



**Hewlett Packard**  
Enterprise

# **HPE Reference Configuration for Citrix VDI on HPE ProLiant Servers**

Providing desktop access from remote locations and  
facilitating collaboration among distributed teams

# Contents

Executive summary.....	3
Solution overview.....	3
Solution components.....	4
HPE ProLiant Servers.....	4
Citrix Virtual Apps and Desktops.....	5
Best practices and configuration guidance.....	7
HPE ProLiant Platform Options.....	7
VMware vSphere.....	10
NVIDIA vGPU (optional for GPU accelerated Workloads).....	11
Citrix Delivery Model.....	12
Workloads and Users.....	15
Summary.....	15
Appendix A: NVIDIA Accelerators.....	16
Appendix B: Bill of material.....	18
Resources and additional links.....	19



## Executive summary

Given the increasing need for user mobility as well as data center space constraints, businesses are continuing to virtualize desktops. Client virtualization is a key initiative for many IT organizations, driven in part by the promise of a flexible, mobile computing experience for end users and consolidated management for IT. IT organizations can consolidate the centralized infrastructure resources efficiently and can apply all the necessary software updates in a shared location. Furthermore, to significantly improve user experience, companies run their engineering, design, and business applications remotely by deploying a GPU-accelerated workload.

Organizations look to client virtualization solutions, such as Citrix® Virtual Apps and Desktops to:

- Reduce distributed maintenance and administration expenses.
- Minimize license operating expenses of their distributed environment.
- Improve security and compliance.

Too often, client virtualization deployments are over provisioned to avoid sluggish performance which results in higher capital expenditure. Therefore, organizations compromise between their solution's performance, resiliency, and cost. We intend to address all three constraints in this paper to provide an optimal TCO.

HPE ProLiant servers, including the HPE ProLiant DL380 Gen10 and the HPE ProLiant DL360 Gen10, provide a new experience in business and IT agility, superior business resilience and protection, and more control on how you consume IT. The HPE ProLiant DL380 is a secure, resilient server that delivers world-class performance and versatility. This is an optimal solution for power-user and high-resolution content which takes advantage of NVIDIA superior API and offloading capability. The HPE ProLiant DL360 Gen10 server provides highest density solution for knowledge-based users. Both HPE ProLiant DL380 and HPE ProLiant DL360 Gen10 servers offer a flexible and forward-looking design that keeps up with business needs and helps to maximize ROI.

This Reference Configuration provides a road map for an enterprise-scale architecture for implementing Citrix Virtual Apps and Desktops 7 1808 hosted desktops on HPE ProLiant Gen 10 servers.

The highlights of this solution are as follows:

- Visualize remotely hosted Windows® and Linux desktops.
- Enable a highly-responsive remote desktop experience with feature-rich experience.
- Reduce infrastructure complexity and costs through virtualization.
- Access centralized collaborative workflows from any mix of on-prem and cloud infrastructure.

**Target audience:** This document is intended for customer IT architects, managers, and administrators together with channel partner engineers, professional services personnel, and other IT professionals who plan to deploy the HPE DL3x0 servers to support Citrix Virtual Apps & Desktop (Formerly XenDesktop).

**Document purpose:** The purpose of this document is to describe an enterprise-scale design, highlighting recognizable benefits to technical audiences.

## Solution overview

Successful client virtualization environments are those that meet the goals set by IT and the business consumer at the outset. At a high level, solutions must be:

- Secure – From the compute infrastructure that hosts end-user sessions to the networks that carry the protocol by which end users receive their experience, security should be top of mind for any implementation of Citrix Workspace Suite or Citrix Cloud.
- Manageable – Migrating new end-user resources into the data center means insuring that the platforms that host those resources are highly manageable.
- Performant – The experience of the end user must be good in order for the implementation to succeed and that means the underlying platforms must be performant.



With the introduction of the HPE ProLiant DL380 Gen10 and HPE ProLiant DL360 Gen10 servers, Hewlett Packard Enterprise has released platforms that provide the security, manageability and performance demanded by client virtualization workloads. This document focuses on helping Hewlett Packard Enterprise's customers understand these platforms and the related options in the context of a client virtualization deployment. It seeks to highlight the features these platforms bring to market that help address the requirements for a successful client virtualization implementation.

## Solution components

This section describes the platforms and software that will power your client virtualization environment including an overview of the functions and benefits of Citrix's suite of products.

### HPE ProLiant Servers

Hewlett Packard Enterprise has launched three new HPE ProLiant Gen10 servers which are a strong fit for a variety of client virtualization environments. Whether your concern is providing high end graphics to remote workers in order to keep data local or you are looking for the most secure platform on which to host your users, HPE ProLiant Gen10 is the right fit.

#### HPE ProLiant DL380 Gen10

The HPE ProLiant DL380 Gen10 platform offers the ultimate flexibility for client virtualization workloads. With a choice of CPUs offering a balance between core counts and core frequencies, very large memory footprints, a broad array of graphics options and a mix of HDD, SSD and NVMe drives with up to 26 disks per host, the HPE ProLiant DL380 Gen10 is an optimal choice for all client virtualization workloads. The HPE ProLiant DL380 Gen10 supports all graphics users from those with simple video needs to workstation class users and does so with support for the NVIDIA Quadro P2000 (x4), Quadro P4000 (x4), Quadro P6000 (x3), Tesla M10 (x2), Quadro RTX6000/8000 (x2).




Figure 1. HPE ProLiant DL380 Gen10



### HPE ProLiant DL360 Gen10

The HPE ProLiant DL360 Gen10 is a performance driven dense 1U server that is ideal for a number of use cases, including VDI for knowledge based or other workloads that do not require high density GPU offload assist. The HPE ProLiant DL360 Gen10 continues support for pass through graphics options including the NVIDIA Quadro P2000 and Quadro P4000. With up to 10 internal disks including NVMe and all-flash configurations as well as a variety of processor and memory options this platform is a fit for many density optimized client virtualization environments. Figure 2 below shows the HPE ProLiant DL360 Gen10.



Figure 2. HPE ProLiant DL360 Gen10

### HPE ProLiant DL560 Gen10

The HPE ProLiant DL560 Gen10 server provides a density optimized, 4 socket server platform that yields scalability compared to two socket platforms while reducing the total number of systems that need to be managed. For large client virtualization environments with limited graphics requirements or for mixed virtualization and client virtualization workload environments looking to limit platform counts, the HPE ProLiant DL560 Gen10 is a good fit. The platform supports up to two (2) NVIDIA Quadro P2000 graphics cards in passthrough mode allowing users with a need for workstation class graphics to receive an experience alongside a mixture of primarily non-graphics enabled end users.




Figure 3. HPE ProLiant DL560 Gen10

### Citrix Virtual Apps and Desktops

Citrix is a pioneer in desktop and app virtualization technology by providing secure, compliant and powerful virtualization solutions. Citrix Virtual Apps and Desktops is an application and desktop virtualization technology that enables users to access their desktops from anywhere on any remote endpoint. Citrix Virtual Apps and Desktops make the seamless delivery of these apps and desktops possible by employing a combination of Citrix product components which includes Citrix Delivery Controller, Citrix StoreFront, Citrix Workspace App, Citrix User Profile Management, Citrix Licensing and Citrix Virtual Desktop Agent.

Active Directory is a key requirement in Citrix Deployments. AD acts as a user repository and controls access to the Citrix management components and the published applications and desktops. The server side management components are Delivery Controller, StoreFront, Studio, Director and Licensing. These components are installed on a Windows Server OS which can be placed on a VM or a bare metal machine. Citrix



Virtual Desktop Agent is a client side component and is installed on the Windows VM hosting applications or functioning as a Desktop. Citrix User Profile Management comes in packaged with Virtual Desktop Agent and runs as a service.

A brief description of each component is provided in the following subsections:

### **Citrix Delivery Controller**

Citrix Delivery Controller provisions and manages Virtual Apps and Desktops. This component is responsible for brokering of apps and desktops, managing user profiles, and managing hypervisor connections. Delivery Controllers act as the second point of authentication when launching an app or desktop and manages the user connection with help of the Virtual Desktop Agent. Each Citrix deployment requires a Microsoft SQL database server to store the environment data. Delivery controller manages the data in this SQL database

### **Citrix StoreFront**

The StoreFront provides an interface for the user to authenticate and login to Citrix environment via Active Directory. It serves as the first point of authentication and can be configured to allow non-domain users to access apps and desktops. StoreFront supports different types of authentication methods including Smart Card authentication and uses Citrix Federated Authentication Service to integrate with Active Directory Certificate services to provide another layer of security during authentication.

### **Citrix Workspace App**

Citrix Workspace app, formerly known as Citrix receiver, is a software component that is installed on the user endpoint. This software component is responsible to provide a secure virtual app or Desktop connection on the endpoint. Citrix Workspace app can be installed on Windows, Linux, Android, iOS, MacOS and embedded thin clients.

There is an HTML5 version of Workspace App which integrated with a web browser that supports HTML5, and performs much of the same core functionality as the platform-specific versions of Workspace app.

### **Citrix User Profile Management**

In a VDI environment, users who do not need high degree of personalization, are provided with non-persistent desktops. i.e., the settings changed by the user are not saved and each time the user logs off and logs back on, they get a refreshed desktop. For such users, user profile can save certain settings which can be applied during the logon. Citrix User Profile Management (UPM) takes up this responsibility of creating and managing the user profile. As compared to Active Directory roaming profiles, Citrix User Profile Management provides better roaming experience and decreases the login and logoff time for the user, it provides more control to the IT team managing the user profiles.

### **Citrix Virtual Desktop Agent**

Citrix Virtual Desktop agent (VDA) is a software component installed on the Windows Server OS, Desktop OS or Linux OS running as a VM or a bare metal machine. This component communicates with the Delivery Controller to make apps and desktops available to the end user. It is responsible to establish and manage the end users connection to the virtual resource.

A Windows Desktop OS supports a single user session. This session can either be an app or a desktop. Windows server OS can be used to create a desktop as well as application session, and can support multiple users under different session IDs. Linux OS can be launched only as desktops.

### **Citrix Cloud**

Citrix Cloud contains many cloud-based services that enables the modern digital workspace with networking and analytics. Citrix Virtual Apps and Desktops Service is an offering on Citrix Cloud and helps an organization in offloading the core components of the delivery infrastructure to the cloud, where Citrix manages the installation, maintenance and upgrades of those components while IT can focus on managing applications, desktops, policies, and user access. Citrix Management components like Delivery Controller, Studio, Director, Licensing and StoreFront can be optionally offloaded to Citrix Cloud. It is a subscription based offering.

### **Citrix Provisioning Services**

By default, Citrix Virtual Apps and Desktops has a built-in disk-based provisioning system known as Machine Creation Services (MCS). Another optional provisioning technology offered by Citrix is Citrix Provisioning (PVS) which takes a network-based streaming approach. It helps in the creation, update and management of VMs. Citrix PVS allows for the streaming of a single shared vDisk image, rather than copying or thin provisioning images to individual machines. It provides more scalability by enabling organizations to reduce the number of disk images that they manage. Even as the number of machines continues to grow, Citrix PVS provides centralized management and offers distributed processing. It also offers better versioning and rollback w.r.t. VM Image management. Citrix Provisioning can also be useful when there are constraints on storage capacity in the environment.



### Hypervisor (optional)

Citrix supports a number of feature rich hypervisors and the customer choice could be driven by licensing, available skillsets, strategic initiatives, or graphics requirements. Citrix and HPE Gen10 servers support both Microsoft Hyper-V and VMware vSphere. In this document, we will focus on VMware vSphere as a hypervisor for virtualization.

The integration of the above components allows for quick provisioning, and of easy maintenance and management of desktops and applications.

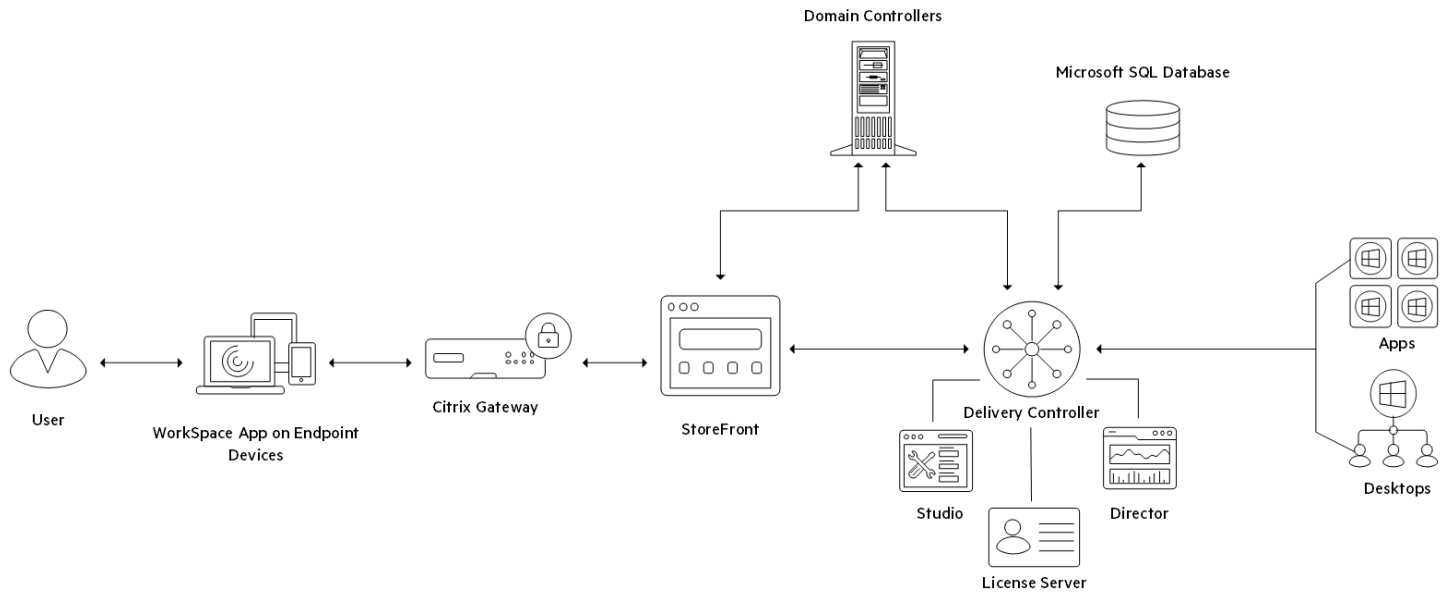


Figure 4. Integration of Citrix components

## Best practices and configuration guidance

In this section, an overview of platform option selections as well as common criteria used in system design are provided. The actual choices might be driven by individual installation factors and a PoC is recommended.

### HPE ProLiant Platform Options

Traditional conversations around the selection of two socket versus four socket systems have centered on a scalability/cost mix, risk mitigation and overall manageability. As client virtualization addresses more graphics centric use cases, the conversation must now include graphics scalability.

HPE ProLiant Gen10 CPU choices are based on a single standard development by Intel® which suggests that overall compute scalability will remain a dead heat between two and four socket platforms. Combine this with cost scalability between the two and four socket platforms is discussed in this document and it points to price/performance being on par across platforms. In selecting a platform, Hewlett Packard Enterprise sales representatives or authorized Hewlett Packard Enterprise reseller can provide a solution with optimal TCO.

As with prior generation platforms, HPE ProLiant Gen10 systems are managed by HPE iLo. This insures that each platform mentioned in this Reference Configuration is as easy to manage and own as the next. Given the common nature of the management tools, overall manageability refers to the desire to manage as few systems as is feasible for a given workload. If this is a priority within your IT environment and the remaining criteria discussed in this section are met then strong consideration should be given to the HPE ProLiant DL560 Gen10 platform as a client virtualization server resource.

Risk mitigation has traditionally involved discussions of minimizing the amount of risk placed on any one system if a hardware failure occurs. When risk aversion is a priority, deploying two two-socket servers versus one four-socket server is prudent.

Majority of systems today deploy graphic intensive applications as well as hypervisor requirements. With VMware® vSphere, there are three modes in which graphics hardware may be utilized by the end user. These modes, Virtual Shared Graphics Acceleration (vSGA), Virtual Graphics Processing Unit (vGPU) and Virtual Direct Graphics Acceleration (vDGA). These modes target different users and use cases. vSGA utilizes a



VMware graphics driver that abstracts the GPU from the end-user VM. This allows a large number of end users who require occasional acceleration capabilities to utilize onboard cards without requiring direct access. vMotion is preserved when vSGA is utilized, but only a limited set of graphics APIs are supported. vGPU allows multiple advanced graphics users to share a single graphics card using native graphics card drivers. OpenGL, OpenCL, DirectX and NVIDIA CUDA are supported and the selection depends on the solution provided. Instant Clones may be utilized to maximize manageability but vMotion is not available. vDGA makes a GPU available directly to the end-user's VM. This offers the best performance for end users that utilize the graphic intense applications and offers the broadest range of API support.

In all HPE ProLiant DL servers mentioned here, it is possible to support both vDGA users and non-graphics (or CPU only graphics users). For both vGPU (Microsoft ) and vSGA (VMware) solutions, the HPE ProLiant DL380 Gen10 is the optimal rack mount solution.

### Selecting options in client virtualization platforms

This section provides guidance for performance, security, reliability and efficiency selection criteria of VDI.

#### CPU selection

CPU selection for VDI requires an understanding of the workload or workloads you will support in your environment. Core count, core frequency and processor price point are all consideration factors. Because environments vary greatly, each organization has to look into the number and types of users, types of workloads and various methods used to utilize apps and desktops, to select the right CPU. The following general VDI guidelines are based on lab testing, customer inputs, and/or partner feedback:

- Heavy graphics users with CPU intensive applications will benefit from higher core frequencies even at the expense of lower core counts. This type of use case is not density optimized on a per server basis due to the nature of heavy graphics users, so the tradeoff is generally advisable to insure the best performance for the user.
- Knowledge workers without requirements for graphics tend to benefit from maximum core counts even at the expense of core frequency.
- Applications with limited graphics requirements that do not drive high CPU utilization should be balanced between core count and frequency with an eye toward maximizing cumulative compute capacity.

#### Storage

End-user computing implementations offer a broad array of local storage options. In a scenario involving SAN connectivity, the local host may have nothing more than a local SD or USB boot device for the hypervisor to reside on. IN contrast, a boot LUN may be presented creating a server with no local boot media (i.e. a disk free server). A more common scenario involves the implementation of software defined storage (SDS). When utilizing SDS it is important to insure you are choosing the right storage from the right vendor.

For VDI scenarios, all flash or tiered storage with at least one flash layer are the best storage options. Boot, virus scan, and user logins are all I/O (read from media) intensive and the storage subsystem has to keep up. However, runtime caching of OS and applications could be heavily write biased. The selection of a write-intensive or mixed-use drive based on the cumulative amount of I/O as well as the underlying RAID configuration of the disks must be considered in the storage architecting phase.

Whichever drive is chosen, Hewlett Packard Enterprise offers features that make it a compelling choice for end-user computing:

- HPE solid-state drives (SSDs) are backed by over 3.35 million hours of the industry's most rigorous testing.
- HPE Smart SSD WearGauge monitors and reports on the lifespan of the SSDs you have implemented.
- Full path error detection checks for data correctness between host interface to the SSD and back.
- All HPE SSDs on Gen10 platforms include Digitally Signed Firmware and best-in-class firmware security features. This prevents malicious firmware modifications that could result in data destruction, manipulation and theft as well as protection from counterfeit drives.

HPE SSDs come in a number of form factors including PCIe based, large form factor and small form factor SAS and SATA with capacities up to 6.4TB. This comprehensive portfolio is designed to meet all customer needs. The HPE SSD Selector Tool ([ssd.hpe.com](http://ssd.hpe.com)) dramatically reduces the time and complexity of selecting just the right SSD for demanding workloads

In addition to drive selection, the controller you choose for your end-user computing implementation should be a focal point when choosing your platform. HPE Smart Array controllers offer features that support a successful end-user computing experience:

- Mixed Mode – Software defined storage is frequently at the core of end-user computing implementations. Many of the SDS vendors require a controller in HBA mode. HPE Smart Array Gen10 controllers offer the ability to use both HBA and RAID mode simultaneously which means using one controller. This frees up a PCIe slot for other uses such as graphics cards.





- Better Performance – When all-flash is the design goal, disk performance is critical. HPE Gen10 controllers deliver up to 65% more performance compared to Gen9 controllers. This can mean faster boot, quicker time to recovery and improved user experience – all of which are used to define a quality end-user computing implementation.
- Less Power – End-user computing implementations move much of the user experience from the desk into the data center or server room. This can mean net new equipment in places where power matters. The Gen10 controllers use up to 45% less power than the Gen9 controllers which means lower power requirements for your implementation.
- Security – Healthcare, Government and Financial Services are big adopters of end-user computing technology and for these entities, security matters greatly. HPE Smart Array SR Secure Encryption is a FIPS 140-2 Level 1 validated enterprise class controller-based encryption solution for data-at-rest on all SAS/SATA drives and provides customers with data security to comply with regulations for sensitive data such as HIPPA or Sarbanes-Oxley.

For the majority of end-user computing scenarios involving software defined storage, including tiered or all flash disk layouts, the HPE Smart Array 400 series of array controllers will be an ideal choice.

### **Platform security**

Reliability and security have been hallmarks of the HPE ProLiant brand and one of the many reasons Hewlett Packard Enterprise continues to lead the industry in quality. HPE Gen10 servers are a continuation of that brand promise with more security features than ever before so you can continue to be confident and in control of your secure server environment within your deployment.

HPE Integrated Lights-Out (iLO) 5 and new Gen10 upgrades allow Hewlett Packard Enterprise to deliver premium security through key innovations that protect your HPE servers from attacks, detect potential invasions and allow you to recover your firmware to the last known good state.

#### **Protect**

Remove vulnerabilities that expose infrastructure firmware to malicious attacks with HPE's exclusive silicon root of trust. HPE Secure Start uniquely ensures that only HPE-signed firmware will boot by validating through HPE's silicon root of trust so you can be confident that your booted firmware is safe. New with Gen10 servers, HTTP/HTTPS boot also offer a secure and reliable replacement for PXE. Exclusively available on Gen10 servers, the new iLO Advanced Premium Security Edition license brings together a unique combination of our iLO Advanced management capabilities and new, premium security features like Commercial National Security Algorithm (CNSA) mode.

Tamper-proof updates also authenticate that firmware updates are accessible only through iLO and are digitally validated.

#### **Detect**

Run time firmware validation ensures that your firmware is checked every 24 hours to identify any potential intrusions that may occur post-boot.

#### **Recover**

Avoid lasting damage to your business by quickly restoring firmware to the factory settings or the last known authenticated safe setting in the unlikely event of a breach.

### **Server networking adapters**

With client virtualization workloads network choices come down to a need for networking performance, redundancy, security and choice, all at an affordable price. HPE's broad portfolio of standard Ethernet adapters addresses the needs of client virtualization workloads with the efficiency for today's data center workload needs. HPE 10GbE and 25GbE adapters offer secure root of trust for authenticating signed firmware.

HPE offers an array of networking options including network interface controllers supporting 1, 10, 20 and 25GbE; converged network adapters (CNA) supporting 10 and 25GbE while supporting Ethernet and Fibre Channel; and host bus adapters (HBA) supporting 8, 16 and 32Gb Fibre Channel. Ethernet connectivity selection may be driven by bandwidth, redundancy and performance. The security needs of IT add to the selection criteria. Storage connectivity drives the selection of the storage adapters. For environments where VM to VM communication needs to be as performant as possible, NICs that enable RDMA over Converged Ethernet (RoCE) can be utilized.

### **Performance**

Implementations of VDI such as those for trader workstations and high end graphics bring requirements for low network latency. For these environments adapters featuring Single Root I/O Virtualization (SR-IOV) capabilities are recommended. SR-IOV offloads the creation of Virtual Functions for access by user VMs which bypasses the hypervisor and virtual switch. This yields lower CPU utilization as well as improved VM performance from a network perspective. Implementing Network Partitioning on an HPE Converged Network Adapter (CNA) can provide QoS at the VM level and further reduce CPU utilization for increased overall performance.



### **Bandwidth and redundancy**

While not mutually exclusive, the trade-off between bandwidth and availability via redundancy is a frequent decision point for design. Many client virtualization workloads, especially those that are not dense from a per server perspective, perform well with a cumulative bandwidth of 10Gb for all functions (storage, production nets, vMotion). But for those IT departments that are looking for always available bandwidth for the various functions within client virtualization environments, HPE 25Gb Ethernet network adapters in HPE ProLiant Gen10 servers can achieve up to 750Gbps max I/O throughput for the highest performance. Frequently these environments are looking to place all traffic on a pair of redundant wires to minimize cost and complexity. For environments where discomfort with risk may drive IT decisions, choosing multiple 10Gb adapters to distribute risk across a number of physical connections may make more sense.

### **Memory**

HPE server memory combines performance, reliability and efficiency to provide the compatibility, capacity and bandwidth to productively manage your expanding workload regardless of computing and budget requirements. Properly sizing memory capacity is important to the success of Citrix Workspace Suite implementations. With HPE ProLiant Gen10 platforms, system memory capacities allow for dense implementations that prevent underutilized memory resources. Individual DIMM capacities are available up to 128GB with speeds of up to 2666 MT/s. With maximum memory bandwidth that is 66% greater than that found in HPE ProLiant Gen9 systems memory performance and sizing issues should be preventable.

Balancing the requirements between CPU and memory requires knowledge of your environment. For environments that are new to end-user computing or migrating to a new operating system contact [HPE Pointnext services](#) to discuss HPE Advanced Client Virtualization Services. For existing environments with in-depth knowledge of user and application behaviors HPE recommends utilizing that data to properly size memory. Selecting the fastest memory available per platform for end-user computing environments is recommended.

### **VMware vSphere**

Server virtualization allows multiple virtual machines (VM) to run on a single piece of hardware. When transforming to a modern digital workspace, server virtualization becomes a key aspect in reducing costs, maintaining high availability and increasing fault tolerance.

VMware vSphere employs a combination of various components to virtualize servers and transform traditional datacenters into modern virtual data centers with ease. These components include VMware vCenter and VMware ESXi. The combination of these two components allows many functions to be active within the virtual datacenter, which are VMware HA, VMware Fault tolerance, and vNetwork Distributed Switch as described in the following subsections:

#### **VMware ESXi**

ESXi is a Type-1 hypervisor that is installed on a physical server hardware. ESXi takes care of the conceptual distribution of compute and memory resources among the various virtual machines running on the actual hardware. ESXi also manages the network and storage needed by VM's.

#### **VMware vCenter Server**

vCenter Server acts as a management hub for all the services in the digital datacenter. vCenter Server is accessible via an HTML5-based interface called VMware vSphere Client. This interface can be used to group different ESXi hosts and available storage into virtual Datacenters and Clusters to manage VMs. VMware vCenter enables the IT administrators to manage and configure fault tolerance, high availability and multiple other rules and policies for VM's.

Generally in a digital workspace the management components like vCenter Server, Citrix Delivery Controllers, Citrix StoreFront Servers, SQL databases are placed in separate clusters which is configured for fault tolerance and high availability. Certain mission critical apps and desktops are also group in clusters are configured according to high data loss criticality of such organizational resources. Citrix best practice states that server OS VDAs should be placed within different clusters from desktop OS VDAs as well to ensure consistent scaling and optimal resource utilization.



Figure 5 illustrates VMware vSphere on HPE ProLiant DL Gen 10 Servers.

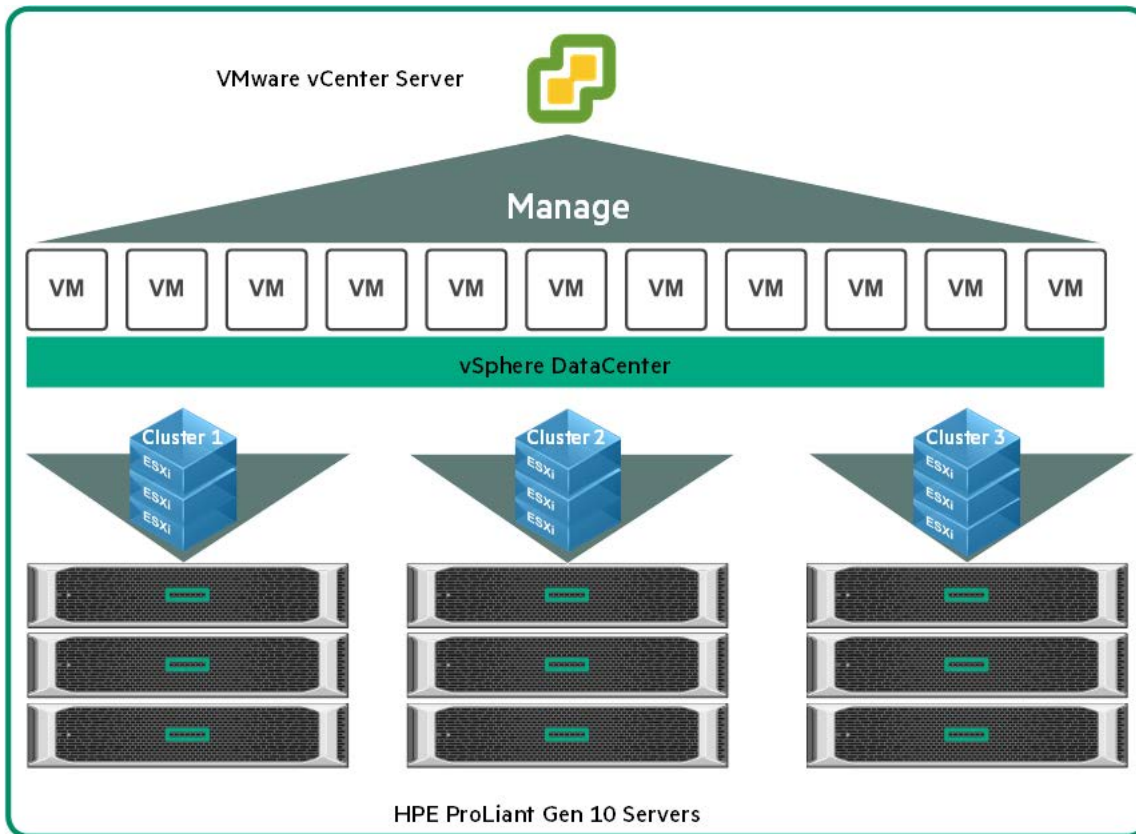


Figure 5. vSphere on HPE ProLiant DL Gen 10 Servers

### NVIDIA vGPU (optional for GPU accelerated Workloads)

Many implementations are now driven by graphics requirements. There are a broad array of options available with HPE ProLiant Gen 10 Servers. This section highlights the GPUs mentioned in this document and offers a brief explanation of their normal usage and selection criteria. For more information about each card, consult NVIDIA's website at [nvidia.com](http://nvidia.com). Further information is provided in Appendix A.

- The NVIDIA Quadro P2000 is the perfect balance of performance, compelling features, and compact form factor delivering incredible creative experience and productivity across a variety of professional 3D applications. It features a Pascal GPU with 1024 CUDA cores, large 5 GB GDDR5 on-board memory, and the power to drive up to four 5K (5120x2880 @ 60Hz) displays natively. This makes it an excellent choice for accelerating product development and content creation workflows that demand fluid interactivity with large, complex 3D models and scenes. The P2000 is available to pass through to individual virtual machines and is supported on all servers discussed in this document.
- The NVIDIA Quadro P4000 is VR ready and combines a 1792 CUDA core Pascal GPU, large 8 GB GDDR5 memory and advanced display technologies to deliver the performance and features that are required by demanding professional applications. The ability to create an expansive visual workspace of up to four 5K displays (5120x2880 @ 60Hz) with HDR color support lets you view your creations in stunning detail. The P4000 is specially designed with the performance that is necessary to drive immersive VR environments. Additionally, you can create massive digital signage solutions of up to thirty-two 4K displays per system by connecting multiple P4000s via Quadro Sync II2. The P4000 is supported on the HPE ProLiant DL360 Gen10 and HPE ProLiant DL380 Gen10 servers.
- The NVIDIA Quadro P6000 is the world's most advanced professional graphics solution ever created, combining the latest GPU, memory and display technologies that result in unprecedented performance and breakthrough capabilities. Professionals across a range of industries can now create their most complex designs, solve the most challenging visualization problems and experience their creations within the most detailed, life-like VR environments. The P6000 features 3840 CUDA cores and is supported as a pass through option on the HPE ProLiant DL380 Gen10 server.

- The NVIDIA M10 is vGPU enabled and focuses on graphics virtualization. IT needs a software-defined approach to virtualization in the data center that delivers scalability, data security, and simplified manageability. However, adoption of desktop virtualization solutions has been limited by their failure to deliver a user experience that's as good as physical devices. The NVIDIA Tesla M10 GPU accelerator works with NVIDIA GRID software to provide the industry's highest user density for virtualized desktops and applications. It supports 64 desktops per board and 128 desktops per server, giving your business the power to deliver great experiences to all of your employees at an affordable cost. The NVIDIA Tesla M10 GPU is available on the HPE ProLiant DL380 Gen10.

### **Citrix Delivery Model**

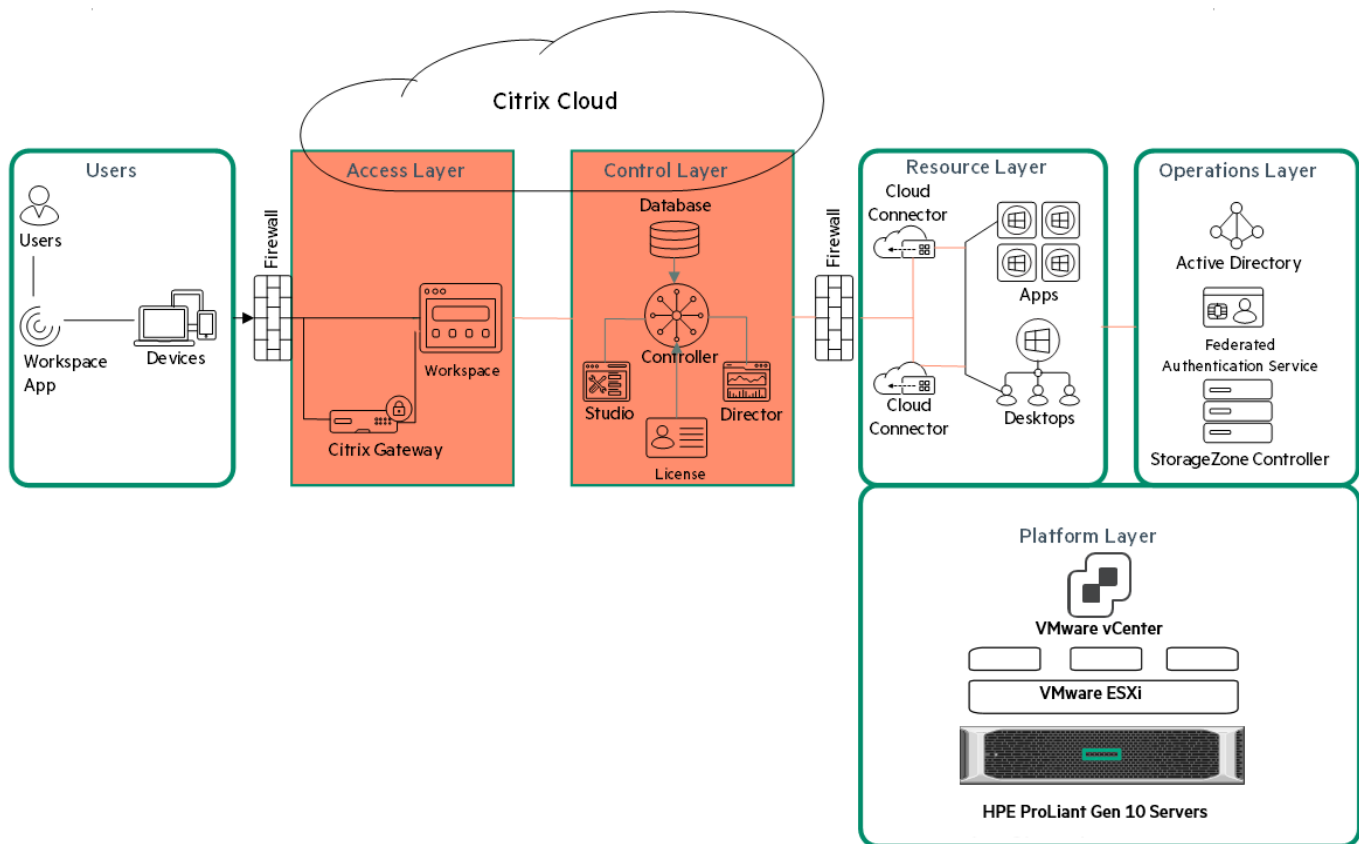
To deliver apps and desktops to the end user, Citrix offers an array of capabilities using different Citrix products called Delivery Models. These Delivery models, using the Citrix Virtual Apps and Desktops Solution, are implemented in Layers:

- User Layer - This layer encompasses users and their devices. Citrix Workspace app is part of this layer
- Access Layer - This layer defines how a user group gains access to their resources while providing secure access to the desktop/application. Citrix StoreFront is a part of this Layer
- Control Layer - This layer represents the Citrix management components and supports users accessing the resources. Citrix Delivery Controllers are a part of this layer and communicate with a SQL server to store and manage the Citrix site data. Active Directory integration and networking is also a part of this layer
- Resource Layer - This layer represents the apps and desktops that can be accessed by the user. The Citrix Virtual Delivery Agent is a component of this layer
- Platform Layer - This layer is a representation of compute, storage, network fabric and the hypervisor solution used for virtualization.



The 2 Citrix Delivery Models defined in this solution use HPE ProLiant servers to host the Resource Layer.

- Citrix Virtual Apps and Desktops Service – In this Delivery model, the Control Layer is deployed on Citrix Cloud. Citrix Virtual Apps and Desktops Service is a cloud based offering which manages the Delivery controllers, SQL database Citrix Studio, Citrix Director and optionally the Citrix StoreFront and Gateway. The resource and platform Layers are managed within the enterprise. This models provides the security and reliability of Citrix Cloud while retaining the sensitive data on premise. The placement of different components in this model is as follows
  - Citrix Delivery Controller, Citrix Studio, Citrix Director, Citrix License Server are placed on the Citrix Cloud. Typically a subscription to Citrix Cloud connection is required with access to “Citrix Virtual Apps and Desktops Service”. High Availability and maintenance of these is managed by Citrix. Local IT teams log into Citrix cloud to manage the resources layer by via Catalogs and Delivery Groups.
  - Citrix Cloud Connectors, Application Servers and Virtual Desktops are placed on premise and the hardware required for these components is managed by the organizations local IT team.
  - Citrix Gateway and Citrix StoreFront can optionally be customer managed and placed on the hardware managed by the enterprise to meet the business needs.



**Figure 6.** Citrix Virtual Apps and Desktops and Service Delivery Model



- Citrix Virtual Apps and Desktops – All the Citrix components, including the Control and Access Layer, are deployed within the enterprise according to this Delivery Model. The control layer acts as management interface to deploy, manage, virtualize and deliver secure Applications and Desktops to the end user. Like the previous Delivery Model, the Access Layer controls the subscription and seamless access of Apps and Desktops to the users via multiple authentication methods. Figure \_\_ depicts different layers of Citrix Virtual environment.

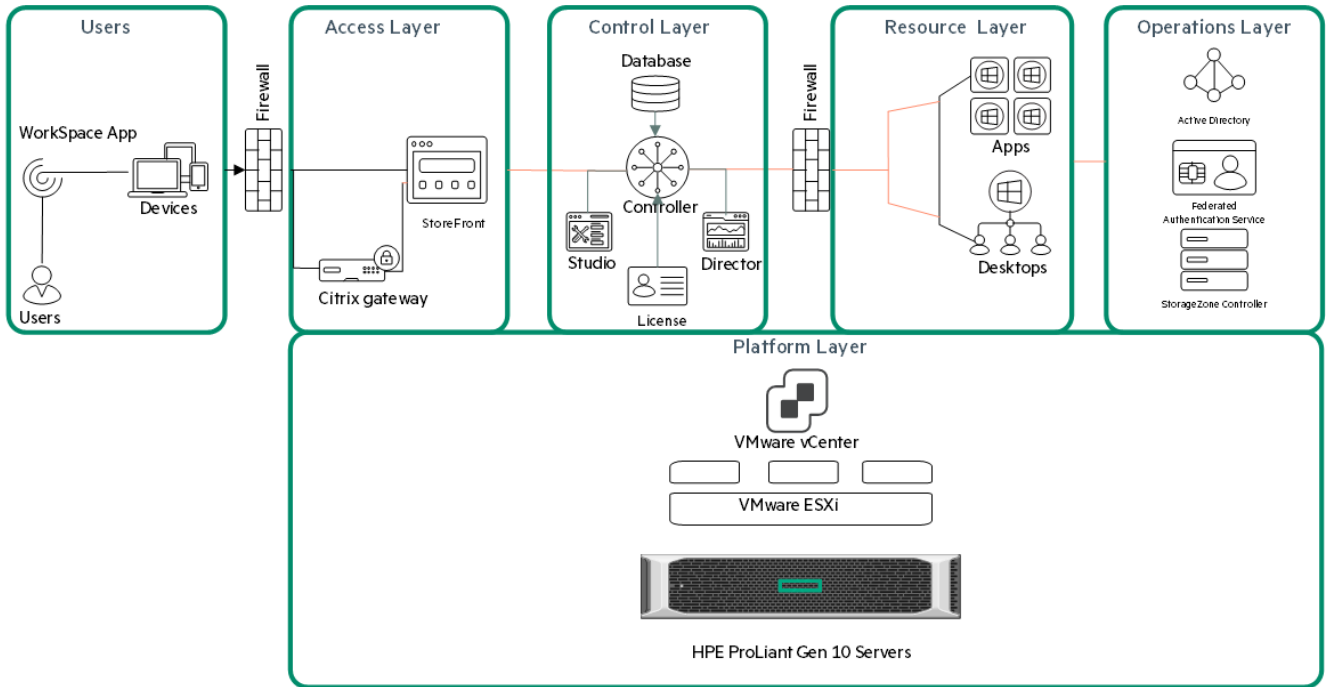


Figure 7. Layered Citrix Environment on HPE ProLiant DL Gen 10 servers



## Workloads and Users

Typically, users in a virtual environment are classified into different categories depending on the job function, type of workload and the computational resources consumed. While Security and mobility are of great significance to deliver content, some users require simplicity and standardization while others require performance and personalization including GPU acceleration. Generally we see the following three classes of users:

- **Task Worker** – A task worker is a user who utilizes the organizational resources in a repetitive manner to generate value in a business. These users need a small set of applications which are light on the Compute and Memory requirements. The tasks performed by them on the digital workspace are pre-defined. This type of user requires speed and simplicity in accessing the apps and desktops in a secure way. The resources for these users are standardized and personalization is not a requirement here since the data generated by this type of users is saved on a backend server and is accessible on demand. Data Loss criticality for this type of users is pretty low. The workstations for these users are static and do not require mobility in most cases. Task Worker are the prime candidates for a digital workspace. Receptionists, Call center employees, retail employees, warehouse workers etc. are the best example of a task worker.
- **Knowledge worker** – A knowledge worker uses their expertise and experience to generate value in a business. The daily actions performed by this type of user are not repetitive and require innovation. These workers typically access internet, use email, create intricate and elaborate documents, presentations and spreadsheets. In most of the cases these users need some degree of personalization. Freedom to access the apps and desktops from anywhere on any device is often sought after by this type of user. As a knowledge worker becomes more specialized the degree of personalization and data loss criticality increases. This type of user is mostly against standardization and appreciates good performance and exceptional user experience. Accountants, analysts, decision makers and managers fall under this category of users.
- **Power User** – Power users are coders, application developers, graphic designers and vfx and cgi experts. With the advent and support of graphics accelerators in a VDI environment it has become easier for this group of users to migrate to digital workspace. With technologies like NVIDIA Grid it has become possible to provide these users with a powerful GPU at a comparatively lower cost. HDX 3D Pro utilizes a graphics processing unit (GPU) for hardware acceleration. Citrix Virtual Apps and Desktops take advantage of this protocol to deliver high quality visual experience. Exceptional Graphics processing power can be assigned to these users with the help of GRID vGPU, making the overall VDI experience smooth and immersive. This kind of users require personalization and data loss criticality is on higher side. Mobility is also an important factor here as the number of employees working from home is increasing day by day

## Summary

HPE ProLiant DL380 Gen10, DL360 Gen10, and DL560 Gen10 server platforms are ideal for VDI workload. With their significant technology enhancements, the HPE Gen10 platforms deliver a more secure, higher performant end-user experience and with Citrix Virtual Apps & Desktop, along with NVIDIA GPU, a wide range of users from Knowledge based to power user can benefit from this solution.



## Appendix A: NVIDIA Accelerators

Organizations are now looking forward to a package of end to end solutions which integrates compute, storage, memory and required software from different vendors. These end to end solutions are designed to cater the needs of all kinds of users. One such category of users are UI designers, video editors, VFX and CGI experts, etc. As more and more enterprises transform to a modern digital workspace, these users are demanding better performance and graphics processing power. To support the demand of such users and to be future ready, HPE ProLiant DL Gen 10 servers support a host of NVIDIA GPUs.

### NVIDIA GRID

NVIDIA GRID architecture enables NVIDIA GPUs to power virtual desktops and applications in a data center environment and accelerate workflows resulting in an optimized user experience and improved productivity. Figure 8 depicts how the NVIDIA GRID software works in conjunction with the hypervisor and virtualizes the GPU cards so that they can be shared across multiple virtual desktops or applications, while the guest OS on the VM accesses the vGPU like a pass-through device.

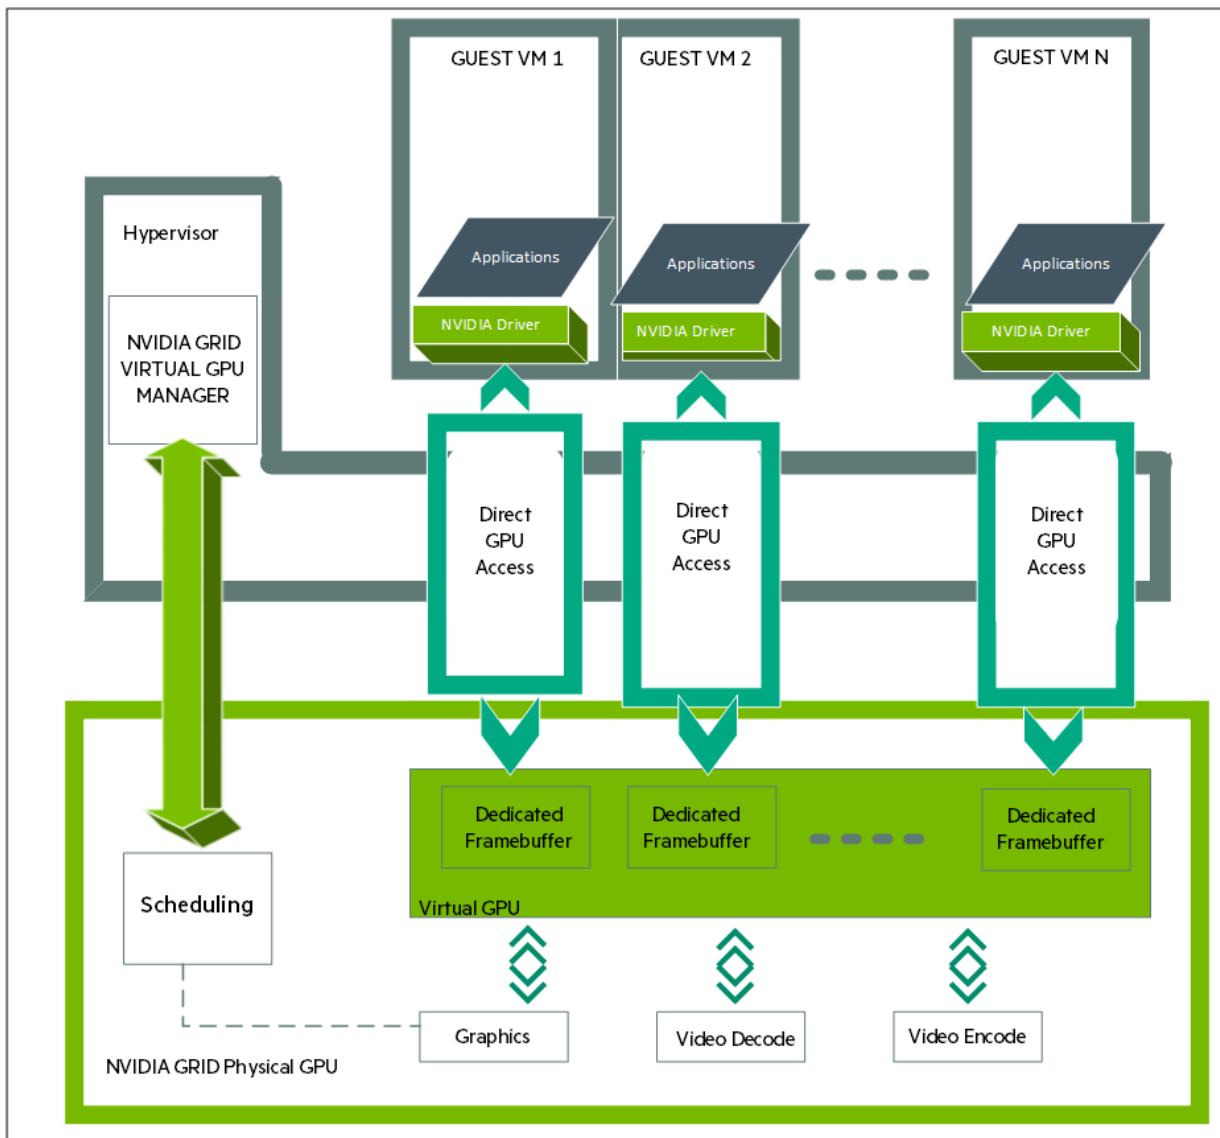


Figure 8. NVIDIA GRID architecture





### NVIDIA Licensing

Different type of NVIDIA GPU cards use different License models. NVIDIA vGPU software allows the use of GRID vPC, GRID vApps, Quadro vDWS and vCompute Server. Depending on the use case and the choice of physical GPU, a choice can be made of the License type. GRID vPC, GRID vApps, and Quadro vDWS are available on a per Concurrent User (CCU) model. NVIDIA vGPU editions can be purchased by enterprises either as perpetual licenses or as an annual subscription. Figure 9 illustrates the NVIDIA Licensing Models and their use cases.

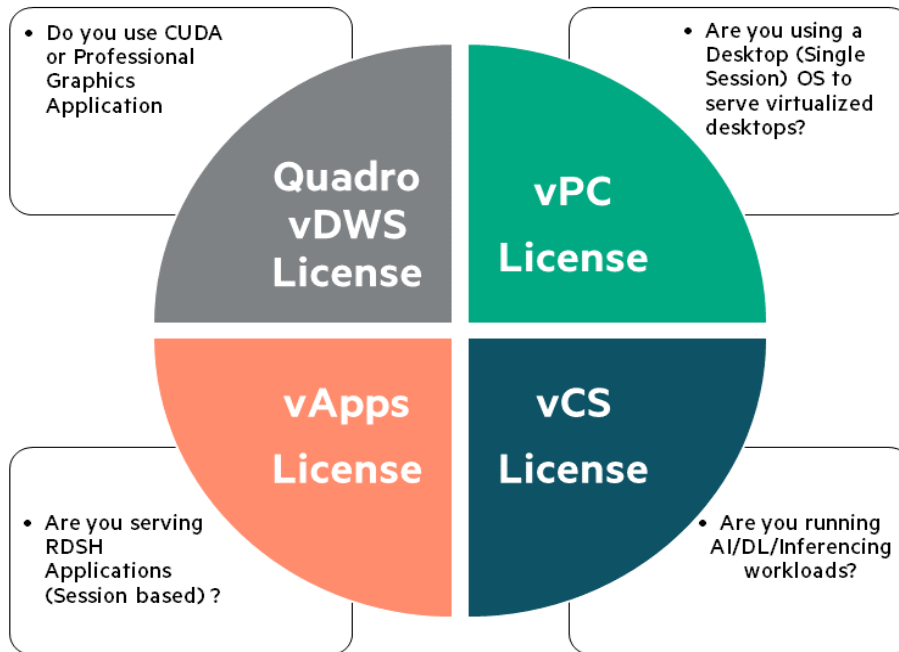


Figure 9. NVIDIA GRID Licensing models and use cases

A perpetual enterprise license allows the hardware to be used indefinitely. Users opt to license using this model are required to subscribe to SUMS for three years. The Support Updates and Maintenance Subscription (SUMS) can be renewed on a yearly basis after the expiration of the initial subscription. Alternatively, NVIDIA GPUs can be licensed with annual enterprise subscription, which is active for a fixed period as defined by the terms of the subscription license. To be kept active, the license will need to be renewed at the end of the subscription period. The subscription license includes the software license and production level SUMS for the duration of the license subscription period.

NVIDIA GRID licensing model supports a broad spectrum of vGPU profiles that helps in aligning resources aptly and avoid under/over provisioning of vGPUs. The servers running on accelerated applications can be configured with 16GB profile, which will be accessible by multiple users. Additionally, NVIDIA offers profile sizes of 1GB, 2GB, 4GB, 8GB, and 16GB for Quadro vDWS. Similarly, the graphics workload can be accelerated with 1GB or 2GB profiles for the accelerated desktops.



## Appendix B: Bill of material

The following BOMs contain electronic license to use (E-LTU) parts. Electronic software license delivery is now available in most countries. HPE recommends purchasing electronic products over physical products (when available) for faster delivery and for the convenience of not tracking and managing confidential paper licenses. For more information, please contact your reseller or an HPE representative.

### Note

Part numbers are at time of publication/testing and subject to change. The bill of materials does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please consult with your HPE Reseller or HPE Sales Representative for more details. [hpe.com/us/en/services/consulting.html](https://hpe.com/us/en/services/consulting.html)

Table 1 provides Bill of materials for reference purpose.

**Table 1.** Bill of materials

Quantity	Part Number	Description
1	P20249-XX1	HPE DL380 Gen10 5218 (16C 2.3GHz) 1P 32G NC 8SFF SVR
1	P02498-B21	HPE DL380 Gen10 Intel Xeon-G 5218 16-Core Processor Kit
11	P00924-B21	HPE 32GB (1x32GB) Dual Rank x4 DDR4-2933 Kit
2	P18420-B21	HPE 240GB SATA 6G Read Intensive SFF Multi Vendor SSD
2	P18434-B21	HPE 960GB SATA 6G Mixed Use SFF Multi Vendor SSD
2	826704-B21	HPE DL Gen10 x16/x16 GPU Riser Kit
1	867810-B21	HPE DL38X Gen10 High Performance Fan Kit
2	830272-B21	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
1	871830-B21	HPE DL38x Gen10 8x 6-pin Cable Kit
1	P24743-XX1	HPE DL360 Gen10 6248R (24C 3.0GHz) 1P 32G NC 8SFF SVR
1	P24487-B21	HPE DL380 Gen10 Intel Xeon-G 6248R 24-Core Processor Kit
11	P00924-B21	HPE 32GB (1x32GB) Dual Rank x4 DDR4-2933 Kit
2	P18420-B21	HPE 240GB SATA 6G Read Intensive SFF Multi Vendor SSD
2	P18424-B21	HPE 960GB SATA 6G Read Intensive SFF Multi Vendor SSD
1	804331-B21	HPE Smart Array P408i-a SR 12G SAS Modular Controller
1	727055-B21	HPE Ethernet 10Gb 2-port 562SFP+ Adapter ( Intel based)
1	P01366-B21	HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit
1	865414-B21	HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
2	R0Z45C	HPE NVIDIA Quadro RTX 6000 Graphics Accelerator <b>(Optional for GPU accelerated workloads)</b>



## Reference Architecture

### Resources and additional links

HPE Client Virtualization Reference Architectures, [hpe.com/info/clientvirtualization-ra](https://www.hpe.com/info/clientvirtualization-ra)

HPE Pointnext Services, <https://www.hpe.com/us/en/services/pointnext.html>

HPE Servers, [hpe.com/servers](https://www.hpe.com/servers)

HPE Server Storage, [hpe.com/us/en/servers/server-storage.html](https://www.hpe.com/us/en/servers/server-storage.html)

HPE Server Networking, [hpe.com/us/en/servers/networking.html](https://www.hpe.com/us/en/servers/networking.html)

HPE Server Memory, [hpe.com/us/en/servers/memory.html](https://www.hpe.com/us/en/servers/memory.html)

HPE Server Power Supplies, [hpe.com/us/en/servers/server-power-supplies.html](https://www.hpe.com/us/en/servers/server-power-supplies.html)

HPE Rack, Power and Cooling, [hpe.com/us/en/integrated-systems/rack-power-cooling.html](https://www.hpe.com/us/en/integrated-systems/rack-power-cooling.html)

Citrix, [citrix.com](https://www.citrix.com)

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

© Copyright 2020 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.

Citrix is a trademark of Citrix Systems, Inc. and/or one more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. Microsoft and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Intel is a trademark of Intel Corporation in the U.S. and other countries.

a50001484enw, May 2020

