Virtualization has had a strong foothold in the healthcare industry for well over a decade, but hospitals are increasingly looking to incorporate the power of graphics processing units (GPUs) into their virtual workloads. Some want to bolster their existing virtualization environment, and explore new ways to streamline and optimize the traditional applications they are already using every day. Others are planning to upgrade to new versions of their software or deploy new applications, and want to ensure they will have the compute and graphics power to reap all the advantages of the latest features these programs should offer. GPUs have also become a standard hardware component for any hospital using picture archiving and communication system (PACS) imaging technology. But regardless of the motivation, every hospital wants to create a virtualization environment that increases their employees' ability to be mobile while still gaining an immersive graphics experience, and do so in a way that keeps patient data protected according to federal regulations.

To deliver this environment, organizations need to move beyond the notion of assigning a single GPU to each user, and instead leverage virtualization to enable multiple workers to share them. Virtual GPU (vGPU) technology accomplishes this by divvying up a server-based GPU among multiple virtual machines (VMs). With virtualized graphics, workers can access their 3D and graphics-intensive programs and applications regardless of their location, while patient data remains more secure in the data center. They gain the freedom to choose which device they want to work on, whether that's a desktop, laptop, tablet, or smartphone. While vGPUs can support the high-performance applications used by the most demanding workstation users, they are now in use by companies in a variety of industries that wish to deliver an exceptional graphics experience for a wide range of users.

Drive mobility while simplifying IT management

Today's highly mobile breed of clinicians rely on the ability to immediately access data from any location in order to better treat patients, and technologies that deliver this rapid, secure data access can drive productivity gains and positively impact patient care. Picture a cardiologist that has just been awakened in the middle of the night and alerted to an emergency with one of his patients. Not long ago, the cardiologist would have had to get up, dress, drive to the hospital, log in and view the latest scans, and only then could he begin developing a game plan to help that patient. Now, using virtualization and virtualized graphics, the cardiologist can access and view the latest cardiac CT scans from a tablet on his nightstand, quickly form an opinion, and have on-duty colleagues set a plan in motion before he has even left his home.

In an environment like this where every second counts, it's critically important to bring the data needed to power these life-saving decisions right to the clinicians' fingertips. vGPU technologies accomplish this by offering doctors the chance to access the most current and accurate documentation that they need to do their jobs, speeding the decision-making process and increasing the effectiveness of collaborative medicine.

From an IT perspective, vGPU solutions dramatically simplify the management of multiple VMs since all resources are managed from a central location, versus having to monitor many different applications running on many different machines. IT teams can manage vGPU solutions end-

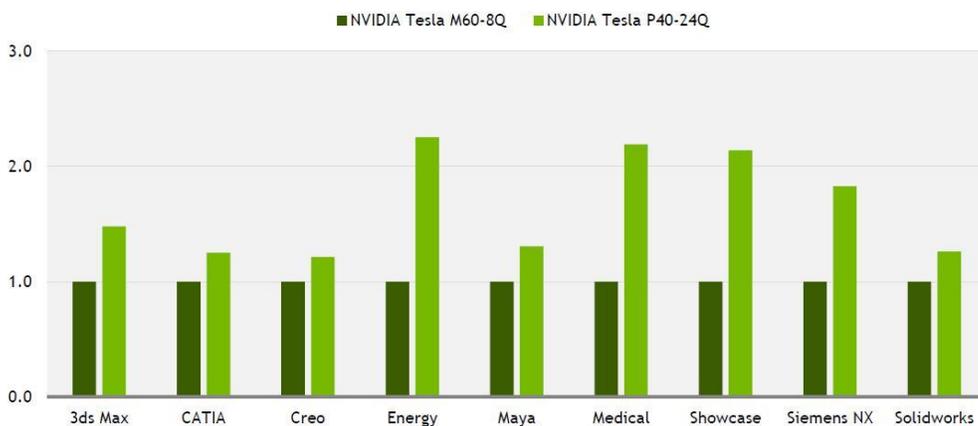
to-end using new monitoring tools that provide complete visibility into their infrastructure at the host, guest, or application level. This enables IT to become more responsive and agile, which ultimately leads to a better end user experience. For IT teams, the ability to centralize version control is very important in a situation where multiple specialists might be conducting different studies simultaneously on the same patient. And from a security perspective, hospitals can confidently enable global collaboration across their workforce while knowing that large, sensitive patient datasets are centrally stored, managed, and secured within the data center.

A consistently exceptional graphics experience for every user

NVIDIA™ virtual GPU solutions extend the power of NVIDIA’s industry-leading GPUs to virtual workstations, desktops, and applications to deliver a consistently exceptional graphics experience for the most demanding graphics-intensive applications. In addition to offering the same immersive, graphics-accelerated experience found in billions of today’s computing devices, NVIDIA virtualization solutions help IT teams centrally manage data and applications, which results in improved security, productivity, and simplified manageability. NVIDIA’s virtual GPU solutions include three products designed to meet the challenges of today’s digital workplace: the Quadro® Virtual Data Center Workstation (Quadro® vDWS), the NVIDIA GRID™ Virtual PC (GRID vPC), and NVIDIA GRID™ Virtual Apps (GRIC vApps). All three offerings run on NVIDIA® Tesla® G_____ PUs based on the NVIDIA Pascal™ and Maxwell® architectures. NVIDIA’s virtual GPU products enable virtual machines to have simultaneous, direct access to a single physical GPU to help users achieve the same performance and graphics experience as on a physical workstation or PC.

NVIDIA® Quadro® professional workstations lead the industry as the most advanced and trusted accelerator of professional workflows by engineers, architects, and scientists. With new Quadro® Virtual Data Center Workstation (Quadro® vDWS) software users can turn NVIDIA® Tesla® GPU servers into powerful workstations, enabling users to get an unparalleled experience running both accelerated graphics and compute workloads on any virtual workstation or laptop. The latest version of Quadro® vDWS supports all NVIDIA® Pascal GPUs (Tesla® P4, P40, P100, and P6) as well as Maxwell® GPUs, (M10, M60, and M6), meaning users can accelerate their productivity and achieve even greater performance by leveraging the latest set of NVIDIA® GPU technologies. In fact, during recent benchmark tests the NVIDIA® Quadro® vDWS with the Tesla® P40 GPU was shown to deliver up to 2X the graphics performance of GPUs based on the previous NVIDIA architecture.

NVIDIA Quadro® vDWS with Tesla® P40 Delivers up to 2X the Performance*



*Comparing a single VM on NVIDIA® Tesla® M60-8Q versus a single VM on NVIDIA® Tesla® P40-24Q and based on SPECviewperf 12.1 benchmark.

Today’s healthcare professionals require an exceptional graphics experience to be productive, but many healthcare organizations still struggle to deliver a user experience that is comparable to what their workers would have experienced on a physical PC. To combat these challenges, NVIDIA recently announced a new version of their NVIDIA GRID® Virtual PC (GRID vPC) software for virtual PC users. GRID® vPC is the ideal solution for users who want the convenience of a virtual desktop but still require an exceptional user experience for Windows applications (Windows 10, Office 365, etc.), browsers, and high definition video. The latest version of GRID® vPC provides improved user density and scalability, providing up to 50 percent more vGPUs for every physical GPU. With GRID® vPC, organizations can deliver a consistent and predictable user experience while leveraging a cost-effective solution that optimizes digital workplace TCO. Depending on the requirements of the virtualization environment, the NVIDIA GRID® Virtual Apps (GRID® vApps) software is an alternative offering designed to deliver Windows applications at full performance on any device, anywhere.

Complementary technologies that deliver the best virtual experience

For over two decades, NVIDIA has been a pioneer in the field of visual computing, operating with a singular focus on bringing the power of GPUs to the largest technology trends and paradigm shifts. Since NVIDIA invented the GPU in 1999, the company has evolved into a specialized platform vendor that targets artificial intelligence (AI), deep learning, high-end graphics visualization, and high-performance computing (HPC) applications – any developing area where GPU power has become essential and uniquely beneficial. Today, NVIDIA's industry-leading portfolio of hardware and software is helping hospital IT departments deliver the graphics and compute virtualization resources needed to meet burgeoning user demands, and empower healthcare workers to operate from anywhere.

To deliver the ultimate virtual user experience, NVIDIA is partnering with HPC leader Hewlett Packard Enterprise (HPE). HPE's rock-solid server platforms, pioneering technologies, and end-to-end support capabilities are a perfect match for NVIDIA's industry-leading GPU technologies. One of NVIDIA's preferred and certified server vendor for its virtual GPU products, HPE has a long track record of leadership in the Virtual Desktop Infrastructure (VDI) and visualization markets as well as the solution expertise to deliver virtualization solutions that are optimized for each customer's unique situation. HPE's portfolio of high performance systems are designed to support higher user densities than traditional systems. For example, HPE's Synergy compute modules can support up to 84 virtual user sessions in a single Synergy blade. And HPE's wide range of advisory, transformation, professional, and operational services are built to help customers leverage the technologies, partners, and operational foundations that are needed to accelerate their digital journey.

With NVIDIA and HPE, customers can choose the vGPU software that best matches their needs, and easily deploy these software products with NVIDIA® Tesla® GPU cards installed on HPE server platforms. NVIDIA's vGPU hardware and software solutions are fully supported across HPE's scalable server portfolio, and the two companies work together closely to test and validate technologies in order to ensure customers will be able to leverage new graphics technology quickly and confidently. In addition, HPE solutions already support the fifth version of NVIDIA GRID®, the latest software upgrade which was announced in September 2017.

Conclusion

At hospitals all over the world, virtualization is increasing data accessibility while reducing IT costs, enhancing collaboration among doctors and specialists, and providing the levels of data security that allow healthcare companies to survive in a highly regulated industry. As these technologies penetrate the mainstream market, it's become vital to incorporate the power of GPUs to accelerate virtualization environments and offer the enhanced, immersive graphics experience that today's medical professionals need to do their jobs. GPU-accelerated virtual desktops and applications can transform the way healthcare professionals work, untethering them from physical PCs, workstations, and offices and enabling them to discover new levels of mobility and productivity. NVIDIA and HPE are teaming up to offer vGPU solutions that deliver graphics acceleration for a variety of modern healthcare workloads, enabling medical professionals to access data more quickly from anywhere in order to make more informed decisions, speed healthcare delivery, and dramatically transform the way they work.

Learn more at
nvidia.com/nvidia-grid



Make the right purchase decision. Click here to chat with our presales specialists.



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).

a00028223ENW, October 2017