



# **HIGH PERFORMANCE DESKTOPS**

Automated, centralized workstation platform for high-end users



# CONTENTS

Executive summary.....	3
Solution overview.....	3
Use case 1: Industry-specific solution architectures.....	4
Use case 2: Automated day 1 operations for high-performance desktops.....	5
Workload-focused sizing sheet.....	5
Design principles.....	6
Workload considerations.....	7
Solution components.....	8
Hardware.....	8
Software.....	9
Solution detail.....	11
Use case 1: Industry-specific solution architectures.....	11
Use case 2: Automated day 1 operations for high-performance desktops.....	13
Solution experience.....	15
Solution deployment.....	16
Use case 1: Industry-specific solution architectures.....	16
Use case 2: Automated day 1 operations for high-performance desktops.....	16
Capacity and sizing.....	16
Summary.....	17
Implementing a proof of concept.....	17
Resources and additional links.....	17



## EXECUTIVE SUMMARY

High-end users expect simple remote access to workstations with the latest and greatest in high-performance compute and graphics capabilities to handle their diverse workloads, including editing multiple streams of high-definition video, analyzing seismic data, low-latency trading, and more. IT organizations have to meet these requirements while maintaining data security and standards compliance, providing robust backup and disaster recovery strategies, and ensuring efficient use of system resources. A shift in approach from more traditional physical and bare-metal workstations to a virtualized and hyperconverged solution built on HPE's industry-leading hardware unlocks elegant solutions to these challenges while maintaining or improving on system performance and user experience. This solution delivers proven configurations for virtualized high-performance desktops in key industries, starting with media, energy and financial. Additionally, there is increasing demand to automate as many of the deployment and management processes as possible to deliver the resources end users need in an agile fashion.

The High Performance Desktops solution from HPE is built on top of our Deployment Automation solution, utilizing the Ansible platform and delivering playbooks, job templates, workflows, and the appropriate inventories and site-specific integrations necessary to automate the deployment and setup of a VDI solution with high-performance desktops. This solution automates some of the most common complicated, repetitive, and pain-point tasks in a VDI deployment with GPUs, including the installation and configuration of host and guest GPU drivers, and configuration and optimization of the host, management stack, and user operating environment, and provides GUI-based VM vending functionality built around industry-specific and customizable user profiles. It is constructed of atomic operations that can be utilized on top of the Deployment Automation solution,<sup>1</sup> which handles infrastructure deployment and configuration automation, or can be consumed as part of a separate Ansible automation environment.

**Target audience:** This document is intended for IT decision-makers and influencers such as chief information officers (CIOs), chief technology officers (CTOs), data center managers, enterprise architects, deployment/implementation engineers, and others wishing to learn more about supplementing their existing automation efforts with a more holistic approach. The reader is expected to have an understanding of the various products, components, and integration points that exist within a data center environment that can be leveraged to effectively use this solution. Working knowledge of server architecture, networking architecture, storage design, and Ansible automation is recommended.

**Document purpose:** The purpose of this document is to describe the High Performance Desktops solution from HPE and best practices for implementing an automation strategy specific to GPU-accelerated VDI. It also describes proven high-performance desktop architectures for specific industry verticals and provides sizing guidance for common user profiles in those industries. Finally, this document serves as a user guide for integrating the relevant automation with the Deployment Automation solution.

## SOLUTION OVERVIEW

This solution is comprised of a few different use cases:

- Use case 1: Proven industry-specific solution architectures
  - Media workstations, primary application focus is Avid Media Composer
  - Energy workstations, primary application focus is Schlumberger Petrel
  - Financial workstations, primary application focus is financial information applications like Bloomberg Terminal and Refinitiv Eikon
- Use case 2: Automated day 1 operations for high-performance virtual desktop deployments
  - Host and guest GPU driver installation and configuration
  - VMware ESXi™ and VMware vCenter® configuration
  - Windows 10 configuration and optimization
  - GUI-based VM vending, built around industry specific and customizable user profiles

Future releases of this solution will include:

- Additional proven industry-specific architectures for high-performance desktops in healthcare, manufacturing, and more
- Day 2 operations/lifecycle management for drivers and configurations installed as part of this solution
- Additional automated workload-specific optimizations

<sup>1</sup> [hpe.seismic.com/Link/Content/DCPo-ydLgd8UeBwsNtPb\\_YGg](https://hpe.seismic.com/Link/Content/DCPo-ydLgd8UeBwsNtPb_YGg)



This solution is not delivered as a package to the customer but instead as a consultative service to identify the specific environment and needs of the customer. Organizations may already have some tasks or areas that are automated, and the goal would be to integrate with their existing flows as much as possible while still bringing the option for a complete ground-up automation environment for those that need it—delivered by our Deployment Automation with High Performance Desktops.

The core automation functionality in Use case 2 is delivered as Ansible playbooks and some ancillary functions in PowerShell and Python. Therefore, specific versions of Ansible, Python, and PowerShell are supported, as listed in Table 3.

### Use case 1: Industry-specific solution architectures

While the migration of high-performance users from dedicated workstations and bare-metal servers to virtualized workstations is well underway, many of these users and their IT organizations are cautious with such transitions, largely because of performance concerns. However, with the latest generation of HPE ProLiant servers and storage options, performance equivalent to or surpassing dedicated workstations for these crucial workloads is achievable on a virtualized platform. Virtualized workstation solutions bring additional advantages, including the following:

- The ability to reallocate system resources as user requirements evolve, as opposed to replacing bare-metal machines that no longer meet user demands
- Seamlessly relocate users across pools of hardware for maintenance windows and more
- Single-pane-of-glass management through VMware vCenter
- Diverse choices for built-in backup and disaster recovery
- The ability to fine-tune user access and permissions and manage user access at scale simply, as opposed to point-to-point remoting solutions

Figure 1 shows a generic virtualized solution architecture to accommodate multiple user types and locations, different delivery protocols depending on workload need and industry compatibility, and different compute and storage options. The specific choices for each virtualized architecture are addressed in the Solution detail section of this paper. The logical flow of the architecture includes the following steps:

1. End users, working either remotely or on-premises, request a VM desktop.
2. The connection manager validates and accepts the request. The connection manager hands off the request to the connection broker.
3. A connection broker handles authentication and then assigns a desktop to a user.
4. Within the Windows 10 image, third-party agents request the appropriate licenses (such as the NVIDIA® vGPU license).
5. End users are connected directly to the remote VMs using the remoting protocol.

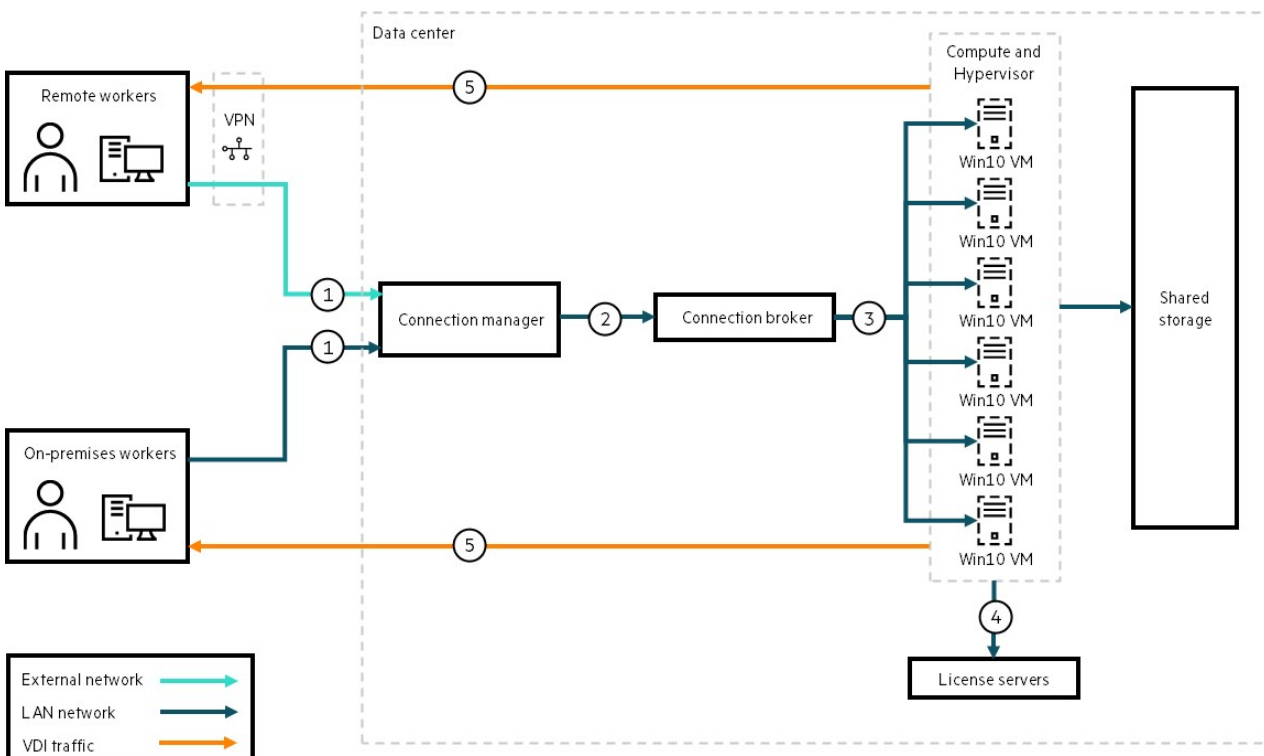


FIGURE 1. Solution architecture



### Use case 2: Automated day 1 operations for high-performance desktops

Another key component in the High Performance Desktops solution is automation, to automate host and guest GPU driver installation and configuration, and OS, hypervisor, and management optimization for graphics workloads. Additionally, there is a GUI front-end provided for all these components, as well as VM vending based on vertical-specific and customizable user profiles. The automation for this solution is built on top of the Deployment Automation solution and utilizes the Ansible AWX, NGINX, and GitLab environment deployed as part of that design.

The benefits of the workload-specific automation developed for this solution include:

- Accelerated time to value with prebuilt Ansible workflows to get up and running faster
- Industry- and workload-specific optimizations at the management, hypervisor, and OS level
- Simple integration with new or existing Ansible automation environments, including HPE Deployment Automation, with tasks delivered as atomic operations through playbooks, job templates, and workflows

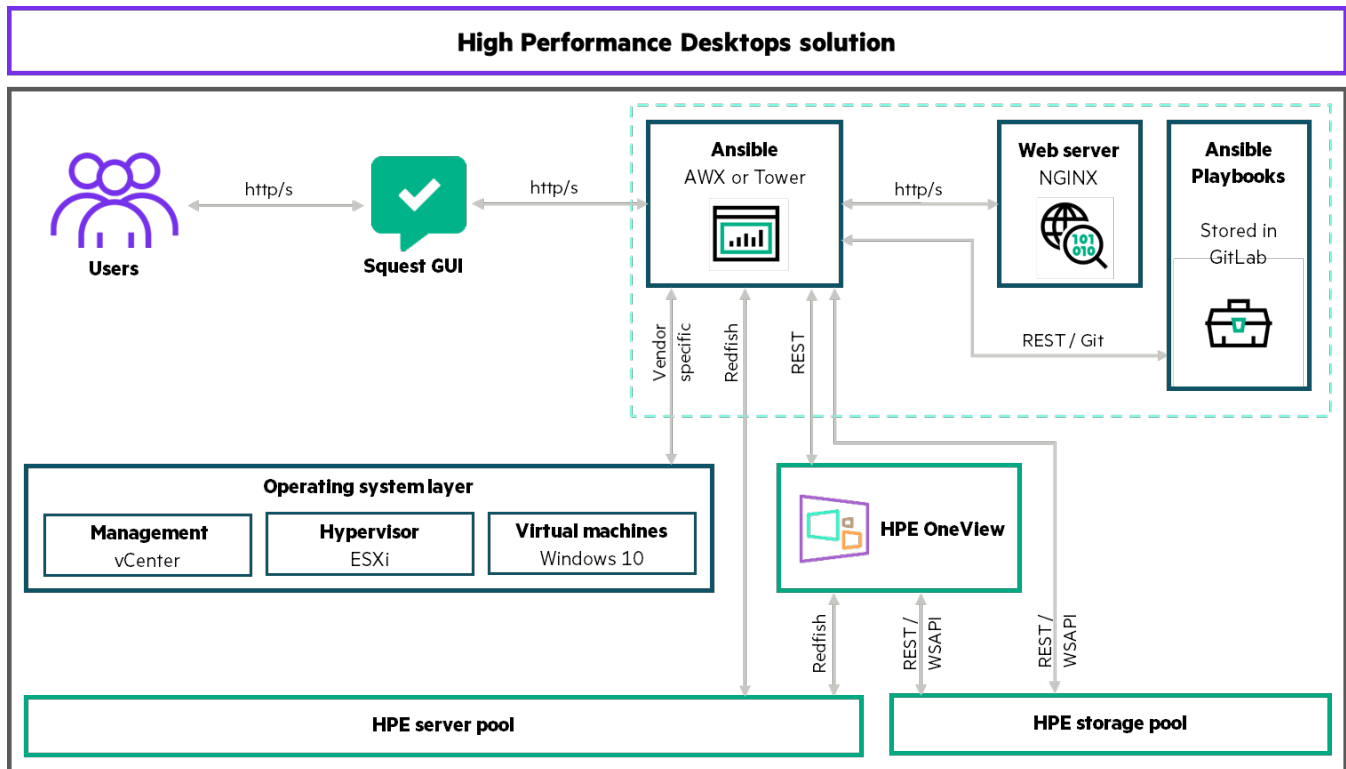


FIGURE 2. Automation architecture

### Workload-focused sizing sheet

The sizing worksheet delivered as a part of this solution makes it simple to get a rough sizing for the solution. It is not meant as a general VDI sizing guide; instead, it is targeted at users in specific industry verticals, and based on actual customer architectures, deployments, and sizing. The hardware choices, such as CPU and GPU, are limited to options that have been proven in real-world deployments, to keep it very simple and to find a sizing starting point.

Many variables influence sizing a VDI architecture, including but not limited to:

- Workload profile size, including the right-sized amount of compute, memory, and GPU
- CPU overcommit ratio
- Balance of clock speeds and core counts in CPUs, specific to the workload under consideration

The sizing worksheet includes several presized media, energy, and financial users based on common user profiles, with the option to edit the specific assumptions about these users if fine-tuning is required.



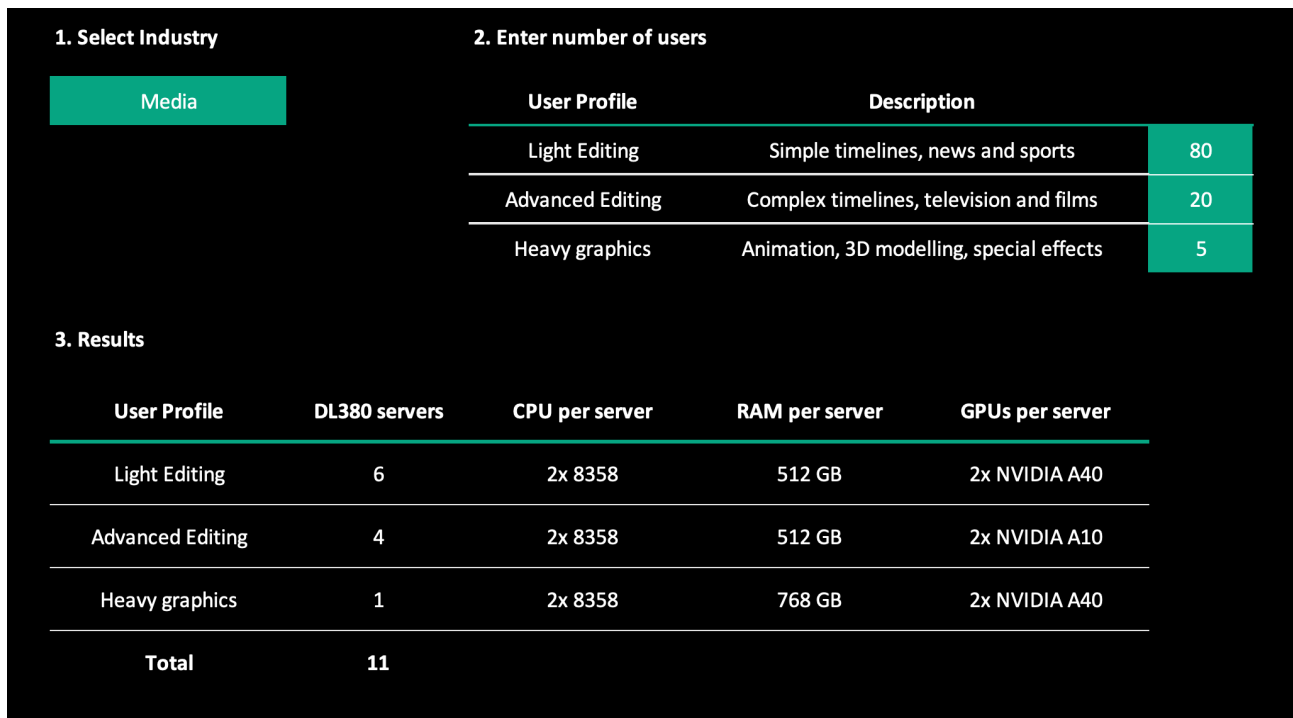


FIGURE 3. Sizing worksheet

The workload-focused sizing sheet is delivered as an Excel spreadsheet, so no deployment is required.

The first tab, labeled **High Performance Desktops**, allows the selection of industry and number of users per profile in the green boxes. Changing these values outputs different numbers and configurations of HPE ProLiant servers in section **3. Results**.

The second tab, labeled **Data**, allows editing the advanced parameters for each user profile so that they can be more accurately sized to a specific customer. These parameters include:

- vCPU count
- pCPU:vCPU ratio or overcommit ratio
- RAM
- GPU frame buffer
- CPU choice

Currently, this spreadsheet has very limited selections, as it is not intended to replace a general VDI sizer. Instead, it provides very focused guidance based on solution architectures and user profiles tested with customers and ISVs.

## DESIGN PRINCIPLES

The automation for this solution is built around Ansible. Ansible is also the foundation of our existing Deployment Automation solution, which provides an established framework for the automated day 0, 1, and 2 tasks for deploying and managing HPE servers and serves as the basis for the specific automation use cases described in this solution.

The underlying hardware for the High Performance Desktops solution is flexible, but recommendations are provided and vary by industry. The recommendations presented in this document are based on joint solutioning efforts between HPE and customers, and include various considerations applicable to the specific workloads that they were designed for, including:

- Ease of setup and management
- Flexibility for graphics options, present and future
- Workload-specific balance of core count and clock speed for CPU choice
- Workload-specific storage requirements



Additionally, it is worth repeating that a core design principle of these architectures is maximizing resource utilization. That is, traditional dedicated desktops or blade servers assigned one to one for individual users often have excess resources for their desired workload, and these resources cannot be reallocated since they are tied to the individual machine. A major goal of these architectures is to right-size the hardware components to allow excellent end-user experience today, with the ability to increase user size and decrease density per server in the future as user requirements grow.

This solution addresses architectures with storage services provided by either HPE Nimble Storage dHCI or VMware vSAN™. These are both hyperconverged infrastructures (HCIs) but use different approaches to where the storage is located and how scaling is accomplished. For increased compute and storage performance, HPE Nimble Storage dHCI is the recommended solution, but that recommendation is made with the understanding that customers will have their HCI preferences and therefore VMware vSAN is supported as well.

## Workload considerations

### Media

This solution is primarily designed around video editors using Avid Media Composer. Two traditional editing profiles can typically be defined, specifically in the context of Avid Media Composer:

**Light editing:** Editing of a relatively simple sequence timeline with minimal video streams, such as for news or sports highlights. This is usually for broadcast companies. A typical light editing user of Avid Media Composer requires 6 vCPU at a 1:1.5 physical to virtual oversubscription ratio, 32 GB RAM, and a 4 GB GPU frame buffer.

**Advanced editing:** Editing of more complex sequences, such as for episodic television or feature films with extensive video stream counts. This is usually for post-production companies. A typical advanced editing user of Avid Media Composer requires 12 vCPU at a 1:1.2 physical to virtual oversubscription ratio, 48 GB to 64 GB RAM, and an 8 GB GPU frame buffer.

Users with application requirements other than Avid Media Composer, such as Adobe Premiere Pro, Autodesk Maya, 3ds Max, and other graphically intensive applications, may require larger frame buffers sizes, up to 16 GB of GPU frame buffer.

It is important to remember that editing requirements differ from customer to customer. A customer environment may be different from the profiles described. It is recommended to use those profiles as starting points, and resize the virtual machine to match the necessary specification for a specific workload or end-user requirement.

For a much more in-depth discussion of sizing media and video-editor high-performance desktop users, see [HPE Reference Architecture for Remote Artist Workstations with Avid, Teradici, and NVIDIA on HPE ProLiant DL380 server](#).

### Energy

This solution is primarily designed around geoscientists and petroleum engineers working with the Schlumberger Petrel application. User profiles for petroleum engineers vary from customer to customer, but the following profiles provide a starting point.

**Low compute:** These users work with existing oil fields and smaller geologic formations, and a typical profile size is 8 vCPU at a minimum of 3 GHz, 64 GB to 128 GB RAM, and a 4 GB GPU frame buffer.

**Medium compute:** These users manage a larger geologic area and do some exploration, and a typical profile size is 16 vCPU at a minimum of 3 GHz, 128 GB or more RAM, and a 4 GB to 8 GB GPU frame buffer.

**High compute:** This category only represents a small percentage of users, who are focused on exploration and large-scale geologic formations. Typical profile size is at least 32 vCPU at a minimum of 3 GHz, 500 GB or more RAM, and an 8 GB or larger GPU frame buffer.

### Financial

This solution is primarily designed around financial traders and developers, both of which have similar high-frequency CPU requirements, but utilize different application stacks.

**Financial trader:** These users are characterized by requirements for extremely low-latency connections and interactions with their sessions, and high monitor counts with high total screen resolutions. Primary applications include Refinitiv Eikon and Bloomberg Terminal, along with very large and complex calculations processed in Microsoft Excel. A typical profile size is 8 vCPU with a high single-core clock speed, 64 GB to 128 GB RAM, and a 2 GB to 4 GB GPU frame buffer, with an additional requirement for low-latency storage access.

**Financial developer:** These users are developing many of the same applications and functions utilized by the financial traders, and thus have similar requirements. A typical profile size is 8 vCPU, 32 GB to 64 GB RAM, and a 2 GB to 4 GB GPU profile.





## SOLUTION COMPONENTS

### Hardware

The underlying hardware for High Performance Desktops is based on HPE ProLiant DL360 and DL380 servers, for the CPU and graphics options available, industry-standard form factor, and ease of deployment as part of different hyperconverged solutions. The automation playbooks in the Deployment Automation solution support HPE ProLiant servers that are Gen10 and later.

#### HPE ProLiant DL360 and DL380 Gen10 Plus servers

The HPE ProLiant DL360 and DL380 Gen10 Plus servers delivers the latest in security, performance, and expandability. They support the 3rd Generation Intel® Xeon® Scalable processor family supporting HPE 3200 MT/s DDR4 SmartMemory. The HPE ProLiant DL380 Gen10 server has an adaptable chassis, including new HPE modular drive bay configuration options with up to 38 SFF, up to 20 LFF, or up to 34 NVMe drive options along with support for up to three double-wide GPU options. There is a choice of OCP networking card or stand-up adapters, which offer a choice of networking bandwidth (1 Gb to 200 Gb) and fabric allowing customers to adapt and grow to changing business needs. The HPE ProLiant DL360 and DL380 Gen10 Plus servers with a complete set of HPE Pointnext Services, delivering confidence, reducing risk, and helping customers realize agility and stability.

Table 1 describes the configuration of HPE ProLiant DL380 Gen10 designed for the High Performance Desktops solution.

**TABLE 1.** Components of HPE ProLiant DL380 Gen10 server

Component	Quantity	Description
<b>CPU</b>	2	Intel® Xeon® Platinum 8358 CPU @ 2.6 GHz, 32-core or Intel Xeon Platinum 8362 CPU @ 2.8 GHz, 32-core or Intel Xeon Platinum 6342 CPU @ 2.8 GHz, 24-core
<b>Memory</b>	16	16 x HPE 64GB 2Rx4 DDR4-3200 Smart Memory Kit
<b>GPU</b>	2	NVIDIA A40 48GB or NVIDIA A10 24GB or NVIDIA Tesla T4 16GB
<b>Storage controller</b>	1 (dHCI) or 2 (vSAN)	1 x HPE Smart Array E208i-p SR Gen10 (boot drives, VMware vSAN, and HPE Nimble Storage dHCI) 1 x HPE Smart Array P816i-a SR Gen10 (only for vSAN)
<b>Disks</b>	2 (dHCI) or 8 (vSAN)	2 x HPE 240GB SATA Read Intensive SFF SSD (boot drives, vSAN and HPE Nimble Storage dHCI) 2 x HPE 375GB NVMe Write Intensive SFF SSD (vSAN cache tier, only for vSAN) 4x HPE 3.84TB SAS Read Intensive SFF SSD
<b>Network adapter</b>	2	2 x HPE Ethernet 100Gb QSFP28 Adapter

#### HPE Nimble Storage dHCI

HPE Nimble Storage dHCI is a disaggregated hyperconverged infrastructure (HCI) platform that delivers the flexibility of converged infrastructure and the simplicity of HCI. It disaggregates compute and storage and integrates hyperconverged control to give enterprises simple infrastructure management on a flexible architecture. It provides the flexibility to scale compute and storage independently for unpredictable growth and the data resiliency and performance needed for business-critical apps. This scalable solution is designed, tested, and documented to address the business requirements, workloads, and applications of HPE customers.

For this solution, HPE Nimble Storage dHCI can be deployed with either HPE Nimble Storage Flash Arrays or HPE Alletra 6000 Series for all-NVMe performance.

#### NVIDIA Quadro RTX GPUs

NVIDIA GPUs can be used for accelerated graphics, as well as a general-purpose graphics processor unit (GPGPU) for high-performance computing and analytical applications. The NVIDIA A40 GPU is a double-slot, 300 watt PCIe form-factor GPU with 10,752 CUDA cores and 48 GB GDDR6 memory. The NVIDIA A10 GPU is a single-slot, 150 watt PCIe form-factor GPU with 9,216 CUDA cores and 24 GB GDDR6 memory. The choice of which of these GPUs to use depends largely on the GPU frame buffer requirements based on the user workload.

Figure 4 shows the NVIDIA A40 GPU.








FIGURE 4. NVIDIA A40 GPU

Table 2 describes the specifications of the NVIDIA A10 and A40 GPUs tested for this solution.

TABLE 2. NVIDIA A10 and A40 GPU specifications

Component	A10 Specification	A40 Specification
CUDA cores	9,216	10,752
Additional cores	288 NVIDIA Tensor Cores, 72 NVIDIA RT Cores	336 NVIDIA Tensor Cores, 84 NVIDIA RT Cores
Memory	24 GB GDDR6 with ECC	48 GB GDDR6 with ECC
vGPU profiles	1 GB, 2 GB, 3 GB, 4 GB, 6 GB, 8 GB, 12 GB, 24 GB, (48 GB)	
System interface	x16 PCIe Gen4	
Form factor	Single-slot	Dual-slot
Max TDP	150W	300W
Thermal solution	Passive	Passive

## Software

### Use case 1: Solution architectures

#### VMware vSphere® and vSAN

VMware vSphere is [the industry-leading compute virtualization platform](#). VMware ESXi is the bare-metal hypervisor that is a part of vSphere. This solution has been tested on ESXi 7.0 U3.

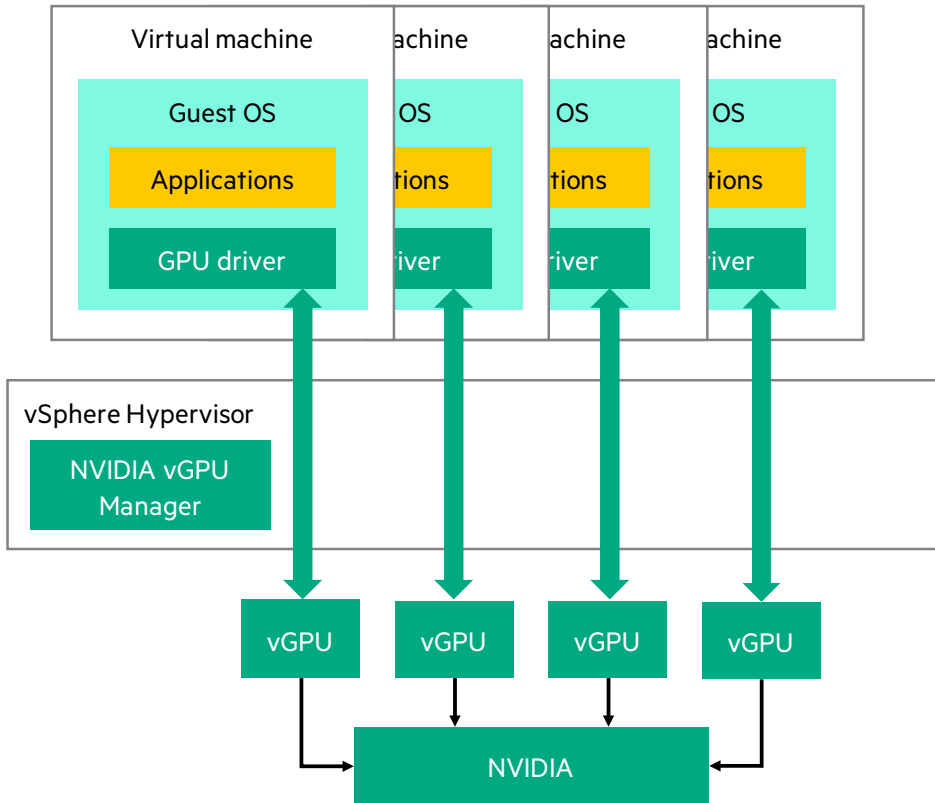
VMware vSAN is storage virtualization software that works together with vSphere to allow management of both compute and storage through a single platform. Instead of traditional, external storage, the solution becomes hyperconverged by using physical storage on the same server that is providing virtualization compute.



**NVIDIA vGPU**

NVIDIA vGPU, or virtual GPU, virtualizes data center GPUs and allows them to be shared across multiple virtual machines. This can allow for significant performance improvements when compared against CPU-only virtual machines, especially on graphics-accelerated applications, such as media workloads. Specifically for this workload, NVIDIA RTX vWS, or RTX Virtual Workstation delivers the highest performance and full capabilities of NVIDIA vGPUs.

Figure 5 shows how an NVIDIA GPU can deliver multiple vGPUs to separate virtual machines through the VMware vSphere® Hypervisor using the NVIDIA vGPU Manager.



**FIGURE 5.** NVIDIA vGPU

**Use case 2: Automation  
Ansible and Ansible AWX**

Ansible is a radically simple IT automation engine that automates cloud provisioning, configuration management, application deployment, intra-service orchestration, and many other IT needs.<sup>2</sup> The automation use cases described in this document are developed in Ansible and packaged as playbooks. Ansible AWX is an open-source upstream product that can scale, secure, and visualize the running of these playbooks, and it is recommended to use this as part of the solution.

**TABLE 3.** Software solution components

Software	Version	Description
<b>Core</b>		
Ansible	2.9.18	Open-source automation engine
Ansible AWX	17.1.0	Ansible playbook orchestrator
Python	3.7	High-level interpreted programming language
PowerShell	5.1+	On the Windows VM to interface with Ansible
VMware ESXi	7.0+	VMware® hypervisor
NVIDIA vGPU Manager	510.47.03	GPU driver for hypervisor; choose latest from NVIDIA
NVIDIA vGPU Software	511.65	GPU driver for Windows; choose latest from NVIDIA

<sup>2</sup> [ansible.com/overview/how-ansible-works](https://ansible.com/overview/how-ansible-works)



The automation built and tested for this solution functions on top of the Deployment Automation solution and includes additional utilized components such as:

- GitLab community edition
- NGINX
- Docker CE
- Docker Compose

This solution is not dependent on specific versions of those components even when they are utilized. For currently supported versions of all software solution components for Deployment Automation, see the [Deployment Automation technical white paper](#).

## SOLUTION DETAIL

### Use case 1: Industry-specific solution architectures

#### Media virtual workstations

This use case focuses on media editors for broadcast television (generally light editing), television shows, and feature films (generally more advanced editing). Media editors working with Avid Media Composer and other media applications expect uncompromising performance, particularly regarding consistent frame rates, low latency, and accurate imagery and colors. These users need the ability to work with multiple Full HD and UHD video streams simultaneously, for which Avid Media Composer demands as many high-clock CPU cores as possible. Other media applications often used by the same user group, such as Autodesk Maya, require large GPU frame buffers and graphic processing capabilities. However, storage requirements are not intense, as many of the media files are stored externally from the desktops, on Avid NEXIS.

In collaboration with Avid and our joint customers, HPE has developed an architecture to meet the needs of these video editor users built on HPE ProLiant DL380 Gen10 Plus servers with Intel Xeon Platinum 8358 processors and NVIDIA GPUs, utilizing VMware vSAN. Additionally, desktops are delivered over Teradici PCoIP Ultra, as the protocol of choice for many media companies. For more detail, see [HPE Reference Architecture for Remote Artist Workstations with Avid, Teradici, and NVIDIA on HPE ProLiant DL380 server](#).

Figure 6 shows the virtualized solution architecture as presented previously but with the specific choices that best fit the media workload filled in.

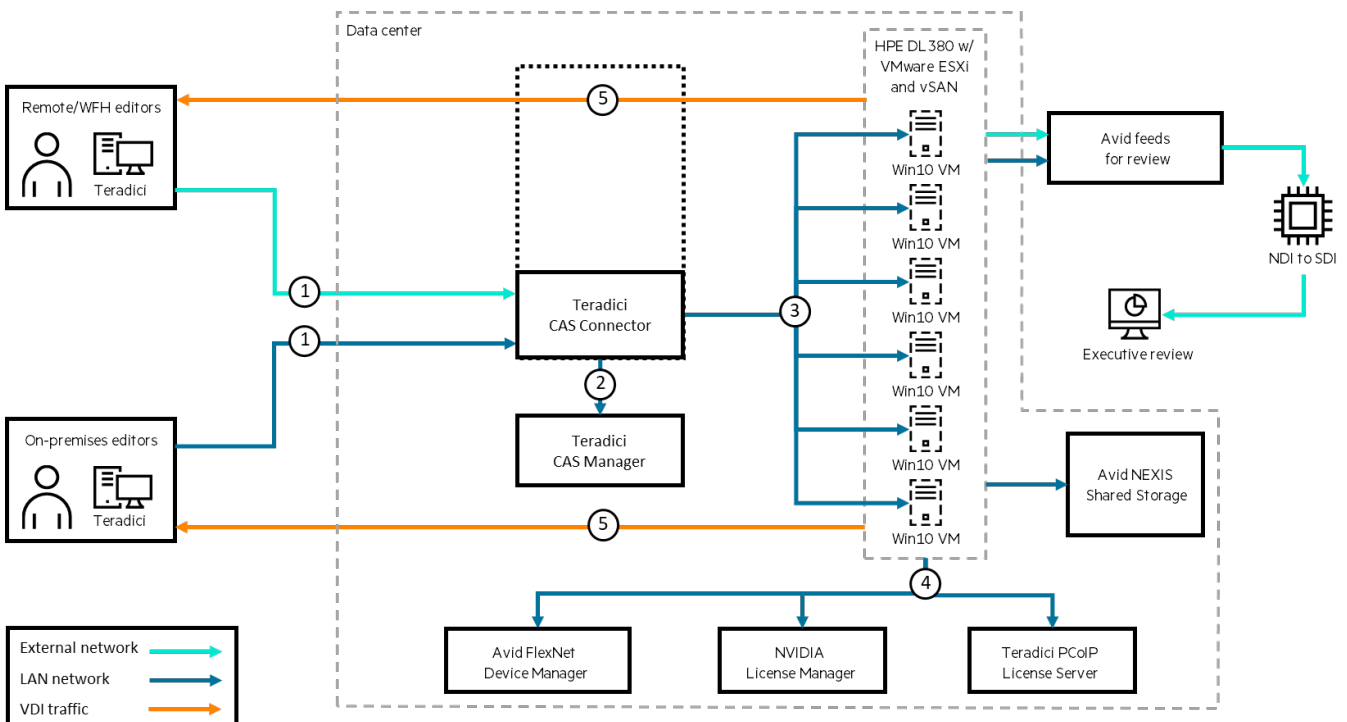


FIGURE 6. Media solution architecture

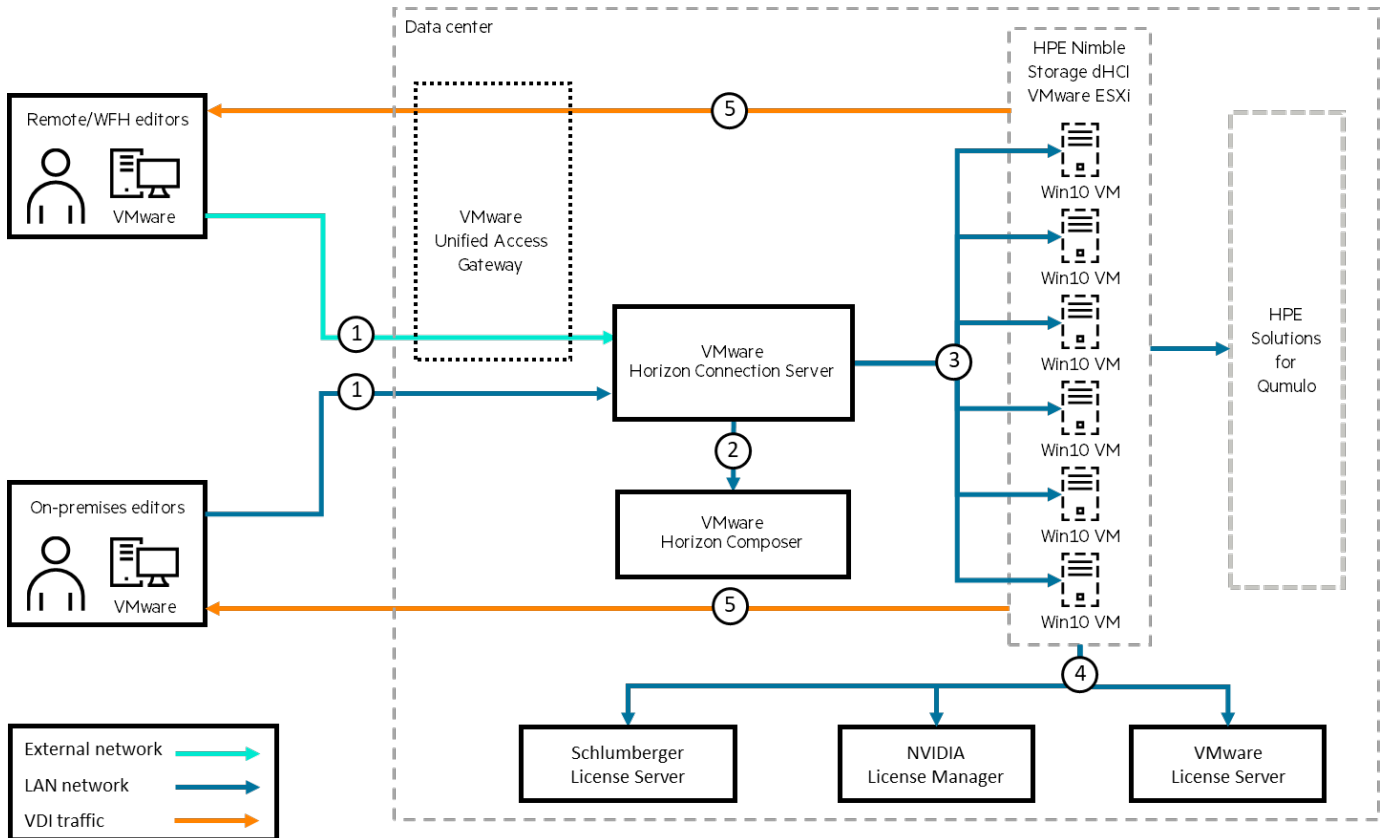


**Energy virtual workstations**

This use case focuses on petroleum engineers working with Petrel and other mapping and imaging applications. They run complex 3D modeling, simulations, and data manipulation on oil fields that require gratuitous compute resources across the board, including extremely frequency-sensitive CPU requirements, RAM footprints starting at 128 GB, and GPU profiles starting at 4 GB to 8 GB. They also work with large datasets and require a large amount of scratch disk space with fast access to the storage system hosting all these files on the back end.

In collaboration with our customers, HPE has developed an architecture to meet the needs of these petroleum engineer users, built on HPE Nimble Storage dHCI, with HPE ProLiant DL380 Gen10 Plus servers, Intel® Xeon® Gold 6342 processors, and NVIDIA GPUs.

Figure 7 shows the virtualized solution architecture as presented previously but with the specific choices that best fit the energy workload filled in.



**FIGURE 7.** Energy solution architecture



### Financial virtual workstations

This use case focuses on financial traders and developers. Financial traders working with key applications like Refinitiv Eikon, Bloomberg Terminal and Microsoft Excel expect extreme low-latency performance to quickly process transactions and complicated calculations, dependent largely on single-threaded performance. These users also need the ability to smoothly drive multiple high-resolution monitors, up to eight monitors total. To help drive this performance, they also need fast storage access.

In collaboration with our customers, HPE has developed an architecture to meet the needs of the high-end trader and developer users in the financial vertical. It is built on HPE Nimble Storage dHCI with HPE Alletra all-NVMe storage, HPE ProLiant Gen10 Plus servers with Intel Xeon Platinum 8362 processors and NVIDIA GPUs.

Figure 8 shows the virtualized solution architecture as presented previously but with the specific choices that best fit the financial workload filled in.

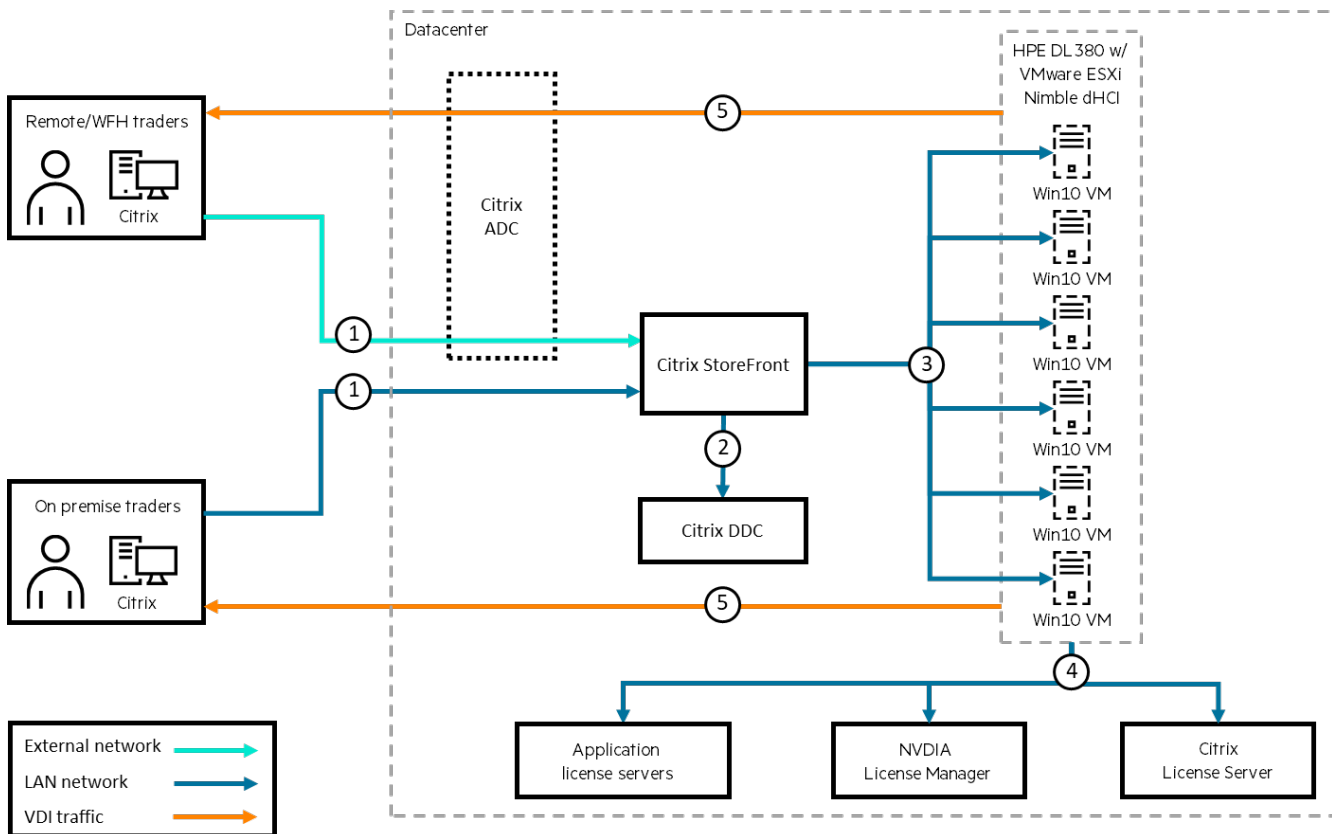


FIGURE 8. Financial solution architecture

### Use case 2: Automated day 1 operations for high-performance desktops

This solution builds on top of the Deployment Automation solution though the Ansible playbooks, job templates, and workflows included can also be utilized independently as part of a separate Ansible automation platform.

The purpose of GPU automation is to deploy the industry-specific architectures, so the appropriate combinations of hardware, hypervisor, delivery protocol, and other software for each industry are taken into consideration within the automation.



**Day 1 operations**

These playbooks automate and orchestrate the GPU driver installation and configuration within ESXi and vCenter once the ESXi hosts are deployed and accessible by IP, optimize and configure a Windows 10 template machine, and deploy customized Windows 10 VMs from that template. Ansible workflows chain the playbooks together and execute the tasks serially over multiple hosts, making it possible to run the tasks against a cluster without ever losing access to resources.

Ansible playbooks for GPU installation and configuration:

- Enable SSH and shell access to ESXi hosts in vCenter
- Get datastore of the target host
- Copy NVIDIA vGPU Manager to datastore
- Enter maintenance mode on the host
- Install NVIDIA vGPU Manager
- Reboot host
- Exit maintenance mode on the host
- Enable VMware vSphere® vMotion® for vGPU in vCenter
- Set GPU mode and assignment policy on hosts
- Disable SSH and shell access to ESXi host in vCenter

Ansible playbooks for Windows 10 template setup and optimization:

- Copy and unpack NVIDIA vGPU Software Graphics Driver
- Install NVIDIA vGPU Software Graphics Driver
- Reboot the Windows guest
- Set Windows Ultimate Performance mode
- Set NVIDIA license registry key
- Copy NVIDIA license token to guest

Ansible playbooks for deploying Windows VMs from template:

- Deploy VM from template
- Set VM-level optimizations
- Power on VM



### Solution experience

These automation workflows are comprised of atomic operations contained in Ansible playbooks that complete specific tasks. These playbooks are grouped in Ansible AWX workflow templates to run multiple tasks in sequence to complete specific use cases. Ansible AWX represents these workflows visually, and the user can follow along with the running of independent tasks in the workflow.

Workflows can be launched from the Ansible AWX portal, under Resources > Templates, where a list of configured workflows that are ready to be run is displayed, as shown in Figure 9.

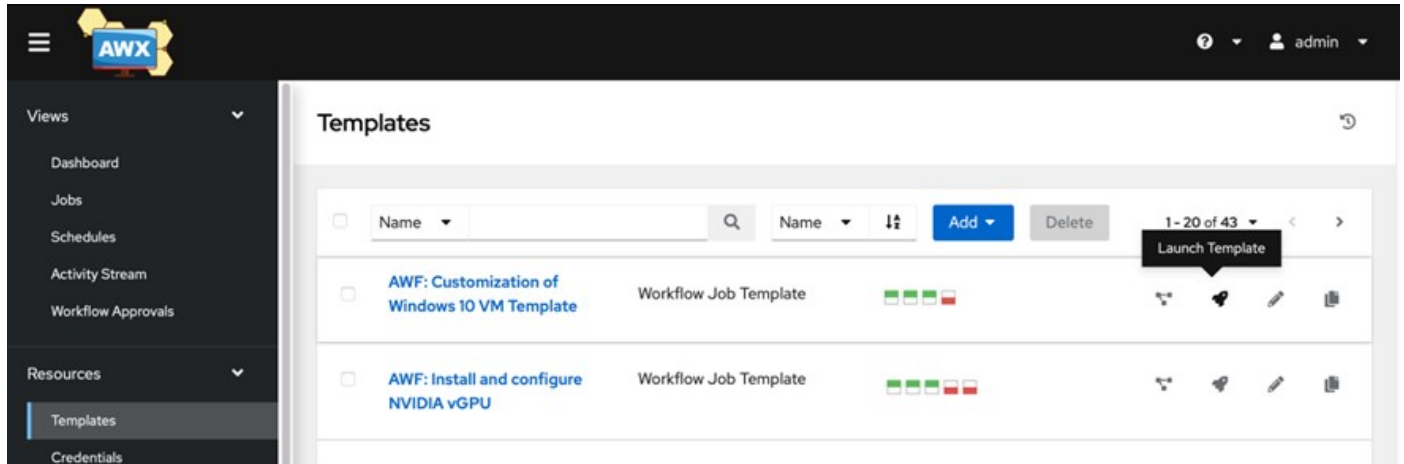


FIGURE 9. Ansible AWX workflow

Workflows may require input before launching. Once the workflow is launched, the workflow implementation screen shows all the job templates that will be run, as well as the current status so that the user can follow along with the implementation.

Additionally, workflows can be launched from the Squest GUI. The Squest interface takes input and passes it to a playbook in AWX, and implements additional functionality such as resource tracking, admin approvals, and more.

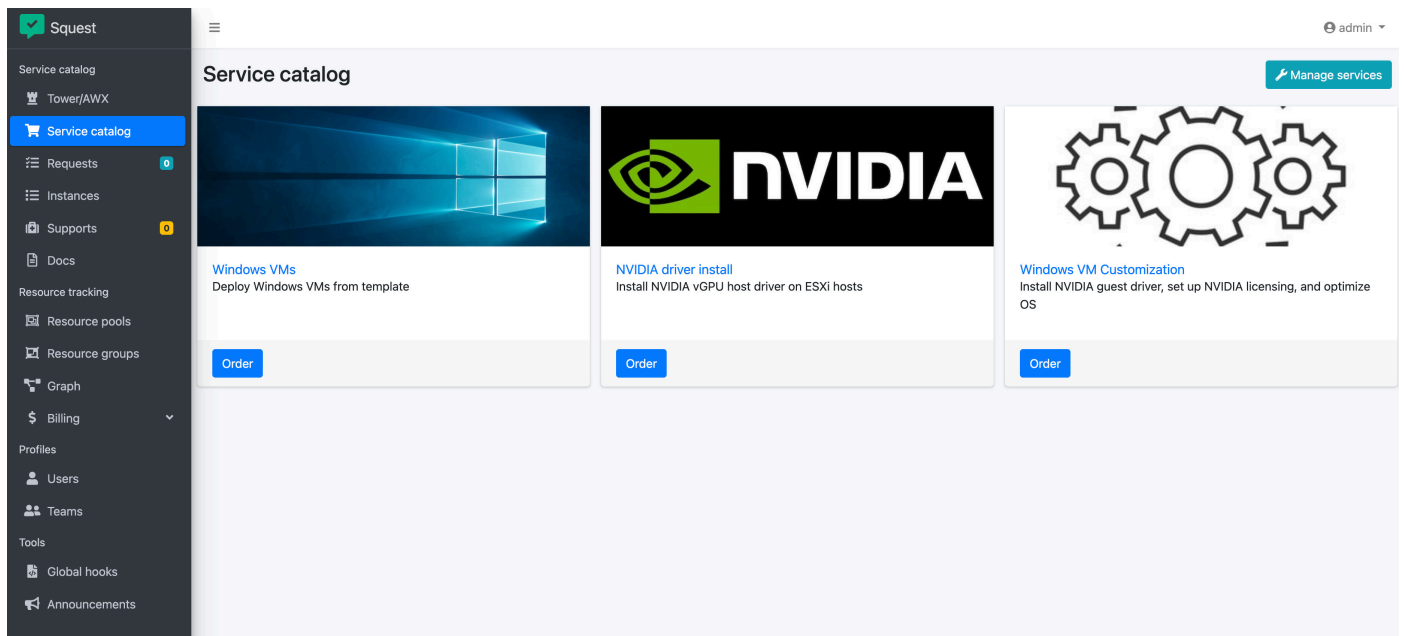


FIGURE 10. Squest service catalog user interface





## SOLUTION DEPLOYMENT

### Use case 1: Industry-specific solution architectures

This section covers the high-level steps to deploy the architectures addressed in this solution.

#### Media

1. Rack/cable/provision network
2. Deploy HPE ProLiant server with our Deployment Automation solution or another method; add to existing vCenter
3. Deploy VMware vSAN through the QuickStart Wizard
  - a. See the [vSAN Planning and Deployment guide](#)
4. For hypervisor and OS GPU automation, see Use case 2

#### Energy

1. Rack/cable/provision network
2. Initialize HPE Nimble Storage array and deploy dHCI with stack setup
  - a. See the HPE Nimble Storage dHCI documentation on [HPE InfoSight](#)
3. For hypervisor and OS GPU automation, see Use case 2

#### Financial

1. Rack/cable/provision network
2. Initialize HPE Alletra storage array and deploy dHCI with stack setup
  - a. See the HPE Nimble Storage dHCI documentation on [HPE InfoSight](#)
3. For hypervisor and OS GPU automation, see Use case 2

### Use case 2: Automated day 1 operations for high-performance desktops

#### Integrating with Deployment Automation

The automation components from this solution, including Ansible playbooks and AWX job and workflow templates, can be integrated into the Deployment Automation solution or another Ansible implementation (AWX or Tower) by creating a new project and importing the relevant code assets into a local repository in the version management software, such as GitLab.

#### Preparing Windows VM to be managed by Ansible

The Windows template virtual machine must be configured to allow Windows Remote Management access. These steps would ideally be integrated into an automated Windows deployment method or can be run manually on the Windows VM. Note that basic authentication isn't very secure, and the best practice would be to use this method to set up the VM and then disable basic authentication before putting copies of the template into production. .NET Framework 4.0 or higher is required.

1. In PowerShell, run `Set-ExecutionPolicy -Scope CurrentUser -ExecutionPolicy Unrestricted`
2. Run the `ConfigureRemotingForAnsible.ps1` script from [docs.ansible.com/ansible/latest/user\\_guide/windows\\_setup.html](https://docs.ansible.com/ansible/latest/user_guide/windows_setup.html)
3. In PowerShell, run `winrm set winrm/config/service '@{AllowUnencrypted="true"}'`
  - a. Note that this will only work if the network is set to private. If it is not, the network can be set to private with the following command:  
`Set-NetConnectionProfile -Name "NetworkName" -NetworkCategory Private`
4. Disable Windows firewall for the private network

## CAPACITY AND SIZING

Sizing for bare-metal and virtual resources to host the underlying Deployment Automation solution is discussed in the [Deployment Automation technical white paper](#).



## SUMMARY

High-performance desktop users require IT organizations to meet their demanding compute and graphical requirements while maintaining data security and standards compliance, providing robust backup and disaster recovery strategies, and ensuring efficient use of system resources. Our High Performance Desktops solution simplifies tackling these challenges for our customers by providing proven industry-specific solution architectures optimized for these workloads, automated day 1 operations for high-performance virtual desktop deployments, and a workload-focused sizing sheet.

This technical white paper describes solution testing performed in 2021 and 2022.

## Implementing a proof of concept

As a matter of best practice for all deployments, HPE recommends implementing a proof-of-concept using a test environment that matches the planned production environment as closely as possible. In this way, appropriate performance and scalability characterizations can be obtained. For help with a proof-of-concept, contact an HPE Pointnext Services representative ([hpe.com/us/en/services/consulting.html](https://hpe.com/us/en/services/consulting.html)) or your HPE partner.

## RESOURCES AND ADDITIONAL LINKS

[Deployment Automation technical white paper](#)

[HPE Reference Architecture for Remote Artist Workstations with Avid, Teradici, and NVIDIA on HPE ProLiant DL380 server](#)

[vSAN Planning and Deployment Guide](#)

[HPE InfoSight](#)

[HPE servers](#)

[HPE storage](#)

[HPE networking](#)

[Advisory and Professional Services](#)

To help us improve our documents, please provide feedback at [hpe.com/contact/feedback](https://hpe.com/contact/feedback)

## LEARN MORE AT

[hpe.com/us/en/solutions/desktop-virtualization-vdi.html](https://hpe.com/us/en/solutions/desktop-virtualization-vdi.html)

---

© Copyright 2022 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.

This document contains confidential and/or legally privileged information. It is intended for Hewlett Packard Enterprise and Channel Partner Internal Use only. If you are not an intended recipient as identified on the front cover of this document, you are strictly prohibited from reviewing, redistributing, disseminating, or in any other way using or relying on the contents of this document.

Docker is a trademark or registered trademark of Docker, Inc. in the United States and/or other countries. Intel Xeon, Intel Xeon Gold, and Intel Xeon Platinum are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Excel, Microsoft, PowerShell, and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. NVIDIA, the NVIDIA logo, CUDA, NVIDIA RTX, Quadro RTX, Tesla, and vGPU are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. VMware ESXi, VMware vCenter, VMware vSAN, VMware vSphere, VMware vSphere Hypervisor, VMware vSphere vMotion, and VMware are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All third-party marks are property of their respective owners.