

THE INCREASING NEED FOR GPUS IN VDI

NVIDIA GPUS IN DELL EMC VDI COMPLETE ADDRESS GROWING PERFORMANCE NEEDS

SUMMARY

Enterprise adoption of Virtual Desktop Infrastructure (VDI) strengthened over the past few years as vendor offerings matured with increasingly consolidated and simplified options. What were once just sets of components coordinated into recommended stacks by the various compute, network, storage, and software vendors are now packaged into fully validated end-to-end solution bundles. Dell EMC took a distinct step further in simplicity of adoption and use with its VDI Complete offering, which consolidated Dell EMC hardware and VMware software into bundles with pricing, delivery, and support from Dell EMC as a single, global enterprise vendor.

As enterprises benefit in their ability to adopt and manage VDI from the simplification of the solutions, a top challenge they face is addressing performance needs from users as the graphical and video intensity of their work increases. There are always target users whose application requirements clearly necessitate GPU use as part of the VDI solution. Typical examples include designers and engineers working on graphics-intensive visual modeling. What has changed over the past few years is the graphics and video intensity of the common productivity applications in use by knowledge workers. Online video content and video conferencing are examples of applications now in common desktop use that drive increased GPU needs. Combined with increased graphics use (OpenGL and DirectX) by the Windows 10 operating system, the processing requirements to serve a large portion of target VDI users have increased to where GPUs are more broadly applicable and necessary in VDI solutions than ever before.

NVIDIA, VMware, and Dell EMC have partnered to add NVIDIA GPU-based virtual desktop offerings as part of the Dell EMC VDI Complete solutions. With NVIDIA Tesla M10 and M60 GPUs integrated into hyper-converged Dell EMC hardware stacks with VMware Horizon in VDI Complete, enterprise IT can tailor the performance they deliver to the various roles they target, from knowledge workers to designers and engineers, using a consolidated, integrated solution stack. Enterprise IT can address increasing graphics and video performance needs with a strong user experience to ensure the success of their VDI deployments. Graphics offloading from CPUs to GPUs affords compaction in the overall infrastructure footprint for savings on total cost, as well as provides the option to expand VDI to more users and more use cases.

THE CHALLENGES OF GROWTH IN VDI GRAPHICS REQUIREMENTS

Graphic and video performance is always a critical consideration in VDI deployment planning. A high-quality user experience is critical to VDI success. No matter how quick, robust, and cost-efficient a VDI deployment may be, substantial user rejection of the solution is one of the easiest ways for a VDI roll-out to fail.

For some target user roles, particularly designers and engineers, 3D graphics requirements are so high that they require use of GPUs for VDI viability as a solution. Applications common to those roles, such as Autodesk AutoCAD, Dassault Systèmes SOLIDWORKS, and Esri ArcGIS, simply will not function or are unusable without the use of a GPU.

Those applications have long had clear-cut GPU needs, whereas typical knowledge worker applications normally did not necessitate GPU use. Three factors have combined to change this assumption when working to maximize the scope of roles addressed in a new VDI deployment or grow an existing deployment.

First, the way knowledge workers conduct their roles has evolved dramatically with the availability of broadened and improved applications and content available to them, particularly in image and video media. The consumerization of video has brought online video, such as YouTube, into everyday business use as an enterprise-relevant knowledge repository for research and training on technologies, market data, and customer information, to name just a few types of information. Similarly, video conferencing applications, such as Microsoft Skype, have become an expected available tool in many roles where relationship-building and/or tight collaboration are critical to success. These applications consume substantial processing resources, and without a GPU, performance can slow to the point where the visual lag results in such a poor experience that they cannot be used.

A second driver of GPU adoption in VDI environments is the evolution of rich media website browsing. For example, use of multimedia websites like Google Earth has become a basic assumed tool to which knowledge workers expect to have access. Browsing business-applicable news sites is a requirement in many knowledge worker roles, and those sites have become dynamic and multimedia-intensive with many images per page and automated video playing upon page access. The increased media involved in these everyday web applications increases the importance of planning for graphics capability in VDI deployments.

The third factor that drives a need to revisit VDI graphics performance requirements of knowledge workers is the graphics use (OpenGL and DirectX) of the latest Windows operating system and common productivity applications themselves. Testing by Lakeside Software¹ found that Windows 10 increased CPU consumption by 32 percent when compared to Windows 7, due to increased graphics requirements – as is common in successive operating system versions where the user experience is improved to take advantage of improvements in the average hardware profile. With Windows 7 approaching Microsoft's end of extended support on January 14, 2020, enterprises planning their upgrade to Windows 10 need to factor this into their VDI decision-making. Plus, graphics performance for supporting Windows 10 is a moving target. As Microsoft continues to roll out updates to Windows 10, adding new features and UI framework updates, graphics support may need to adjust as well.

It's not just the Windows 10 operating system that is more graphics intensive, but also the common productivity applications. For example, Lakeside Software found that the Microsoft Office Outlook, PowerPoint, and Excel applications all increased in graphics demands by 53-85% in their use on Windows 10 versus their use on Windows 7.

This means that without a GPU, a VDI user conducting common knowledge worker tasks is likely to have a slower, less responsive experience. Their ability to scroll, view tables and images, watch video, and access websites will feel rough. On the other hand, the addition of GPUs to the VDI infrastructure along with GPU virtualization software can emulate a native, PC-like experience, providing users the performance and freedom they expect.

NVIDIA GPU INTEGRATION FOR VDI GRAPHICS PERFORMANCE

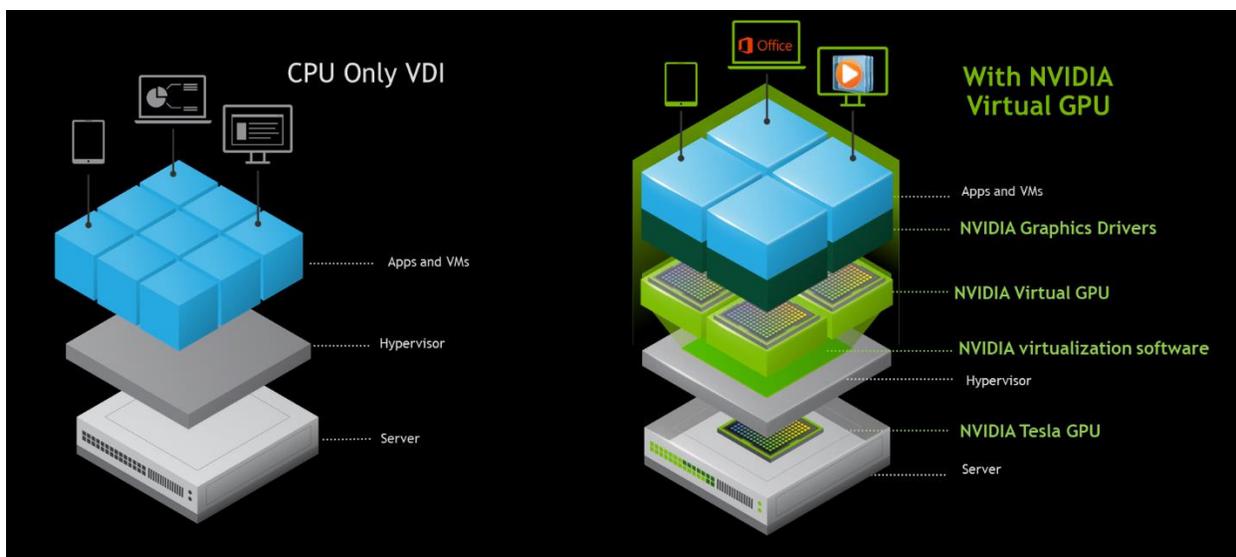
NVIDIA, VMware, and Dell EMC have partnered to add NVIDIA Tesla M10 and M60 GPUs to the Dell EMC VDI Complete portfolio to help enterprise IT address these graphics and video performance needs.

In a VDI environment powered by NVIDIA virtual GPU, the NVIDIA virtual GPU software is installed at the virtualization layer along with the hypervisor. The NVIDIA virtual GPU software creates virtual GPUs enabling every virtual machine (VM) to share the physical GPU installed on the server. The NVIDIA virtualization software includes a graphics driver for every VM. Sometimes referred to as server-side graphics, this method

¹ [Lakeside Software White Paper](#). Elevating User Experience Through GPU Acceleration: A Windows 10 versus Windows 7 Analysis. 2017

enables every virtual machine to get the benefits of a GPU just like a physical desktop. Offloading work typically done by the CPU to the GPU, such as DirectX and OpenGL workloads as well as video streaming, delivers an improved user experience with more effective use of CPU and GPU resources. This allows support for demanding engineering and creative applications in virtualized (and even cloud) environments as well.

FIGURE 1. NVIDIA VIRTUAL GPU DELIVERS A GPU EXPERIENCE TO EVERY VIRTUAL DESKTOP



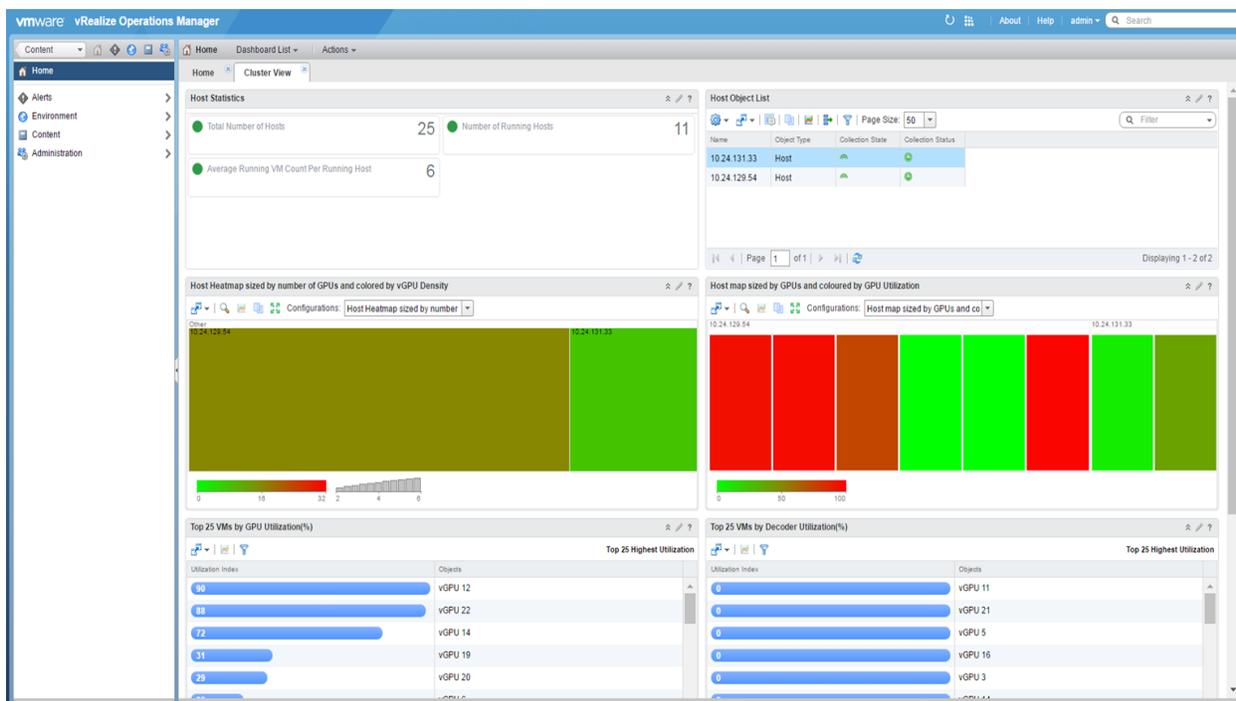
Source: NVIDIA

The NVIDIA GPUs combined with NVIDIA GRID Virtual PC (GRID vPC) or GRID Virtual Apps (GRID vApps) software enable strong graphics and video performance for rich web and productivity applications, while NVIDIA GPUs combined with NVIDIA Quadro Virtual Data Center Workstation (Quadro vDWS) deliver the performance required for graphics-intensive design and engineering applications. Dell EMC also offers higher end NVIDIA GPU options with Tesla P40 GPUs but recommends M10 and M60 for VDI use short of very high-performance needs, such as larger GPU memory allowance or virtualized use of CUDA with VDI. Tesla M10 is the recommended GPU for Windows 10 VDI use for highest density deployments, and M60 is available for higher performance users, particularly for serving a mix of knowledge workers and designers or engineers from a unified VDI infrastructure deployment.

On top of the NVIDIA GPU integration into the easy to adopt and manage model of VDI Complete is the added value of NVIDIA virtual GPU management and monitoring capabilities to provide end-to-end visibility across your IT infrastructure with real-time insight into GPU performance. Broad partner integrations ensure that IT can use the tools they are already familiar with today.

For example, NVIDIA virtual GPU dashboards in VMware vRealize Operations (vROps) provide end-to-end GPU insights for the most comprehensive view of your virtualized infrastructure, supporting IT architects, administrators, and help desk use cases. IT now has a broader view of what is going on in their VM, host, or cluster and the ability to quickly glance at real-time utilization metrics and drill down into individual GPUs or virtual GPUs. Color-coding as shown in Figure 2 below indicates the intensity of GPU resource consumption, directing IT attention to where there may need to be adjustment in a user's GPU resource allocation.

FIGURE 2: NVIDIA VIRTUAL GPU CLUSTER VIEW WITHIN VMWARE VREALIZE OPERATIONS MANAGER



Source: NVIDIA

This level of ongoing visibility eases IT management of the multi-media experience for users. In addition, as individual users and/or roles change in the applications they use,

IT can adjust resource allocations to meet their needs based on available capacity before needing to order additional capacity.

From a cost efficiency perspective, the offloading of processing from the CPU to the GPU, combined with the ability to control and maximize utilization of GPU capacity, allows IT to deploy more VDI users on what is otherwise the same compute, network, and storage infrastructure.

CALL TO ACTION

The performance needs of users have grown substantially based on the increasing use of rich media in common productivity applications, as well as from increased graphics use by Windows 10 and the common knowledge worker applications run on it. In evaluating target roles and organizations for deploying VDI, Moor Insights & Strategy recommends that enterprise IT leaders consider the use of GPUs to address these increased graphics performance needs and deliver a strong user experience necessary to ensure success of VDI deployments.

Moor Insights & Strategy recommends that enterprise IT leaders consider NVIDIA Tesla GPUs with NVIDIA GRID vPC, GRID vApps or Quadro vDWS software within the Dell EMC VDI Complete portfolio for the unique combination of end-to-end solution ease of adoption and management from Dell EMC and the performance and control of graphics acceleration. Dell EMC offers the option of custom VDI solution stacks, for example with alternative options to NVIDIA GPUs and VMware Horizon, which could be of interest to organizations that highly value specific capabilities in those alternatives. It is just tough for the value of any different feature capabilities to override the value of the consolidated single-vendor offer, software tools, and support package from Dell EMC with NVIDIA Tesla GPUs and VMware Horizon.

IMPORTANT INFORMATION ABOUT THIS PAPER

AUTHOR

[Rhett Dillingham](#), Vice President & Senior Analyst at [Moor Insights & Strategy](#)

PUBLISHER

[Patrick Moorhead](#), Founder, President, & Principal Analyst at [Moor Insights & Strategy](#)

INQUIRIES

[Contact us](#) if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

DISCLOSURES

This paper was commissioned by NVIDIA. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

©2018 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.