Remote User Experience with Graphics Software for Higher Education on Virtual Desktop Infrastructure

Technical White Paper

Abstract

This technical white paper describes performance considerations and best practices for optimizing the remote user experience with several professional graphics-accelerated applications that are typically used in Higher Education environments.

Dell Technologies Solutions

Dell Technologies
Validated Design

Part Number: H18992
January 2022

DELL Technologies

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

1

Introduction

Topics:

Executive summary

Organizations in the Higher Education and Career and Technical Education sectors are adopting virtual desktop infrastructure (VDI) to benefit from centralized management, better security and compliance, and worker mobility. Users in these sectors typically run professional-grade graphics applications that require advanced graphic acceleration by dedicated GPU hardware.

Running professional graphics software on virtualized platforms provides many benefits, including:

- The ability to work or learn from anywhere
- A better return on infrastructure investment
- A quicker time-to-market with faster software development life cycles
- The ability to dynamically meet the development requirements of educational users
- More favorable user experiences

With the changing dynamics of present-day education and the diversity of user situations (such as working from campus or from home), it has become challenging for IT teams to provide a consistent user experience. For example, a user viewing 3D geographic scenes at the end of a remote internet connection — which could be subject to a wide range of distance and performance impairments — might not experience the same image quality as someone viewing them on campus.

This technical white paper describes performance considerations and best practices to optimize the remote user experience on professional graphics-accelerated applications for Higher Education, with a focus on geographic information systems and engineering computer-aided design (CAD) applications. This paper provides details of performance test results for both baseline local area network (LAN) and remote wide area network (WAN) user experiences with virtual desktop implementations of Esri's ArcGIS Pro and an Engineering CAD application, on Dell Technologies VDI with NVIDIA virtual GPUs, using NVIDIA's nVector performance monitoring tool.

Document purpose

This document describes the business challenge, approach, and benefits of implementing an efficient virtual workstation infrastructure for remote working with graphic-enhanced workloads. It describes the specific benefits of using an NVIDIA graphic-enhanced Dell virtual workstation running Esri's ArcGIS Pro and Engineering CAD. The document provides guidance on the expected remote end-user experience as measured by NVIDIA's nVector monitoring tool for several remote users' connectivity scenarios. It also discusses various technical considerations, including GPU and remote protocol options.

Audience

This document is intended for decision-makers, managers, architects, developers, and technical administrators of IT environments in the Higher Education and Career and Technical Education sectors who want to understand Dell Technologies' virtual workstation solutions. Readers should be familiar with Dell VxRail hyperconverged infrastructure (HCI), Dell PowerEdge servers, Esri's ArcGIS Pro and Engineering CAD software, NVIDIA GPU technologies, and VMware Horizon virtualization.

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by email.

Author: Colin Byrne, Solutions Architect

Contributors: Scott Powers, Technical Marketing

2

Business challenge

Topics:

- Market environment

Market environment

Traditionally, designers and engineers requiring 2D/3D visualization and analysis used high-end physical workstations. Deploying, upgrading, and maintaining these workstations is difficult and time-consuming for IT.

Using VDI with virtual workstations instead provides an opportunity to overcome many of the challenges of IT management and maintenance as well as offering better collaboration, security, and flexibility.

Virtualization enables IT teams to deploy and present multiple virtual workstations to a user's endpoint device. The virtual workstations reside in a data center and IT can manage and maintain them centrally. Upgrading operating systems and applications, and applying security patches is easier than in the case of physical workstations.

As the volume and resolution of visual data continues to increase so does the IT administrative burden. PowerEdge servers and NVIDIA virtual GPU (vGPU) technologies provide a stable and robust graphics-accelerated virtualization solution that is designed to meet the increasing needs of mobile users. Using this solution, IT departments in the education sectors can reduce costs and administration, decrease provisioning times, and increase security, while delivering high-quality GPU-enhanced user experiences. Virtualized graphic-enhanced instances can be provided in both on-premises and cloud-based offerings.

A further complication to delivering a high-quality and consistent user experience in the modern higher education sector is the need to support remote learners with the same standards as are available for on-campus students, even when they are subject to a wide variety of distance and performance impacts in their LAN and WAN environments.

Approach

To address these challenges, we simulated and performance-tested several remote-user internet connectivity scenarios using two representative graphics-intensive applications, Esri's ArcGISPro and a leading Engineering CAD software application, running on Dell Technologies VDI.

3

Technology overview

Topics:

Graphics-intensive workloads

Measuring user experience at the VDI desktops or workstation level is informative if the end-users' endpoints are in proximity to the hosting data center, for example the same LAN or campus in the same building or site. However, these measurements are less representative if the user is potentially at the end of a remote internet connection. The ability to measure end-user experience at the actual endpoint device gives a better indication of the remote user experience than relying on data captured at the VDI desktop or workstation.

This document describes the results and analysis of a remote user experience for several professional graphics-accelerated applications and multiple simulated remote network categories as measured by NVIDIA's nVector Lite benchmark performance tool. The following table shows the graphics-intensive workloads that we tested:

Table 1. Graphics-intensive workloads

Graphics application	Description	General user representative workload	Reference
ArcGIS Pro	ArcGIS Pro is a desktop-based Geographic Information Systems (GIS) application that allows users to visualize data, run analysis, manage and maintain data, and work with 2D and 3D data.	ArcGIS Pro Performance Assessment Tool (PAT) benchmarking tool Level 2	ArcGIS Pro GIS with ArcGIS Pro on Dell Technologies Virtual Desktop Infrastructure
Engineering CAD application	A software suite for producing 3D graphics models for the automotive and manufacturing industry.	Engineering CAD benchmarking tool running GSRealworld workload	Engineering CAD Dell Virtual Workstation Solutions for Automotive and Manufacturing
NVIDIA nVector Lite	nVector Lite is a benchmarking tool designed to provide insight into the quality of the VDI end-user experience. The tool focuses on the experience of the user on the endpoint device rather than the response time at the virtual desktop.	nVector Knowledge Worker	NVIDIA: Quantifying the impact of virtual GPUs Dell Technologies: nVector performance test results and analysis

Geographic Information Systems

A Geographic Information System (GIS) is a system that includes a type of database containing geographic, location-based data and which has software tools for managing, analyzing, and visualizing the geographic data in various useful ways.

ArcGIS Pro

Esri's ArcGIS Pro is a leading example of GIS software. ArcGIS Pro is a Windows-based application requiring enhanced graphics support. It is similar in computing and graphics requirements to professional Engineering CAD applications. ArcGIS Pro is a desktop-based GIS application that enables users to visualize data, run analysis, manage and maintain data, and work with 2D and 3D data. It connects directly to ArcGIS Online by using the ArcGIS Enterprise Portal.

ArcGIS users broadly fall into three usage profiles or personas, which can be described as Fundamental User, Intermediate User, and Advanced User. These categories put common ArcGIS Pro tasks and activities into groups. These activities range from map visualization and the use of 2D vector data and light editing, to advanced map visualization using complex datatypes such as raster data and applying advanced spatial analytics with geo-processing tools. These three user personas loosely correlate to the amount of CPU and GPU resources that are required to perform these associated activities.

The following table describes the user personas:

Table 2. User personas

User persona	Typical activities	Licensing
Fundamental User	Interactive map visualization in 2D and 3D, single user data editing and database spatial analysis using vector data, and raster imagery visualization.	GIS Professional Basic
Intermediate User[a]	All capabilities of the Basic license level with the addition of multi-user editing and database management with larger datasets, and higher-end spatial and raster analysis tools.	GIS Professional Standard
Advanced User	All the capabilities of the Standard license level with the addition of high-end cartography and the full range of ArcGIS Pro's spatial analysis capabilities.	GIS Professional Advanced

a. The persona used for the typical user's representative configuration for testing.

For more information, see User types, roles, and privileges.

Engineering CAD applications

Engineering CAD is software that architects, engineers, and construction professionals rely on to create precise 2D and 3D drawings. The software enables users to:

- Draft, annotate, and design 2D geometry and 3D models with solids, surfaces, and mesh objects
- Automate tasks such as comparing drawings, counting, adding blocks, creating schedules, and more
- Customize the software with add-on applications and APIs

Users can enjoy a connected design experience across Engineering CAD from various remote devices. Remote users can also view, edit, annotate, and create drawings using the Engineering CAD mobile app or from a web browser using the Engineering CAD web app.

Virtual workstations

Dell Technologies VDI Solutions offer virtual workstation solutions that provide performance and experience levels that meet or exceed the levels of physical workstations. The solution combines leading technologies from Dell Technologies, VMware, and NVIDIA. The virtual workstations run on Dell VxRail HCI and PowerEdge servers.

The following table describes the benefits of virtual workstations:

Table 3. Virtual workstation benefits

Feature	Benefit
Centralized management	With faster provisioning of virtual workstations the timeline of projects is reduced. IT can focus on issues that are more important to the business while preventing and managing threats and vulnerabilities efficiently.
Ease of maintenance	Virtual workstations recover faster if there is an outage compared to a physical workstation. Less hardware to manage means reduced maintenance.
Ease of access	Virtual workstations are in the data center, unlike physical workstations, which are typically outside the data center network. Access to large datasets using virtual workstations from global, mobile, or third-party users requires less network traffic because only pixel streams from workstation displays are transmitted to end-point devices over the network.
Accelerated time-to-market (TTM)	Employee productivity is improved through ease of access to graphics-enhanced applications.
Flexible workplace	The ability to work from multiple devices and locations combined with the freedom from being tethered to physical workstations provides flexibility.
Secure your IP	Securing intellectual property within the perimeter of your organization provides data security. Virtual workstations reside in the data center and only the encrypted pixel stream from their displays travels outside the data center perimeter.
Improved version control	Centralized data eliminates the need to synchronize the distributed datasets. Syncing issues are reduced and everyone is working on the same up-to-date model.
Easy life cycle management	Virtual workstation solutions support vMotion, which enables migration between hardware. Ability to perform maintenance on hardware and workstations without affecting production hours.

NVIDIA nVector

NVIDIA nVector is the test tool we use for virtualized, graphics-related, end-user experience performance analysis and characterization (PAAC) testing for Dell Technologies VDI solutions. It allows for high-scale density testing with many simulated clients connecting to many desktop VMs.

NVIDIA nVector Lite is a scaled-down version of the full NVIDIA nVector test tool that is targeted towards smaller-scale testing with single client to desktop VM relationships. It can also be used to measure the end-user experience metrics of custom workloads. There are currently two workloads (Knowledge Worker and Spec View) enabled within the nVector tool.

We used Knowledge Worker as a representative typical workload for testing using NVIDIA's recommended configuration. The Knowledge Worker workload uses a standard collection of desktop application software on each VDI desktop, such as Microsoft Office and Adobe Acrobat Reader. It uses GPU-enabled launchers or endpoint devices to connect a specified user to a virtual desktop or virtual workstation.

The tool focuses on the measurement of the user's experience on the endpoint device rather than the response time at the virtual desktop or workstation.

WANem

WANem is a wide area network (WAN) emulator. This remote networking simulation software consists of a Linux-based device with customizable kernel settings that enable the configuration of numerous remote network parameters (such as bandwidth, latency, loss and jitter, and so on) and a web-based UI that simplifies user interaction. For more information, see WANem.

4

Infrastructure and test topology

This section describes the key hardware and software components that we used for testing.

Topics:

Dell VxRail hyperconverged infrastructure

VxRail HCI uses VMware vSAN technology, running on PowerEdge servers, and is further enabled by Dell management software. VxRail HCI simplifies IT operations and provides a highly scalable and agile VDI environment for running your virtual workstations. VxRail HCI is an integrated, preconfigured, turn-key system that makes the deployment and configuration of virtual workstations easier and faster. The infrastructure is sustained as a single product and comes with automated life cycle management capabilities. Upgrades and patches are nondisruptive.

The VxRail V Series is optimized for VDI workloads. It supports GPU hardware configurations that can run the graphics-intensive applications commonly used in the automotive and manufacturing industries. The recommended model in the V series is the VxRail V670F (all-flash). The VxRail V670F is a 2U model with a broad range of configuration options. This model is based on PowerEdge R750 rack servers.

GPU acceleration

NVIDIA vGPU

NVIDIA RTX Virtual Workstation (vWS) software delivers powerful virtual workstations from data centers to an endpoint device. The vWS software provides easy access to GIS, 3D Engineering CAD, and Computer-aided Engineering (CAE) applications in a VDI environment. The ease of accessing datasets enables greater collaboration among professionals in the GIS sector. From an IT standpoint, vWS provides centralized management of virtual workstations and exceptional monitoring features. The following figure shows an NVIDIA vGPU with a vWS license selected:

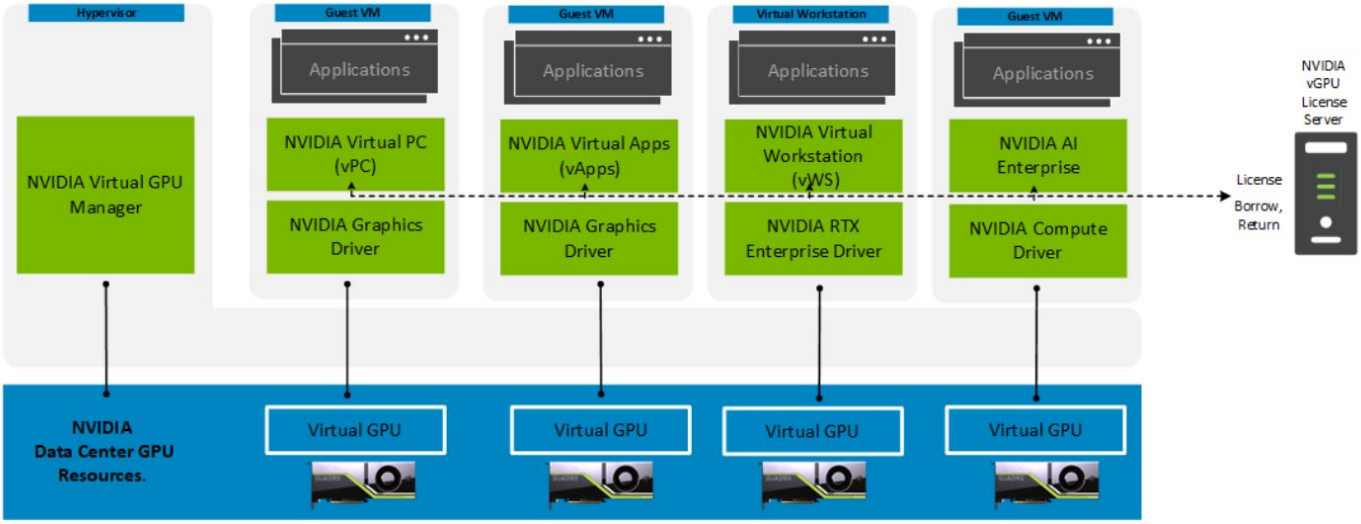

Figure 1. NVIDIA vGPU with vWS license selected

NVIDIA vGPU software creates virtual GPUs that are mapped to underlying physical GPUs. These virtual GPUs are assigned to each virtual workstation. NVIDIA vWS licenses enable the powerful graphics driver for professional 3D applications. The driver, which is installed on each virtual machine (VM) host, enables graphics commands from the VMs to pass to the GPU hardware. NVIDIA vWS comes with multiple profile size options that cater to different use cases and varying graphics requirements. You can select the correct GPU and vGPU profile (frame buffer) size for your professional graphics use cases. For GPU, vGPU profile, and density recommendations, see NVIDIA's RTX Virtual Workstation Sizing Guide.

The vWS software supports applications with the highest resolution graphics and enables professionals to work with multiple 4K, 5K, and 8K displays. The software supports up to four 5K displays or two 8K displays, enabling professionals in the GIS sector to work on the details of graphics with ease.

At the time of testing, NVIDIA T4 GPUs were used to support the vWS CAD and GIS use case. NVIDIA now recommend their latest Ampere GPU technology, specifically the NVIDIA A40 GPU, for these workloads. Details and T4 upgrade options are shown in the following sections.

NVIDIA T4 GPU

The NVIDIA T4 GPU is a flexible GPU that serves various workloads. The NVIDIA T4 Tensor Core GPU is based on the Turing architecture and comes with 2,560 CUDA cores and a 16 GB DDR6 memory. The T4 GPU operates at 70 W, providing higher energy efficiency and lower operating costs than its predecessors. It has a single-slot PCIe form factor. You can configure up to six T4 GPUs in a single PowerEdge R750 server. The T4 GPU comes with an enhanced NVIDIA NVENC encoder that can provide higher compression and better image quality with H.264 and H.265 (HEVC) video codecs. The NVIDIA T4 NVENC encoder provides up to 25 percent bit rate savings for H.265 and up to 15 percent bit rate savings for H.264.

NVIDIA A40 GPU

The NVIDIA A40 is a flexible GPU recommended for light to high-end vWS user use cases. The A40 is the recommended upgrade path for RTX8000/6000 and T4 GPUs. It is based on the Ampere architecture and comes with 10,752 CUDA cores, 336 Tensor cores, 84 RT cores, and 48 GB of DDR6 memory. The A40 operates at 300 watts and has a dual slot form factor. You

can configure up to two A40s in a single PowerEdge R750 server. The A40 comes with an enhanced NVIDIA NVENC encoder that can provide higher compression and better image quality with the H.264 and H.265 (HEVC) video codecs.

NVIDIA A16 GPU

The NVIDIA A16 is a flexible GPU recommended for entry level vWS and lightweight user use cases. It is the recommended upgrade path for T4 and M10 GPUs. The NVIDIA A16 GPU is based on Ampere architecture and comes with 4 x 1,280 CUDA cores, 4 x 40 Tensor cores, 4 x 10 RT cores, and 4 x 16 GB DDR6 memory. The A16 operates at 250 watts and has a full height, full length (FHFL) dual slot form factor. You can configure up to two A16s in a single PowerEdge R750 server. The A16 comes with an enhanced NVIDIA NVENC encoder that can provide higher compression and better image quality with the H.264 and H.265 (HEVC) video codecs.

VMware Horizon 8

VMware Horizon 8 provides the centralized management, agility, and simplicity that is required for your virtual desktop infrastructure. With Horizon 8, your workstations reside in the data center premises, which makes the provisioning, maintenance, and recovery of virtual workstations easier. Horizon 8 can provision and deliver virtual desktops and applications in a fast, flexible, and personalized manner. It leverages Instant Clones technology for ultrafast provisioning of desktops, App Volumes for real-time application delivery, and Dynamic Environment Manager for contextual policy management to provide an experience with the simplicity of non-persistent management.

When you implement VDI using Horizon 8, the intellectual property (IP) of your organization is secure inside your data center premises. A lightweight Horizon 8 client is installed on the endpoint devices that communicates with virtual workstations in the data center. The Horizon 8 client is compatible with most devices on the market. It provides flexibility for employees working from home, as well as contractors and partners working from any device anywhere, thus increasing collaboration and productivity.

For graphics-accelerated VDI workloads, Dell Technologies recommends using the Blast Extreme Display protocol, which provides an enhanced remote session experience for professional graphics applications, even in a low latency network. The Blast Extreme protocol supports the H.264 and H.265 codecs, which can encode the graphics content from a virtual workstation display. NVIDIA T4 GPUs come with an advanced NVENC encoder. This encoder can offload H.264 or H.265 encoding from server processors, providing lower latency and better performance for the professional graphics applications that are used in higher education.

Dell Technologies-optimized configurations for virtual workstations on VxRail

Dell Technologies VDI Solutions offer virtual workstation solutions for running the professional graphics applications that are used in the GIS sector. The solution configuration based on the VxRail HCI platform is recommended for running professional graphics applications such as Esri's ArcGIS Pro and Engineering CAD, as well as Revit, Inventor, and 3DS Max.

The compute layer uses VxRail V670F HCI, which is VDI-optimized and offers the highest processor speeds and graphics capability. VMware vSAN software-defined storage technology powers the storage layer. The network layer consists of Dell PowerSwitch S5248 (25 GbE ToR switches). Virtual workstations run on a VMware ESXi hypervisor. VMware Horizon 8 is the virtual desktop brokering software. NVIDIA vWS virtualizes the GPU hardware to provide virtual GPUs for each virtual workstation.

For graphics-intensive desktop deployments, Dell Technologies recommends the VDI-optimized 2U servers that support GPU hardware. We have designated common configurations as Management-optimized and Virtual Workstation. The following table describes these configurations:

Table 4. Common graphics-intensive desktop configurations

Configuration	CPU	RAM	Disk	GPU	Description
Management-optimized	1 x Intel Xeon Gold 6330 (28 core @ 2.0 GHz)	256 GB (16 x 16 GB @ 2933 MHz)	4 TB + (Capacity)	None	Offers a scalable and value-targeted configuration that meets the required compute and I/O demands

Table 4. Common graphics-intensive desktop configurations (continued)

Configuration	CPU	RAM	Disk	GPU	Description
Virtual Workstation	2 x Intel Xeon Gold 6354 (18 core @ 3.0 GHz)	128 GB-768 GB @ 3200 MHz	8 TB + (Capacity)	Up to 6 full length, single width (FLSW) GPUs(for example the T4) Up to 2 full height, full length (FHFL) double width GPUs (for example the A40)	Offers even higher performance at the tradeoff of user density. Typically used for high-end graphics workloads.

For more information, see the VMware Horizon on VxRail and vSAN Ready Nodes Design Guide.

The following figure shows an overview of the solution architecture of the virtual workstation environment on a VxRail platform:

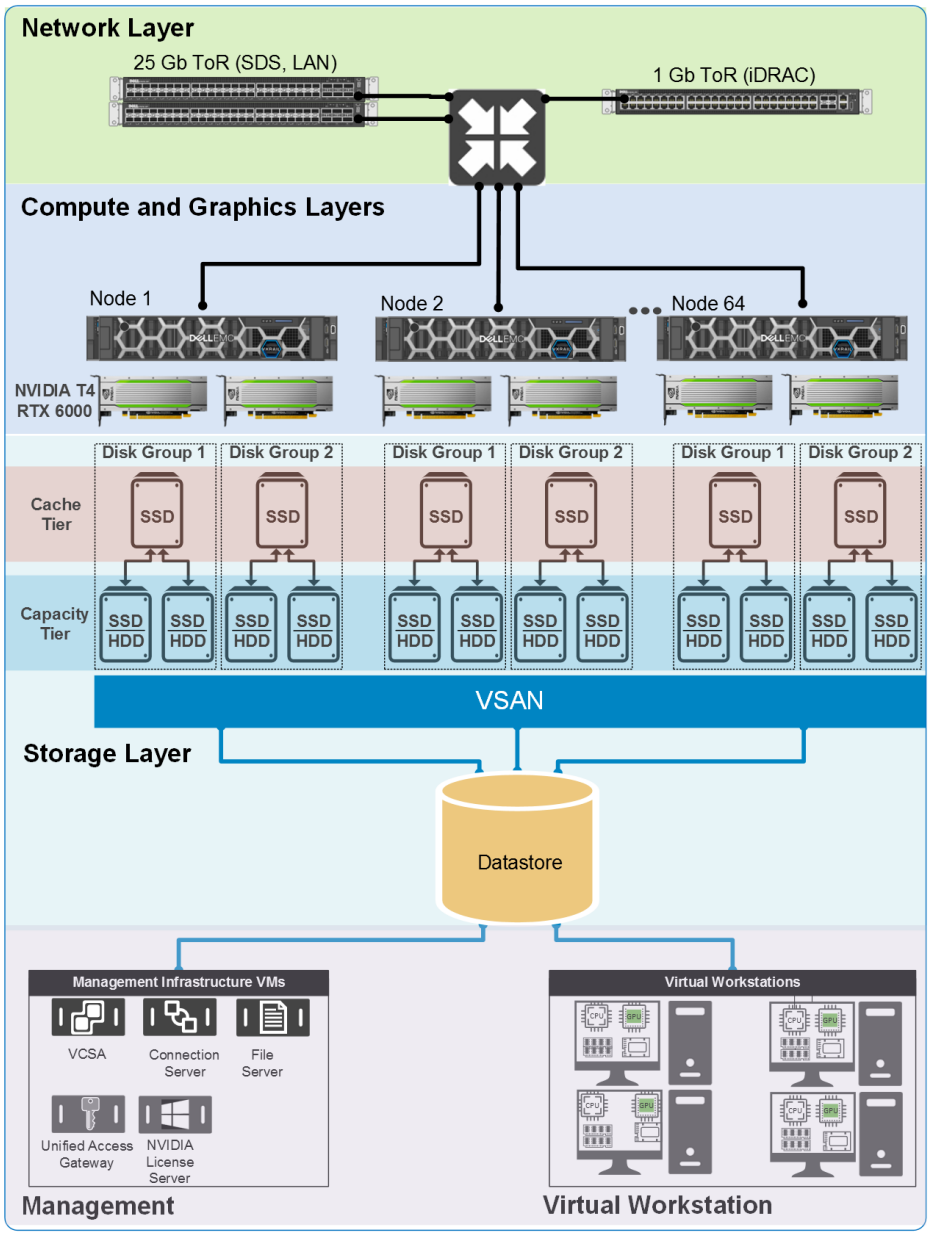

Figure 2. Virtual workstation on VxRail solution architecture

VMware vSAN software-defined storage

VMware vSAN is available in all-flash or hybrid configurations. After vSAN is enabled on a cluster, all disk devices presented to the hosts are pooled together to create a shared datastore that is accessible by all hosts in that vSAN cluster. You can then create VMs and assign storage policies to them. The storage policy dictates availability and performance. vSAN provides the following configuration options:

- All-flash configuration: Uses flash for both the cache tier and capacity tier to deliver enterprise performance and a resilient storage platform. In this configuration, the cache tier is fully dedicated to writes, allowing all reads to come directly from the capacity tier. The cache device protects the endurance of the capacity tier. All-flash configured solutions enable data reduction features to extend the capacity tier.
- Hybrid configuration: Uses flash-based devices for the cache tier and magnetic disks for the capacity tier. Hybrid configurations are ideal for clients looking for higher volume in the capacity tier. The performance of SSD and magnetic spinning disks is comparable in VDI applications if you use enough magnetic spinning disks.

The following table shows the hardware configuration that Dell Technologies recommends for the virtual workstation solution on the VxRail platform:

Table 5. Recommended hardware configuration

Component	Configuration
Chassis	V670F
CPU	Intel Xeon Gold 6354
Memory	256 GB — 1024 GB
Storage adapter	HBA 355i
Network	25 GbE
Boot device	BOSS S2 + M.2 SATA SSDs
Cache tier	400 GB WI NVMe SSD
Capacity tier	1.92 TB or 3.84 TB RI SAS or vSAS SSD
GPU	NVIDIA A40 x 2

The degree to which higher scalability is achieved depends on the typical day-to-day activities of the graphical-application users, and the concurrent nature or otherwise of their work. Several configuration attributes directly affect density scalability, such as virtual workstation parameters, GPU scheduling policy, and vGPU profile selection. It is recommended that you test and validate the appropriate configuration options to meet the needs of your users.

Virtual workstation VM profile configurations

The following table shows virtual workstation VM profile configurations:

Table 6. Virtual workstation VM profile configurations

Workload	CPU	RAM	Reserved RAM	Desktop video resolution	Operating system
nVector Knowledge Worker (1B vGPU profile)	2	4 GB	4 GB	1920 x 1080	Windows 10 Enterprise 64-bit
ESRI ArcGis (8Q vGPU profile)	8	16 GB	16 GB	1920 x 1080	Windows 10 Enterprise 64-bit
Engineering CAD (8Q vGPU profile)	6	24 GB	24 GB	1920 x 1080	Windows 10 Enterprise 64-bit

Frame buffer or vGPU profile size is a major factor to consider when you size the virtual workstations. Choosing a vGPU profile for the user depends on the type of workflow and the size of the model with which they work. Users with more advanced graphics requirements and larger datasets require a larger frame buffer. For these users, assign a larger vGPU profile that can cater to the frame buffer requirement of the application. To understand the frame buffer requirement, perform appropriate tests in the pilot phase with tools such as GPU Profiler.

You can configure a VxRail V670F with up to two NVIDIA A40 GPUs or six NVIDIA T4 GPUs, providing 96 GB of frame buffer per node. We recommend that you size the vGPU profiles based your users' workflows. In general, if an application's use of the frame buffer nears 100 percent, you should assign the next profile size to that user.

For guidance, see the NVIDIA RTX Virtual Workstation Sizing Guide, the NVIDIA Quadro Virtual Data Center Workstation Application Sizing Guide for ArcGIS Pro and Dell Virtual Workstation Solutions for Automotive and Manufacturing. For our testing we based a typical user's representative workload for graphics-intensive applications with virtual workstations on these two documents.

Network performance categories

A good internet speed is a subjective topic influenced by several factors. For more information, see What is a Good Internet Speed? and How do I Know if My Internet Connection is Good?

We used a broad range of remote broadband configurations for our testing. We labeled these configurations with four subjective broadband categories for analysis purposes, as shown in the following table.

The LAN configuration represents the campus, office or local user and was used to benchmark against the remote network categories. We added the "Very Poor" broadband category after testing to aid with contrasting and delineating end-user experience results.

Table 7. Network categories

Network category	Very Poor (E-1, DS-1 Europe)	Poor (Standard ADLS downstream)	Standard (Standard ADLS downstream)	Good (OC-3, STS-3, high-speed ADSL downstream)	LAN
Bandwidth	2.048 Mbps	6.155 Mbps	44.736 Mbps	155 Mbps	1825 Mbps
Latency	80 milliseconds (ms)	30 ms	20 ms	4 ms	0.046 ms
Loss	6%	2%	1%	<1%	0.50%
Jitter	30 ms	10 ms	5 ms	1 ms	0.046 ms

Test topology

The following figure shows the VDI solution stack based on the VxRail platform that we used for testing:

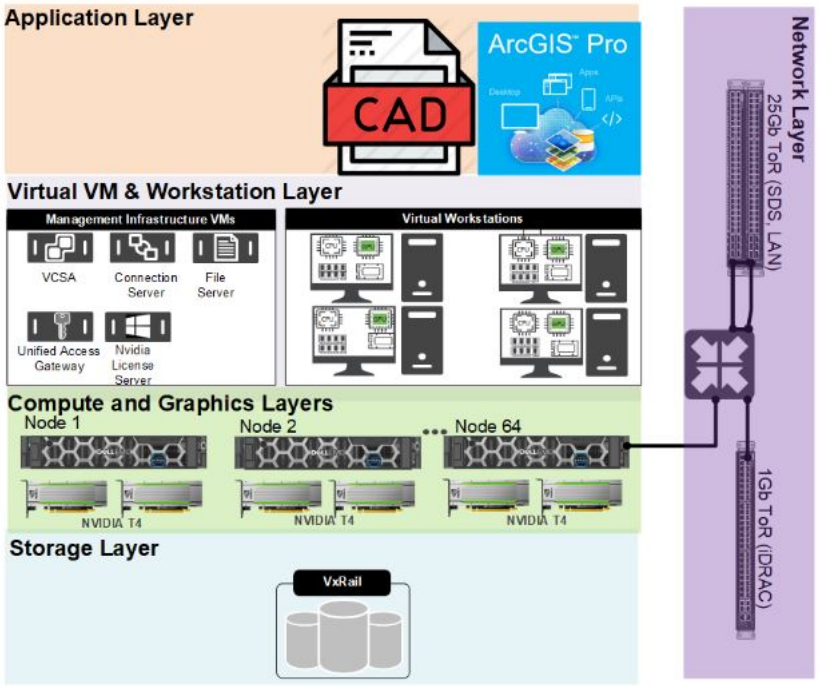

Figure 3. VDI solution stack

This solution typically resides inside a data center and the brokered VDI desktops or workstations are made available to remote users over a suitable network infrastructure.

The following figure shows the test topology we used for measuring remote user experience. The topology consisted of a brokered VDI desktop/workstation VM connected to an endpoint device. In this instance the connection was WAN-emulated. We used several different network configurations during testing.

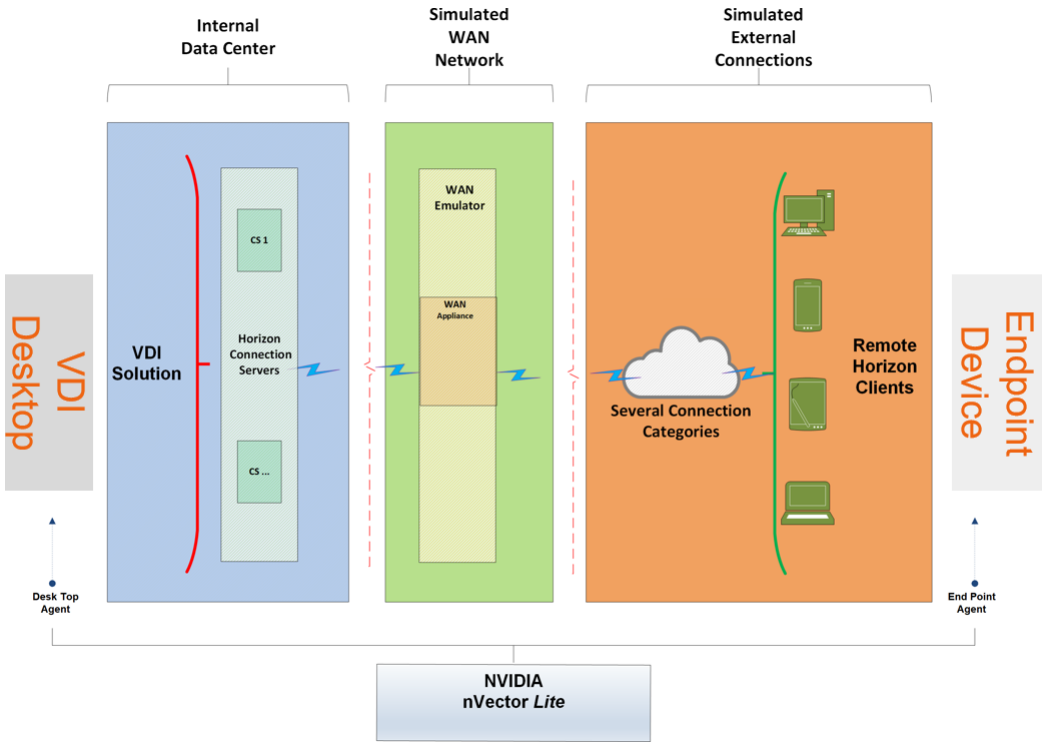

Figure 4. Test topology (remote user experience)

NOTE: A real-world VDI topology might contain load balancing and redundancy components, but these components are optional and use-case dependent. They are not required for end-user testing.

5

Test methodology and results

The following section provides our test results based on three metrics: end-user latency, frame rate, and image quality.

Topics:

Methodology

Our testing consisted of 45 test cases based on the three graphics workloads (shown in Table 1) and the five network categories (shown in Table 7). We carried out three tests for each workload and network category combination (each test had 15 combinations) and averaged the results.

NVIDIA nVector Lite performance attributes

We ran the NVIDIA nVector Lite tool to assess the end-user experience. The tool measured the following remote user experience attributes:

- End-user latency: This metric defines the level of response of a remote desktop or application. It measures the duration of any lag that an end-user experiences when interacting with a remote desktop or application. The metric is based on the graphics driver on the user's endpoint device and is measured along the whole graphics pipeline from the VDI desktop to the remote screen.
- Frame rate: This metric is a common measure of user experience and defines how smooth the experience is. It measures the rate at which frames per second (fps) are delivered on the screen of the endpoint device.
- Image quality: This metric defines the impact of remoting on image quality. It uses the Structural Similarity Index (SSIM) to compare an image that is rendered on the target desktop/workstation VM with the image that is displayed on the endpoint device. The average SSIM index of all pairs of images is computed for a single point in time for the remote VDI session. The index score is calculated once during the workload so a single value score is given for each workload.

WANem parameters

We configured the WANem emulator with the parameters for each network category as shown in Table 7 and configured the VDI desktop/workstation profiles as shown in Table 6. We applied the required network routing paths (routing using the WAN emulated network) on the appropriate infrastructure devices (such as the VDI desktop, the Horizon connection server, and the endpoint device) as shown in Figure 4.

Workloads

We selected representative user workloads for each graphics-intensive application as shown in Table 1. For all test instances, the workload we selected was chosen to represent the requirements of a medium user.

Test results

End-user latency

The end-user latency metric defines how responsive a remote desktop or application is. The following figure confirms the common assumption that the better the network connection, the lower the end-user latency. The chart shows that average end-user latency more than doubles for a remote user with a “Very Poor” network connection when compared to an office user with a local LAN connection.

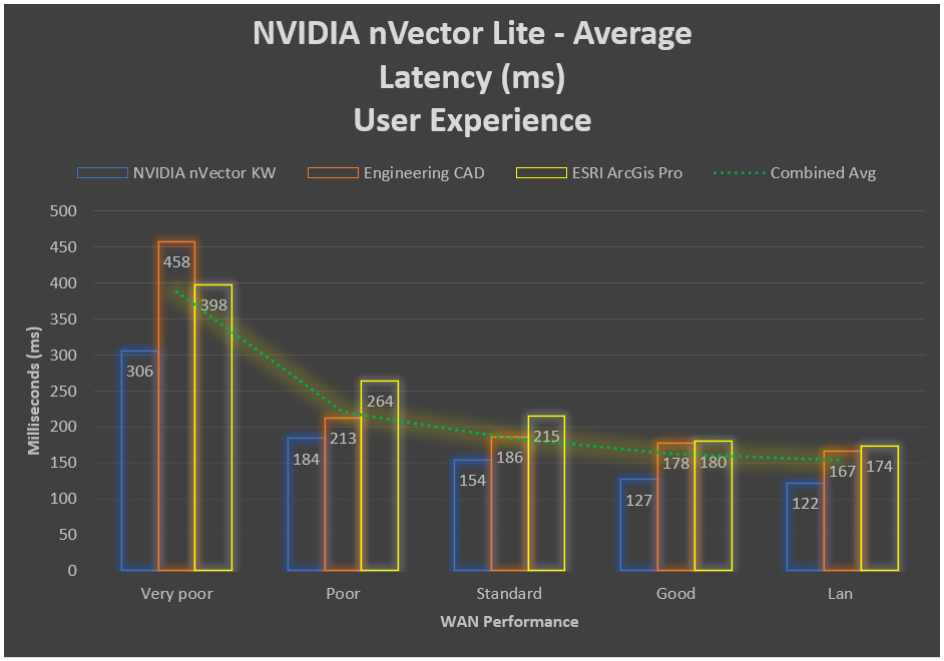

Figure 5. Average end-user latency

The following figure shows the end-user latency variability between the graphics-accelerated workloads.

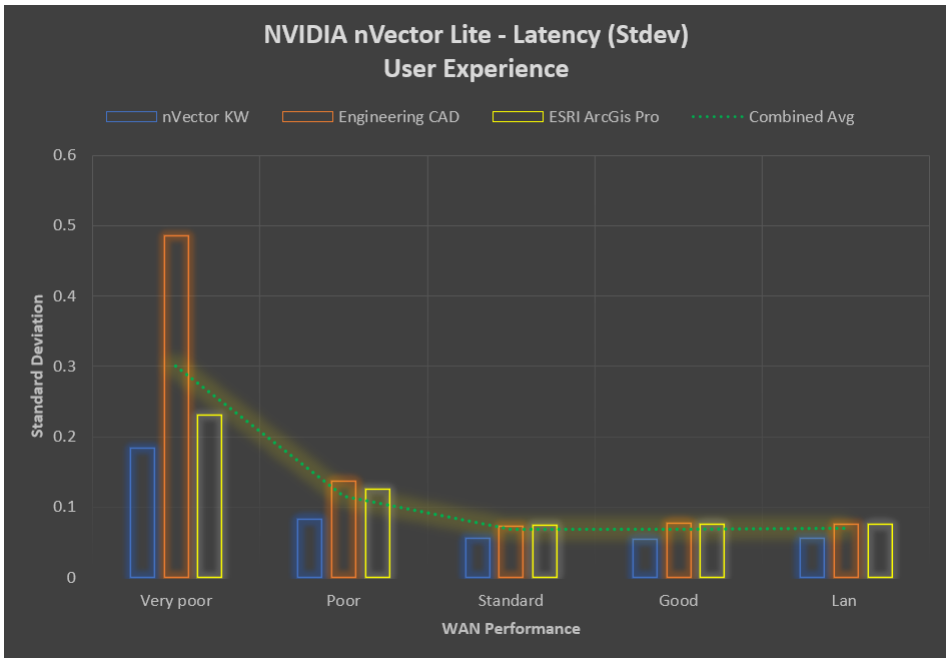

Figure 6. Stdev end-user latency

Remote users with a “Standard” to “Good” network connection can expect to see the same consistent user experience, with end-user latency distribution well below 500 ms. Latency consistency decreases as the network performance decreases, with outlying values in excess of 2 seconds for specific applications. The following figures show the user experience latency distribution for different workloads.

NOTE: Box plots help display the variability or dispersion of data. They give a good indication of how the values are distributed, which can help give a more holistic view when comparing performance data from different workloads. These plots enable the viewer to see several useful statistics on one graph, such as mean and median values, while also identifying outliers (values that fall outside the standard distribution).

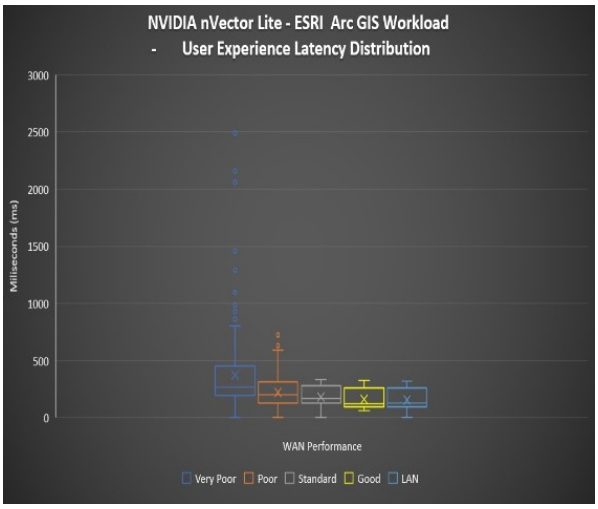

Figure 7. ArcGIS Pro workload

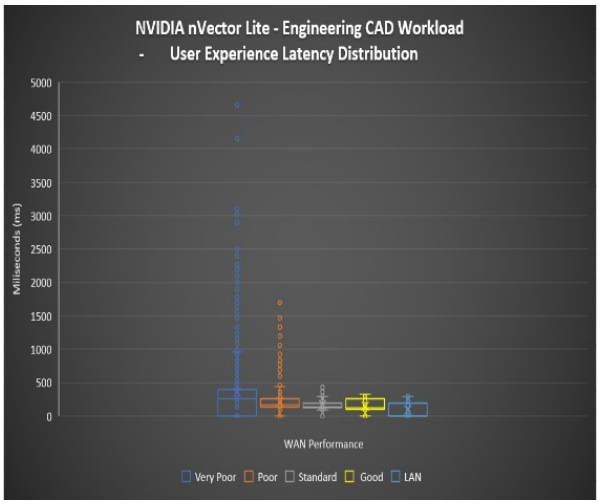

Figure 8. Engineering CAD workload

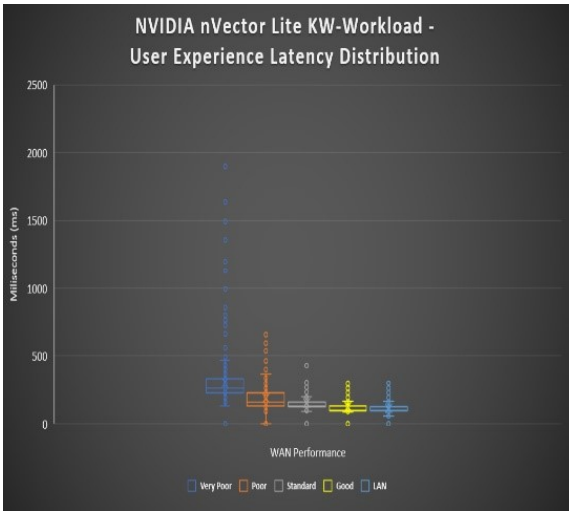

Figure 9. Knowledge Worker workload

Frame rate

The frame rate metric defines how smooth the end-user experience is. It is a measurement of "smoothness" at an endpoint device when a user is interacting with a remote VDI desktop or application. The nVector tool samples frame rates at 5-second intervals for the duration of the workload.

NOTE: By default, some graphic-accelerated applications such as ArcGIS Pro use the DirectX rendering engine to limit the frame rate to 60 frames per second. Users may choose to lower the rendering detail and anti-aliasing settings to optimize performance on virtual workstation instances. For more information, see Set display options.

The higher the number of endpoint frames per second, the smoother the end-user experience. The average fps number show a marginal increase with network performance, with minimal variability shown between the application workloads, as shown in the following two figures.

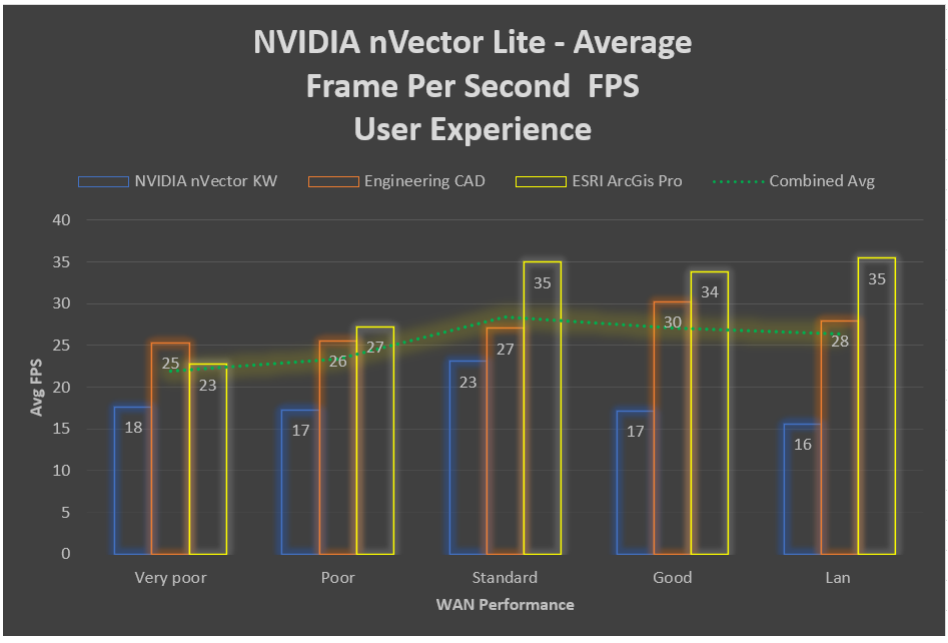

Figure 10. Average fps

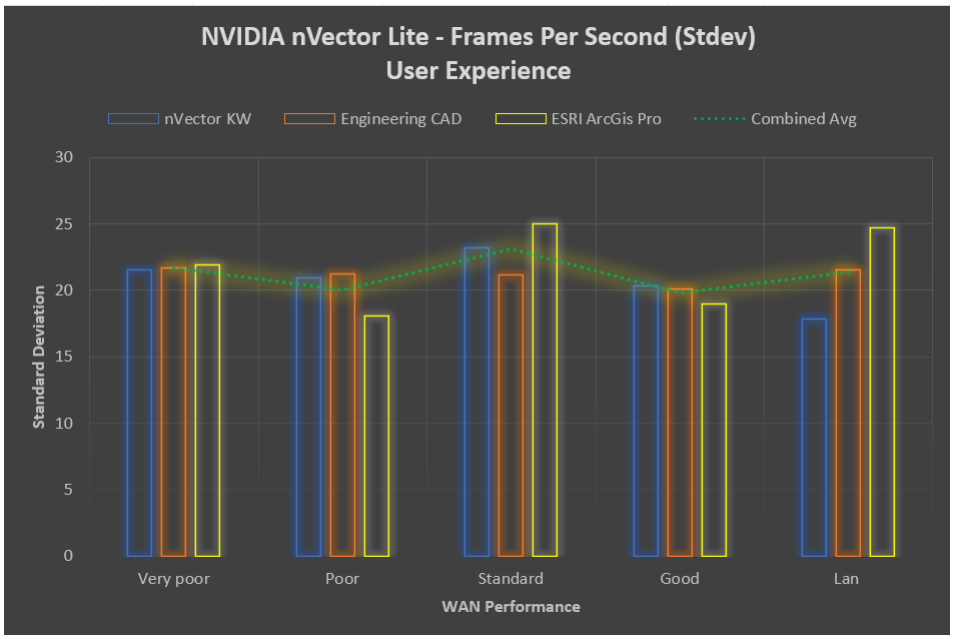

Figure 11. Stdev fps

There was no consistent trend for fps distribution among the workloads, as shown in the following figures. Consistency appeared to be application or workload dependent. Frame rate limiting can be clearly seen for some applications, as shown in the following figure:

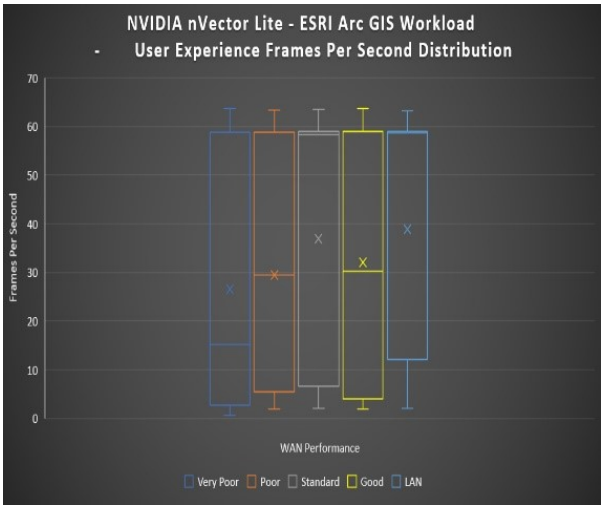

Figure 12. ArcGIS Pro workload

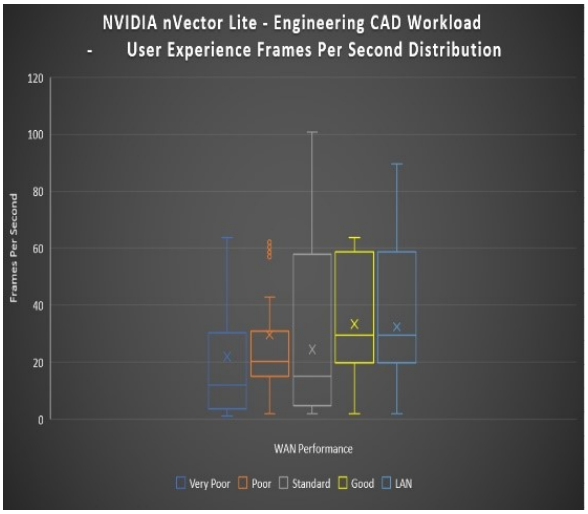

Figure 13. Engineering CAD workload

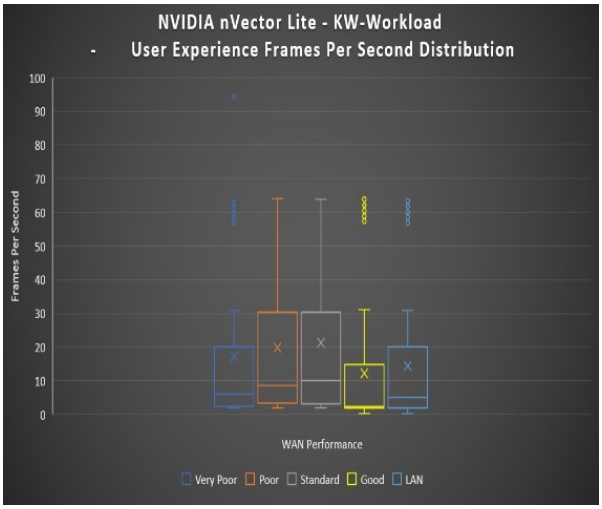

Figure 14. Knowledge Worker workload

Image quality

While image quality remained highly consistent across all network categories for some workloads, it showed a noticeable drop in the “Very Poor” network category for others. Remote image quality appears to be related to a combination of the workload and the network performance. The following figure shows the average image quality:

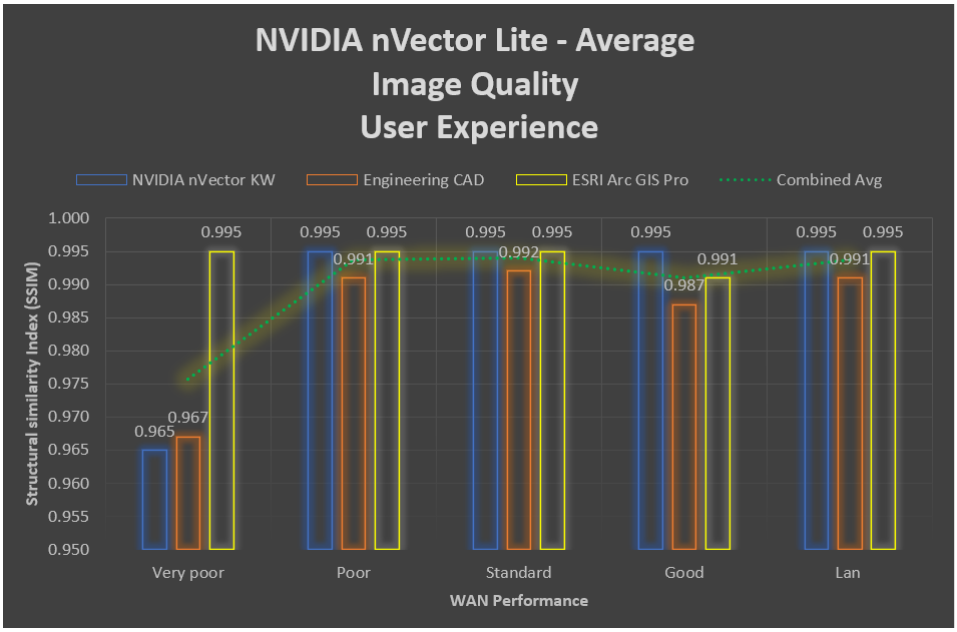

Figure 15. Average image quality

6

VDI end-user experience findings

Based on the performance testing that this study describes, when configuring a VMware Horizon virtual workstation on Dell Technologies hardware the Dell Technologies VDI Solutions team recommends that you consider the following factors:

On the physical host

- Set the Hosts BIOS System Profile to “Performance”.
- If the ambient temperature is not cool enough, consider changing the iDRAC Thermal Profile to Maximum Performance. This configuration reduces the probability of thermal throttling of the CPU and GPUs. For more information, see the Integrated Dell Remote Access Controller 9 Version 3.15.15.15 User's Guide.
- Size the virtual CPUs of the workstation VMs within the NUMA node boundaries of the physical processors. For more information, see Virtual Machine vCPU and vNUMA Rightsizing Guidelines.
- Choose and set the appropriate GPU scheduler for your workload and use case. The fixed share scheduler option provides the most consistent dedicated performance at all times. The default deployment is the “Best effort” GPU scheduler policy, which typically supports more users per server and provides better total cost of ownership (TCO) per user. Available GPU scheduler options include Best Effort (Default), Equal Share, and Fixed Share. For more information, see NVIDIA’s RTX Virtual Workstation Sizing Guide.
- Specific NVIDIA GPU scheduling guidelines are described in the NVIDIA Quadro Virtual Data Center Workstation Application Sizing Guide for Esri ArcGIS Pro. Performance results for both the “Fixed Share” scheduler and the “Best Effort” scheduler options are evaluated, with recommendations based on benchmarking data and customer best practices for Esri's ArcGIS Pro deployments.

VDI display protocol recommendations

While the choice of VDI display protocol is at the user's discretion, at the time of publication Dell Technologies recommends the VMware Blast protocol.

On the virtual workstation VM

VMware Blast display protocol

- Set the VMware Horizon Blast Max Frame Rate to 60 using the registry key or GPO. For more information, see VMware Blast Policy Settings.
- Using the VMware Horizon Performance Tracker in the Virtual Workstation VM, ensure that the protocol is Blast and the encoder name is NVIDIA NVENC H264. The recommended protocol for running graphics workloads in a VMware Horizon environment is Blast Extreme. The following figure shows the settings:

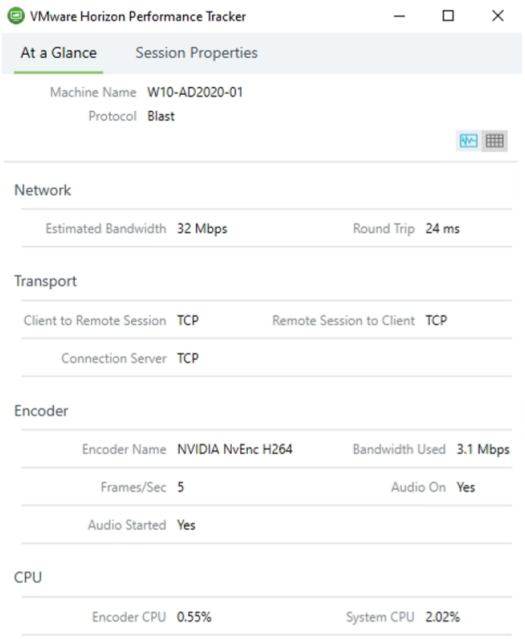

Figure 16. VMware Horizon Performance Tracker settings

- Ensure that the monitor refresh rate is set to 60 Hz in the NVIDIA Control Panel in the virtual workstation VM.

PCoIP

To ease configuration, VMware provides an active directory ADMX template file containing the related PC over IP (PCoIP) display protocol settings. See PCoIP Policy Settings.

Settings are segregated into the following categories: PCoIP general settings, PCoIP clipboard, and drag and drop settings, PCoIP bandwidth settings, PCoIP keyboard settings, PCoIP Build-to-Lossless Features. The relevant settings for this solution are:

PCoIP general settings:

- Configure PCoIP image quality levels (range 30-100, default 40): A lower value allows higher frame rates, but with a potentially lower quality display. A higher value provides higher image quality, but with potentially lower frame rates when network bandwidth is constrained. When network bandwidth is not constrained, PCoIP maintains maximum quality regardless of this value.
- Configure frame rate compared with image quality preference (range 0-100, default 50): A higher value (maximum 100) provides high image quality even if the frame rate is choppy. A lower value (minimum 0) provides a fluent experience with aggressive image quality.

PCoIP bandwidth settings:

- Maximum PCoIP Session bandwidth: This setting applies to Horizon Agent and the client. If the two endpoints have different settings, the lower value is used. Setting this value prevents the agent from attempting to transmit at a higher rate than the link capacity, which causes excessive packet loss and a poorer user experience. The default value is 900 Mbps.
- Minimum (Floor) PCoIP Session bandwidth: This setting applies to Horizon Agent and the client, but the setting only affects the endpoint on which it is configured. This setting configures the minimum expected bandwidth transmission rate for the endpoint. The default value is 0, which means that no minimum bandwidth is reserved.

VDI Client Broker

Use the latest version of the VMware Horizon Client. At the time of publication of this document, the latest version was 8.3.0.

7

Conclusion

With the changing dynamics of modern-day education, for example, the diverse user situations such as working from campus compared to working from home, it has become challenging for IT teams to provide a consistent user experience.

While the general intuition that "The better the network performance, the better the end-user experience" holds true, our test results show that a typical remote ArcGIS Pro and Engineering CAD user with a "Standard" to "Good" broadband connection can expect the same end-user experience as an on-campus or office LAN user, as measured by NVIDIA's nVector performance monitoring tool.

Virtual workstation solutions from Dell Technologies Validated Designs (DTVD) for VDI accelerated by NVIDIA GPUs and NVIDIA vWS software offer an excellent platform to run the professional graphics applications used in higher education. These virtual workstations provide an experience that is on a par with physical workstations. Our solutions simplify the provisioning and management of virtual workstations so that your IT teams can focus on innovation. Creative and technical professionals get a high-quality, interactive graphics platform that fosters collaboration and provides flexibility. The virtual workstations are secure and the organization's IP will not be compromised.

We provide virtual workstation configurations that are based on Dell Technologies Reference Architecture Guides for VDI for GIS and virtual workstation solutions for automotive and manufacturing — see the Dell VDI Info Hub. The solution configurations in these guides show the best way to integrate several Dell hardware platforms with NVIDIA GPUs and virtual workstations powered by NVIDIA vWS while ensuring predictable performance.

The DTVD for VDI team recommends the solution configuration described in this paper for running professional graphics applications. However, we recommend that you perform pilot tests to verify end-user experience and to correctly size the virtual workstation environment according to your application requirements and user needs.

8

References

Dell Technologies documentation

The following links provide additional information from Dell Technologies. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- VDI Design Guide—VMware Horizon on VxRail and vSAN Ready Nodes
- Virtual Desktop Infrastructure for Geographic Information Systems
- Virtual Workstation Solutions for Automotive and Manufacturing
- PowerEdge R750 Spec Sheet
- Dell Technologies VDI Solutions Info Hub
- TechBook—Dell VxRail System
- Planning Guide—Dell VxRail Network Planning
- Planning Guide—VMware Cloud Foundation 3.x on VxRail
- Dell PowerSwitch S-Series 25GbE switches
- Dell PowerSwitch S-Series 10GbE switches
- Dell Latitude Laptops and Two-in-Ones
- Dell OptiPlex Business Desktops and All-in-Ones
- Dell Precision Workstations
- Wyse Thin Clients
- Dell PowerProtect DD Series Appliances
- Dell PowerStore Scalable All-Flash Storage
- Dell PowerMax NVMe Storage
- Dell PowerScale Family

NVIDIA documentation

The following links provide additional information from NVIDIA:

- NVIDIA RTX Virtual Workstation Sizing Guide
- NVIDIA Quadro Virtual Data Center Workstation Application Sizing Guide for Esri ArcGIS Pro
- NVIDIA Virtual GPU Software Quick Start Guide

Esri ArcGIS Pro documentation

The following links provide additional information from Esri on Geographic Information Systems and ArcGIS Pro:

- Esri ArcGIS Pro Overview
- Esri ArcGIS Pro Getting Started Guide
- Esri ArcGIS Pro Display options guide
- Esri Delivering ArcGIS Pro Virtually in High Resolution

Engineering CAD documentation

The following links provide additional information from Autodesk, whose AutoCAD product was used by Dell Technologies in this project as the representative Engineering CAD graphics application. However, while installation of AutoCAD in a VDI environment is possible, Autodesk does not provide technical support for virtual environments.

- Engineering CAD overview

- Product documentation for Engineering CAD example
- Autodesk virtual installation guidelines

VMware documentation

The following links provide additional information from VMware:

- VMware Horizon 8 Configuration Limits
- VMware Horizon 7 Sizing and Limitations and Recommendations
- VMware Workspace ONE and VMware Horizon 7 Enterprise Edition On-premises Reference Architecture
- VMware Horizon 7 Enterprise Edition Multi-Site Reference Architecture
- VMware Horizon Resources Page
- VMware Blast Policy Settings
- VMware PCOIP Policy Settings
- VMware Horizon License FAQ
- VMware EVC and CPU Compatibility FAQ
- VMware Update to VMware's per-CPU Pricing Model
- VMware Data Protection for a VMware Horizon VDI Environment using Dell Data Protection Suite Operations Guide