

CLOUD & GPUS EMPOWER BROADER VDI DEPLOYMENT

INCREASED FLEXIBILITY & PERFORMANCE ENABLE VIRTUAL DESKTOPS FOR MORE OF YOUR ORGANIZATION

EXECUTIVE SUMMARY

IT leaders have deployed virtual desktop solutions for more than a decade in user environments where reliable network connectivity enables the solutions' benefits of increased data security and streamlined management of client computing resources. These leaders wanted to expand the deployments as they grew confident in the consistency of the user experience and comfortable with deployment and management, but they faced two big constraints.

First, legacy virtual desktop infrastructure (VDI) wasn't designed to meet the requirements of the modern digital workplace. Knowledge workers today increasingly rely on graphics and multi-media rich applications on Windows 10, such as streaming video, dynamic (WebGL) websites, and graphics-rich Microsoft PowerPoint and Excel files. These applications and use of multiple monitors can exceed VDI compute capability and significantly impair performance. Plus, VDI historically was not an option for artists, designers and engineers who rely on professional workstations to run 3D graphics applications.

Second, the acquisition and deployment constraints of private infrastructure did not allow ease of trial deployment for additional users or dynamically adjusting deployment sizes for sudden staffing size changes.

The introduction of public cloud-based delivery and GPUs addresses these challenges in a way that naturally complements existing deployment options to unlock much broader deployment and a native-PC experience that meets the expectations of the modern knowledge worker. Solutions deployed on elastic public cloud infrastructure enable low-cost user trials and immediately available capacity for sudden shifts in staffing. Integration of GPUs into virtual desktop solutions – available in both public cloud and private infrastructure – enables IT to address users running even the most challenging of 3D graphics-intensive applications and deliver a native PC experience.

To address these challenges, Dell EMC has partnered with VMware and NVIDIA to offer the full benefits of hybrid infrastructure and GPUs to unlock deployment of VDI for more of your organization and to deliver a robust user experience. Dell EMC VDI Complete Solutions offer Dell EMC VxRail or vSAN Ready Nodes with VMware Horizon and NVIDIA virtual GPU software and GPUs that empower IT to address the full range of user performance needs on private infrastructure. The private deployment can be supplemented with the same VMware and NVIDIA solutions that are running on public cloud servers today from providers such as Amazon Web Services (AWS) and Microsoft

Azure. In this paper we examine how this combination of the flexibility public cloud infrastructure and the cost-efficiency of private infrastructure solutions from Dell EMC enables IT to deliver cost-optimal resource allocation per user with a virtualized user experience that can be indistinguishable from the experience of running natively on a PC.

THE OPPORTUNITY TO EXPAND VIRTUAL DESKTOP USE

The emergence of cloud-based delivery as well as GPUs for virtual desktop deployments offers such an expanded set of options for IT leaders that it demands they revisit where to adopt virtual desktops within their organization. Before assessing the application of these technologies, it is important to consider their emergence in the context of broader technology and user trends faced by IT leaders.

Addressing Workforce Transformation with VDI

IT leaders are commonly facing workforce transformation challenges from a mix of evolving user expectations and formal organizational initiatives. User expectations have shifted substantially with the growth in mobile device capability and access anywhere to their applications and data. This has fueled growth in a cultural norm of more interspersed work and personal time, particularly embraced by the millennial generation that has grown up in an always-connected lifestyle. New initiatives should aim for results that deliver increased mobility in application access with a focus on the safeguarding of data wherever work is done.

Mobility initiatives have previously been focused on knowledge workers, giving access to email on mobile devices or other light workloads, but this work culture shift has carried over to a broader set of roles, up through and including designer and engineer roles requiring high-performance computing capability. In the past, these types of users needed to perform their work on a high-performance workstation in the office without the flexibility of working remotely from home or another work environment. With so many other applications now available across their environments, they want similar flexibility for their core work applications. The same is true for many knowledge workers whose application set has grown more demanding of graphics performance with increasing use of 3D graphics and streaming video (as detailed later in this paper). Senior leadership is typically supportive of the requests based on expected productivity and workforce retention benefits. The cost of multiple high-performance clients means it is infeasible to provide such endpoints across their multiple environments, so centralized virtualization of their applications via VDI is a natural solution.

Another workforce challenge for IT is to better support staffing size flexibility. This is most commonly related to contractors, as their use increases with the trend of outsourcing work not core to an organization's mission and differentiation. It is also applicable to temporary staffing, such as in cyclical or seasonal businesses. This involves not only sudden increases and decreases in the

number of users but continuing turnover in the users within the same roles. This is much easier to address with a centralized, elastic pool of VDI resources complemented by low cost, limited configuration, easy-to-manage clients than a fleet of full PC clients, which need to each be customized and loaded with the software and data that user needs.

The Long-Term Future of the Virtual Desktop

The ever-increasing popularity of browser-based cloud delivery of applications via software-as-a-service (SaaS) is a market trend that runs counter to the value of VDI. Desktop applications converted to SaaS delivery reduce the percentage of applications users need delivered through a virtual desktop, particularly the Windows applications that SaaS delivery tends to replace. This can seem like a reason to question the long-term value of virtual desktops versus a browser-centric client since SaaS can address workforce transformation needs via centralization just as VDI does. However, when reviewing a typical organization's mix of client applications, SaaS adoption is only addressing a portion of the overall set. This is because while SaaS delivery – and even heavily adopted alternative platforms to Windows, such as Android and iOS for mobile devices – are popular for new and even very broadly used enterprise applications, there is not a broad reinvestment in existing Windows applications to move them to those platforms.

In other words, adoption of a SaaS model for office productivity applications and common enterprise applications of CRM and ERP may reduce a substantial number of an organization's applications to be delivered via virtual desktop, but this rarely gets anywhere close to 100%. This is due to the tens of enterprise applications each user role needs from the commonly hundreds to thousands of Windows applications within large organizations. This will continue to evolve but at such a slow pace as to commonly not be a factor on even a five-year infrastructure investment time horizon.

The emergence of SaaS and applications developed for non-Windows application platforms, such as Android, Chrome OS, iOS, and macOS, is more of a motivation to consider tools to manage application delivery consistently across multiple platforms and the virtual desktop, such as VMware's Workspace ONE. This unified endpoint management software complements a VDI platform, such as VMware Horizon, by automating the software configuration of new Windows clients, for example with the VDI client software, and rounds out Windows application delivery options. These tools also address similar security and device management needs across non-Windows platforms, including policy management and a seamless user access interface.

Cloud-based Virtual Desktop Delivery Enables the Flexibility of Hybrid Infrastructure

Public cloud has become a significant factor in IT decision-making for VDI. Initially considered an alternative to private VDI investment based on proprietary service offerings, such as Amazon

Workspaces, public cloud is now recognized as a strong complementary option when applying the same VDI software stack from VMware or Citrix across public cloud and private infrastructure. This also extends to use of NVIDIA GPUs that can span public and private cloud infrastructure. Whether a user logs on to a session hosted in public cloud or on-premises, they get a seamless VDI user experience.

The first reason is the flexibility on-demand cloud provisioning provides for trial use and proof-of-concept deployments when testing VDI with new users. On-demand provisioning allows spinning up virtual machines (VMs) to test with users where adjustments can then be made in real time to the VM resource mix (e.g. more memory, CPU cores, GPU resource, network bandwidth and storage performance) to optimize the performance and consistency of the user experience for the target role. Managed hosting-based DaaS (Desktop-as-a-Service) lacks this flexibility and typically requires significant commitments to justify trial use deployments. Public cloud-based DaaS, such as Amazon Workspaces, introduced this level of flexibility, but the delivery software and protocols are completely different from VMware Horizon and Citrix Virtual Desktop, which risks substantial differences in the user experience and resource profile when moving users from public cloud to private infrastructure, as well as presents difficulty in managing separate platforms.

With VMware Horizon and Citrix Virtual Desktop supported on AWS and Microsoft Azure and with NVIDIA virtual GPU technology now available from those cloud providers, users can reliably be moved from a public cloud deployment to a private infrastructure deployment with minimal disruption. The public cloud providers offer various NVIDIA GPU options, which eases matching private infrastructure GPU selection based on the needs of users to the optimal public cloud. This allows IT leaders to de-risk investment in growth of VDI deployments through trial use on public cloud.

The second reason is the opportunity for cost optimization. Public cloud infrastructure's elasticity of provisioning enables cost savings from optimized utilization. For instance, a VDI deployment can be dynamically scaled with aggregate usage or at the granularity of an individual user VM. While this is possible to the first degree on private infrastructure, particularly when automated on a private cloud platform based on VMware vCloud Suite, OpenStack (e.g. from Red Hat), or Microsoft Azure Stack, public cloud provides much greater elasticity when significantly adjusting capacity. In contrast, private infrastructure is often more cost effective than public cloud for consistent VDI usage patterns under two conditions. First, the aggregate usage profile must be predictable enough to yield high utilization of the infrastructure that is purchased and depreciated over a multi-year timeframe. Second, the organization must have enough maturity and scale in private infrastructure operation that operational costs do not override the cost savings of owning the infrastructure. These conditions are common in large enterprise and government organizations but can vary in smaller enterprises depending on their VDI deployment scale and operational maturity.

A highly complementary hybrid infrastructure model between public cloud and private infrastructure results, enabling ease of trial virtual desktop deployments with optimization of provisioned resources for new users, elasticity to address significant bursts in hiring, and an optimal long-term cost structure of a VDI base on private infrastructure complemented by more variable utilization in public cloud.

GPUs Enable Virtual Desktops for Today's Knowledge Worker and Professional Workstation Users

For professional designers and engineers in industries such as manufacturing, architecture and oil & gas, 3D graphics and compute requirements are so high that they require use of GPUs for VDI viability as a solution. Applications common to those roles, such as ANSYS Discovery Live, Autodesk AutoCAD, Dassault Systèmes SOLIDWORKS, and Esri ArcGIS, will not function or are unusable without a GPU. Adding GPUs to a VDI deployment unlocks IT's ability to address the needs of these high-performance users to broaden VDI's value in data security and reduced client management cost while addressing workforce transformation goals across the full spectrum of users.

Those highly graphics-intensive applications have long had clear cut GPU needs; whereas, typical knowledge worker applications have only more recently evolved to require GPU use. The shift in knowledge worker requirements has been driven by three factors. First, knowledge workers conduct their roles with much more image and video media than in the past. For example, online video, such as YouTube, and online collaboration, such as Microsoft Skype, have become expected available tools in many roles. 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.

Second, dynamic websites heavy on multimedia content and 3D graphics using WebGL, such as Google Earth, have become a basic assumed tool for knowledge workers to access. Business-applicable news sites are commonly referenced in many knowledge worker roles, and those sites are increasingly dynamic and multimedia-intensive, with many images per page and automated video playing upon page access.

Third, Microsoft Windows 10 and common Windows productivity applications, such as Microsoft Office, have become more graphics-intensive to deliver a better user experience by assuming broad commonality of GPU capability per client. Testing by Lakeside Software¹ found that Windows 10 increased CPU consumption by 32 percent when compared to Windows 7 due to increased graphics requirements. Lakeside Software also found that the Microsoft Office Outlook, PowerPoint, and Excel applications increased in graphics demands by 53-85% in their use on Windows 10 versus their use

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

on Windows 7. With Windows 7 approaching Microsoft's end of extended support on January 14, 2020, enterprises planning to move to Windows 10 need to factor in a considerable increase in their planned GPU usage.

In addition, these workers more commonly use these applications in workspace environments with multiple and/or high-resolution (HD 1080P or UHD 4K) monitors, adding to the graphics performance requirements. Based on this combination of factors, IT leaders need to plan for increasing GPU use in VDI deployments going forward. This is necessary for broadening VDI deployments to address more high-end users of graphics-intensive applications, whether designer and engineer types of roles or additional knowledge workers. It is also necessary in order to continue to deliver a strong user experience for knowledge workers already using VDI whose work is growing increasingly graphics-intensive. Fortunately, NVIDIA GPUs are now available in the top public clouds, including AWS and Microsoft Azure, in compute VM types specifically targeted for use with VDI workloads, so much of the value of public cloud trial and hybrid optimization is applicable to VDI deployments with GPUs as well.

DELL EMC AND VMWARE DELIVER A COMPREHENSIVE APPROACH TO VDI

To take full advantage of the VDI deployment options of hybrid infrastructure and GPUs, IT leaders need VDI solutions that offer ease and flexibility in deployment. Whether the starting focus for high performance virtual desktops with GPUs is on private infrastructure or public cloud, and as part of broader application delivery management beyond Windows or standalone, flexibility to grow into multiple deployment types should be the goal. Dell EMC has partnered with VMware to uniquely deliver this VDI flexibility for IT leaders by combining VMware Horizon's hybrid infrastructure portability with the differentiated simplicity of Dell EMC VDI Complete as a single vendor solution for private infrastructure.

Dell EMC and VMware Provide Ease of Portability Across Hybrid Infrastructure

VMware supports Horizon in public cloud on AWS and Microsoft Azure. This common Horizon platform allows ease of VDI deployment portability by using the same software stack across public cloud and private infrastructure. Contrasting this with challenges of different IT management and performance – and therefore user experience – when moving a VDI deployment from AWS's proprietary Workspaces service to Horizon on private infrastructure, it is easy to understand why proprietary DaaS adoption has been driven mostly by smaller organizations less concerned with long-term VDI flexibility or cost efficiency at scale. Leveraging the portability benefit of the Horizon platform, Dell EMC's VDI Complete solution simplifies deployment to private infrastructure by uniquely consolidating pricing, purchase, deployment, and support of pre-validated VxRail-based infrastructure and software stacks.

Dell EMC and NVIDIA Simplify Deployment and Management of GPUs in VDI

Dell EMC offers NVIDIA GPU technology, as shown in Figure 1, just like those available on public cloud. In order to virtualize the physical NVIDIA GPU resources, an NVIDIA virtual GPU software license is required. There are three NVIDIA virtual GPU software editions priced and designed to meet specific use cases:

1. NVIDIA Quadro Virtual Data Center Workstation (Quadro vDWS) software for creative and technical professionals accessing virtual workstations running CAD, CAE and creativity applications.
2. NVIDIA GRID Virtual PC software for knowledge workers accessing modern productivity applications and Windows 10
3. NVIDIA GRID Virtual Applications software for users accessing virtualized applications

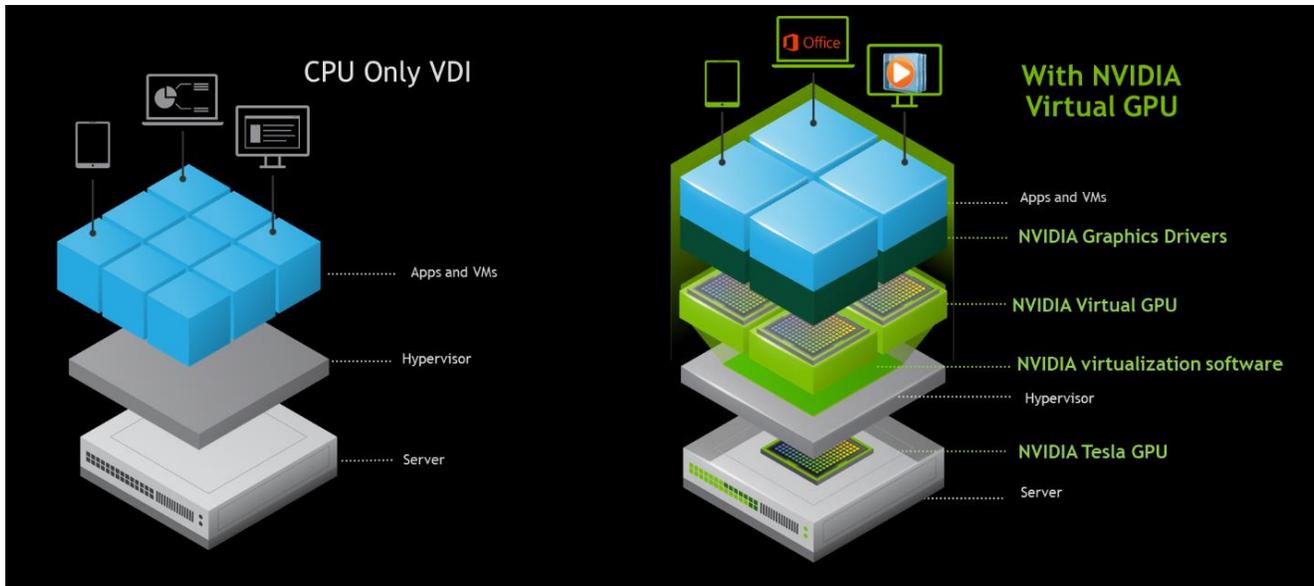
Dell EMC VDI Complete currently offers each of these editions of NVIDIA virtual GPU software, along with NVIDIA M10, P40 and M60 GPUs, offering IT leaders a variety of GPU performance options. NVIDIA M10 GPUs offer the best density for performance needed by knowledge workers. NVIDIA M60 GPUs address entry to mid-range professional graphics users, while NVIDIA P40 GPUs address the mid-range to high-end graphics users.

Dell EMC has recently introduced server configurations offering NVIDIA T4 GPUs that address the same use cases as M60 GPUs with up to double the performance. T4 GPUs also address the same use cases as M10 with the same density in a single slot form factor (compared with dual slot M10) and with significantly lower power consumption. Plus, in cases where the infrastructure is shared between VDI and other workloads, T4 GPUs offer added flexibility² to power machine learning and deep learning workloads across training and inferencing use, as well as rendering workloads.

VDI Complete also uniquely pairs these server-side options with client-side offers of Dell Wyse Thin Clients configurable to address the full client performance and feature range up through high-performance environments for users wanting multiple displays capable of 4K UHD resolution.

² NVIDIA T4 GPU applicability across workloads is detailed here: https://www.nvidia.com/content/dam/en-zz/Solutions/design-visualization/solutions/resources/documents1/TechBrief_T4.pdf

FIGURE 1. NVIDIA VIRTUAL GPU TECHNOLOGY



Source: NVIDIA

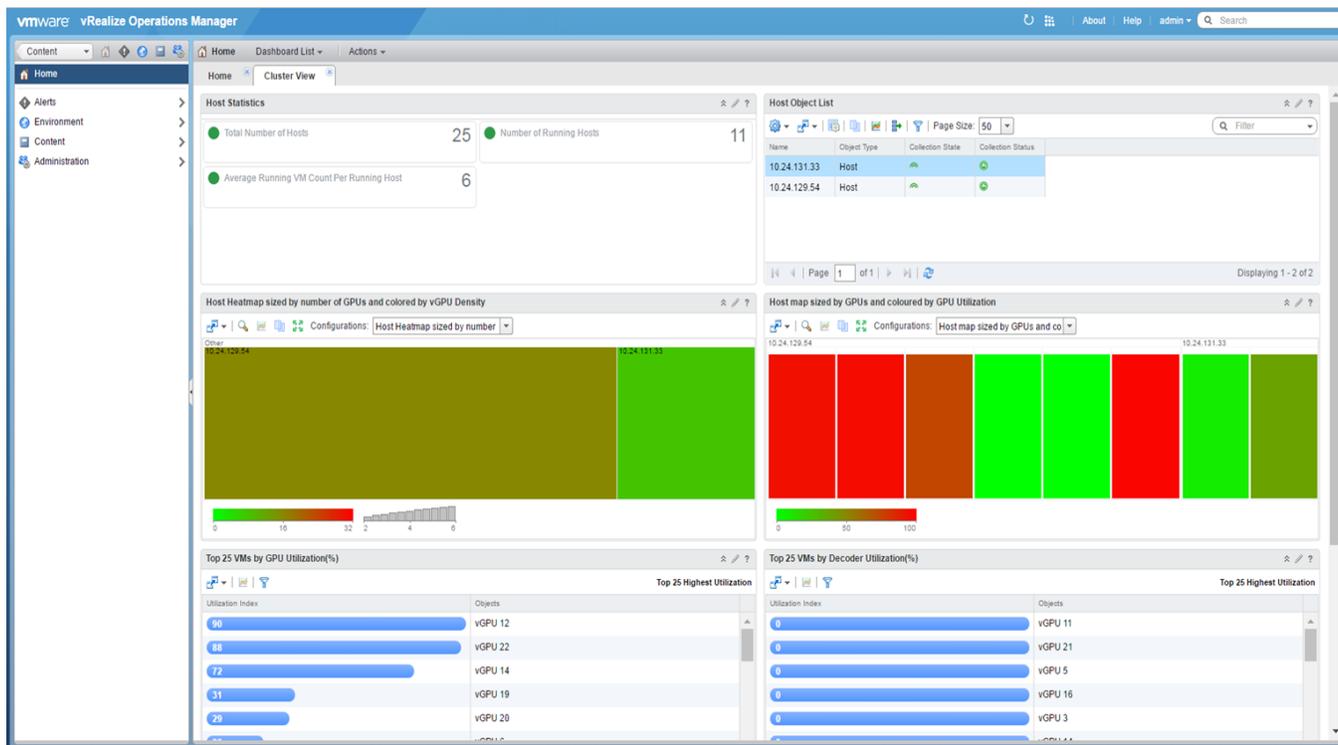
Because NVIDIA virtual GPU software is supported with VMware, the offering complements critical software capabilities VMware added to Horizon in recent years to address these graphics-intensive workloads. This includes its Blast Extreme feature, which delivers improved user experience consistency in conditions of intermittent network connectivity and its vMotion support for live migration of GPU-enabled VDI sessions for fleet management and optimization of capacity utilization. With NVIDIA virtual GPU metrics integrated into VMware vRealize Operations, IT gets end-to-end visibility into the entire virtualized infrastructure, with insights that enable proactive monitoring and quick resolution of issues.

NVIDIA virtual GPU software allows IT to divide up the physical NVIDIA GPU resources across multiple users. When working with larger models and increasingly demanding workflows, NVIDIA Quadro vDWS software enables multiple NVIDIA GPUs to be assigned to a single VM. This ability to optimize allocated graphics resources enables IT to deliver a user experience that can be indistinguishable from running the applications natively on a PC, which is critical to winning user acceptance in broader VDI deployment trials.

NVIDIA virtual GPU software also provides virtual GPU management and monitoring capabilities for real-time insight surfaced in VMware vRealize Operations. IT gets a broad operational view of its 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 VIEW IN VMWARE vREALIZE OPERATIONS MANAGER



Source: NVIDIA

CALL TO ACTION

Workforce transformation and the opportunity to deliver VDI across much more of an organization justifies an in-depth revisit of client computing strategy by IT leaders. In the past, VDI's value of improved data security and reduced cost of client endpoint management has been limited to too few roles across knowledge workers, designers, engineers and others with graphics-intensive applications. Extending VDI into a hybrid infrastructure with public cloud provides IT leaders with great flexibility to trial, optimize and dynamically scale deployments. Adding GPUs to a VDI deployment enables IT leaders to broaden deployment across much more of an organization while establishing compute resource flexibility to address continuing user growth in graphics-intensive processing with GPU resources. Accessing these benefits from public cloud and GPUs requires the right tools to achieve this flexibility.

Moor Insights & Strategy recommends that enterprise IT leaders consider Dell EMC VDI Complete solutions with Dell EMC VxRail servers, VMware Horizon and NVIDIA GPUs with NVIDIA virtual GPU software as a VDI solution to provide this flexibility. VDI Complete offers Dell EMC's unique, single vendor, end-to-end solution, ease of adoption and management of VDI on private infrastructure combined with VMware and NVIDIA's solution stack portability for hybrid infrastructure extension to public cloud. VDI Complete solutions empower IT to broaden delivery of the value of VDI in their organization through trial and optimization to deliver users an excellent virtual desktop experience and their organization a highly cost-efficient infrastructure solution.

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 Dell Technologies. 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.

©2019 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.