Hewlett Packard Enterprise | NVIDIA

HPE High Performance Remote Visualization Solution

High-performance, density-optimized, and scalable solutions for scientific and engineering applications

Technical white paper

Contents

Executive summary

Designers and engineers in today's environments are working with increasingly complex models and simulations. Whether it is increased reality of an animated movie, or an in-depth fly-through of the engine compartment of a prototype vehicle, the need for higher performance, more data, and increased realism in visualization is the order of the day. Recent advances in graphics processing unit (GPU) technology, virtualization software, and remote protocols support this quality and responsiveness, enabling remote visualization solutions that deliver more performance, robustness, and improved cost of ownership than ever before.

At the same time, corporations face challenges in delivering performance and functionality as engineers become more mobile. Sustaining a geographically dispersed team requires the ability to support secure, real-time collaboration around the globe—no matter how big or small the data set is. Yet, working with a geographically distributed workforce can lead to increased infrastructure costs, slow file load times, and security implications and risks.

How can corporations solve these issues? Losing the latest semiconductor chip design file, or disallowing access to remote engineers, or implementing a complicated and expensive on-premise infrastructure is all out of the question. What is needed is a way to securely and cost-effectively allow engineers access to graphic-intensive files from remote locations on any device from which they choose to be productive.

To address these challenges, remote visualization was introduced into the marketplace. Remote visualization transmits complex 2D and 3D images from a sender system across standard computer networks to remote users. The remote users interact with the host sender system and its applications as if they were using a local workstation, while the host sender systems and application data remain securely in the data center. It enables display of 2D or 3D graphics, full motion video, and multidisplay on a system that only has 2D graphics card(s), and allows multiple users to interact with the same model in real time.

With its comprehensive portfolio of high-performance hyperscale systems and extensive alliances with leading GPU suppliers, HPE is able to deliver a remote visualization solution optimized for each customer's situation. This paper describes remote visualization solutions tailored for demanding compute- and visualization-intensive applications, like those in the energy, manufacturing, and life sciences industries. These solutions combine HPE Apollo density-optimized systems with integrated high-performance GPUs, like those provided by NVIDIA®, to deliver performance, storage capabilities, and density to best support simulation and 3D modeling applications. The resulting remote visualization solutions are tightly integrated and tested to provide reliable performance, flexibility, and compelling efficiency and affordability.

Value creation with remote visualization

Remote visualization brings value and improved cost of ownership to high-performance computing workflows in a variety of ways, including security, productivity, optimal resource utilization, and expanded application access from geographically dispersed sites.

- **Security:** Remote visualization allows companies to keep their critical data within the data center; only the pixels from images are transmitted over the network, and that pixel stream is encrypted. This substantially reduces the risk of losing critical corporate information.
- **Productivity:** Remote visualization drives increased productivity by permitting anytime, anywhere access to graphics-intensive models and data. Collaboration becomes seamless, regardless of physical location. Time spent moving data between remote workstations and the data center is eliminated.
- **Manageability:** IT managers that manage tens or hundreds of applications with perhaps thousands of clients can simplify their jobs if their users employ remote visualization, especially in virtualized environments. Load-balancing, better monitoring capabilities, and better control and management of application licenses are examples of improved manageability.
- **Optimal resource utilization:** Virtual desktop environments with embedded GPUs allow IT to provide a flexible and scalable resource to their businesses to meet demand, optimizing hardware, software, and human resources to allow better utilization and cost savings. Cost savings can be particularly striking, especially considering the normal situation where every workstation must have licenses for all applications in use; using remote visualization, the servers are assigned licenses, not the remote users.
- **Expanded application access:** Highly skilled workers increasingly value work-life balance and the ability to work from the location of their choosing, untethered to a fixed office location. The ability to work from a preferred location allows companies to retain their key scientists and engineers available for a longer period prior to retirement by providing them with a great place to work.

HPE remote visualization strength

Hewlett Packard Enterprise has deep expertise in remote visualization, gained through decades of experience working with leading companies across multiple industries. This depth and breadth of experience is leveraged to develop optimized systems for high-performance 3D remote graphics in demanding industries. HPE has extensive capabilities to design and implement solutions for specific customer environments, maximizing the value created by remote visualization systems. `

A broad portfolio of best-in-breed graphics solutions, server platforms, software, and partnerships are key components of remote visualization solutions from HPE. Each solution is optimized and tested to ensure it meets the customer's operational and performance-price objectives. Complementing these integrated solutions is a comprehensive suite of service capabilities, which assure optimum system performance and high availability.

Finally, HPE has deep industry expertise, supporting customers in manufacturing, energy, and other industries with technology, solution architecture, and account management teams that focus on high-performance computing (HPC).

HPE solutions for remote visualization

IT organizations continue to leverage general-purpose servers for their traditional virtual desktop infrastructure (VDI) with the goal to manage their corporate IT environments centrally. Unfortunately, this type of infrastructure does have its limits, especially as demands increase for more network bandwidth and performance to support both compute and graphics.

To address these limits initially, companies may choose to implement workstation-class machines within their data centers. End users can then connect to these servers in a one-to-one model through an endpoint device. This type of implementation gives the most demanding users the power and performance required to process and interprets sizeable output files. However, for those organizations who need to support multiple users, this type of implementation comes at a cost.

Hewlett Packard Enterprise bridges the gap between the classical virtual desktop infrastructures and bare-metal implementations with a variety of solutions across its product portfolio (Figure 1). Each solution plays a part in the HPE end-to-end infrastructure that supports virtual desktops and accelerated graphics by meeting the needs of specific users across the organization.

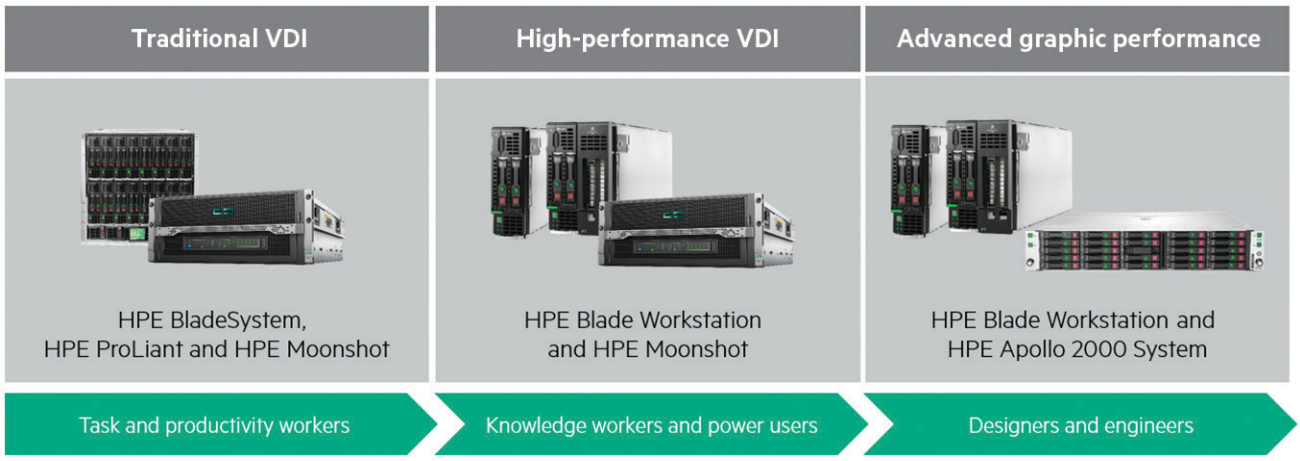

Figure 1. Hewlett Packard Enterprise provides a variety of remote visualization solutions across its product portfolio

IT organizations can choose from a variety of methods to enable remote visualization, including virtualized and bare-metal implementations. Each implementation offers different graphical performance at the endpoint device, varies in cost, places different demands on network bandwidth, and requires different hardware components.

For companies that need to support engineers and designers around the world, user experience is the gold standard. These users expect the ability to access 3D graphics and simulations via an interface that quickly loads data files and provides smooth model manipulation. This user experience is expected whether the engineer is located close to the compute resource or halfway across the world.

The HPE High Performance Remote Visualization Solution delivers encrypted 3D graphics and simulations from a central compute resource, enabling global collaboration and user experiences that rival those on a high-end workstation—no matter the user's location. With the HPE High Performance Remote Visualization Solution, companies can also realize the cost-effectiveness from sharing GPUs across multiple users.

Based on the HPE Apollo 2000 System, with its ability to handle 3D graphics and resource-intensive computer-aided engineering workloads, the HPE High Performance Remote Visualization Solution can be deployed across a number of industries.

Technology overview

Remote visualization is available across Hewlett Packard Enterprise's scalable portfolio. HPE works closely with developers of GPU cards, including NVIDIA, to test and validate technologies, and ensure customers can leverage new imaging technology confidently and quickly.

There are multiple approaches to enable remote visualization, offering different graphical performance at the endpoint device, a range of costs, and placing different demands on network and hardware capabilities. Recent advances in GPU technology and remote protocols have enabled performance capable of remotely supporting multiple engineering or scientific computing users.

For compute- and data-intensive simulations, rich and dynamic renderings benefit greatly from low-latency, high IOPS-capable servers, in addition to dedicated GPU allocations. For example, GPU-enabled Apollo 2000 Systems are easily integrated into the HPC infrastructure, and support the key technologies needed for remote visualization with high-density and attractive price performance.

Remote visualization methods

Three models of remote visualization are commonly deployed and are discussed in the following sections:

- Virtualized GPU—Allowing several clients to simultaneously share GPU resources
- Pass-through GPU—Allocating each GPU to single client use
- Dedicated Hardware (also known as bare metal), in which the visualization server is dedicated to one client at a time—essentially serving as a remote workstation tightly integrated with a central data center

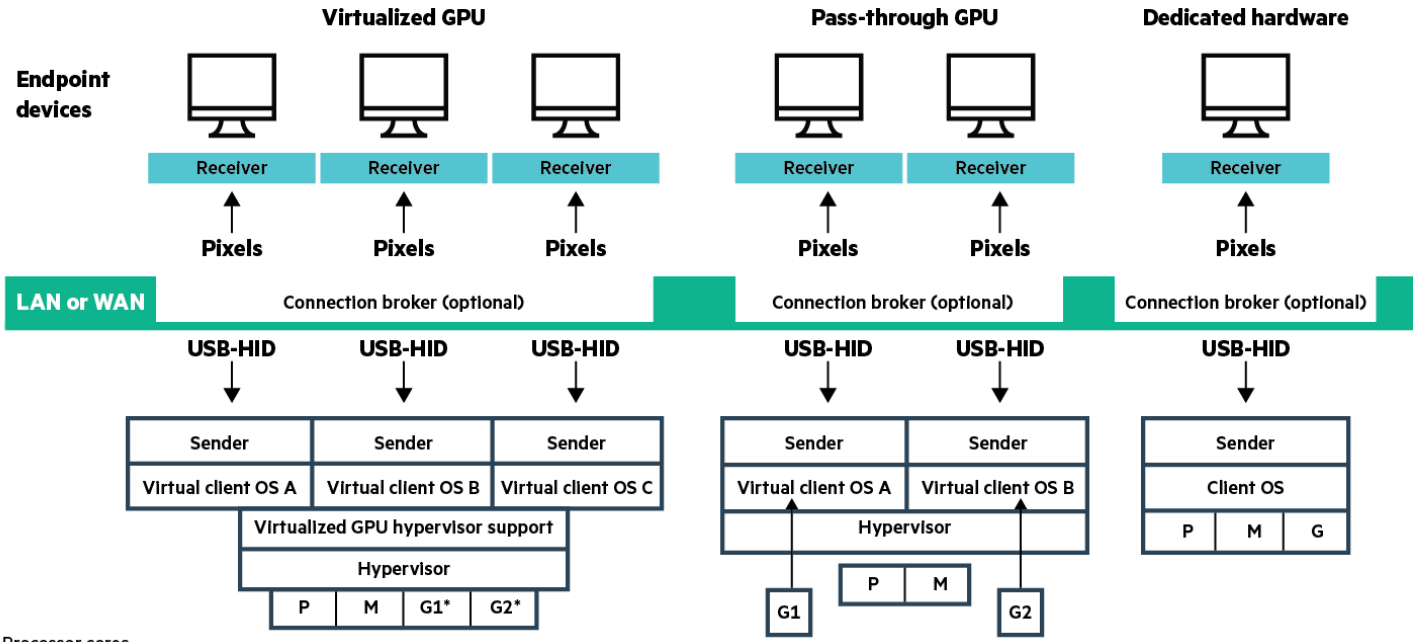

P—Processor cores
M—Memory
G—GPU
USB-HID—USB human interface device(s)

Figure 2. Three models of remote visualization

The following sections describe each of these solutions in more detail.

Virtualized GPU—Virtual machine (VM) with virtualized GPU
The virtualized GPU takes advantage of specialized GPU hardware and software, such as the NVIDIA Tesla M60 combined with NVIDIA GRID software, which is designed to support fully virtualized environments. In such an architecture, each VM can be assigned a fixed portion of the graphics card resources, giving each virtual desktop (VM) multiplexed access to the hardware GPU installed in the host. Additional software includes a hypervisor that is capable of supporting the virtualized GPU, virtual GPU management software in the hypervisor, as well as an optional connection broker. This model enables a high-performance, immersive virtual workspace, with a user experience that rivals that of a physical device, accessible from anywhere, on any device. Virtualized GPU also provides improved economics with multiple users (up to 16, in the case of the NVIDIA Tesla M10) simultaneously sharing each GPU, resulting in lower per-user cost.

Pass-through GPU
Also referred to generically as direct attached GPU or vendor-specific vDGA (VMware®) and GPU pass-through (Citrix®). This method allows discrete GPU devices to be directly mapped to a virtual machine for dedicated 1:1 use by the VM. The virtual machine has full and direct access to the GPU, including the native graphics driver, allowing for full workstation class graphics and GPU compute performance in a virtual machine. Typically intended for high-end 3D and GPU compute users, the GPU device is directly owned and managed by the VM operating systems just as in a desktop workstation. The GPU driver is loaded within the virtual machine. Dual GPU cards can support two simultaneous clients in this mode. For high-performance applications, pass-through GPU delivers the highest performance client or end-user experience available in a virtualized implementation.

Dedicated hardware
This method is the classic hosted-workstation and workstation blade remote visualization technology. The client OS is installed directly on the server hardware, and system resources are dedicated to a single user (at a time). End users connect to the graphics-enabled server or workstation via remote protocols via an endpoint device. This method is still used today by customers that prefer a bare-metal implementation or demand the power and the performance of dedicated hardware.

HPE portfolio for high-performance Remote Visualization

HP Remote Graphics Software

HP Remote Graphics Software (RGS) is an industry-leading remote visualization solution, developed by HP Inc., the market leader in workstations. With HP RGS, users can securely access graphics-rich applications located on a remote system and collaborate in real time from any PC, thin client, or Windows® tablet. All applications run natively on the remote system (traditional workstation, graphics server, or graphics blade), taking full advantage of the system's graphics resources, including 3D graphics support for the latest versions of OpenGL and DirectX. HP RGS is designed to provide a seamless, local-like feel for even the most demanding 3D professional applications. HP RGS is ideal for connecting to virtual workstations hosted in the data center or even the cloud. HP RGS was first released to the market in November 2003. One of the early RGS codecs was developed by Hewlett Packard Labs and used with the NASA's Mars rover program.

HP RGS employs several innovative technologies developed by HP Inc. Components such as a patented compression/decompression algorithms and HP-developed software form a modularized software product that can be installed on a wide variety of HP and non-HP systems. High-performance compression enables real-time visualization for graphic-heavy applications, and facilitates remote visualization over lower bandwidth or poorer quality network connections.

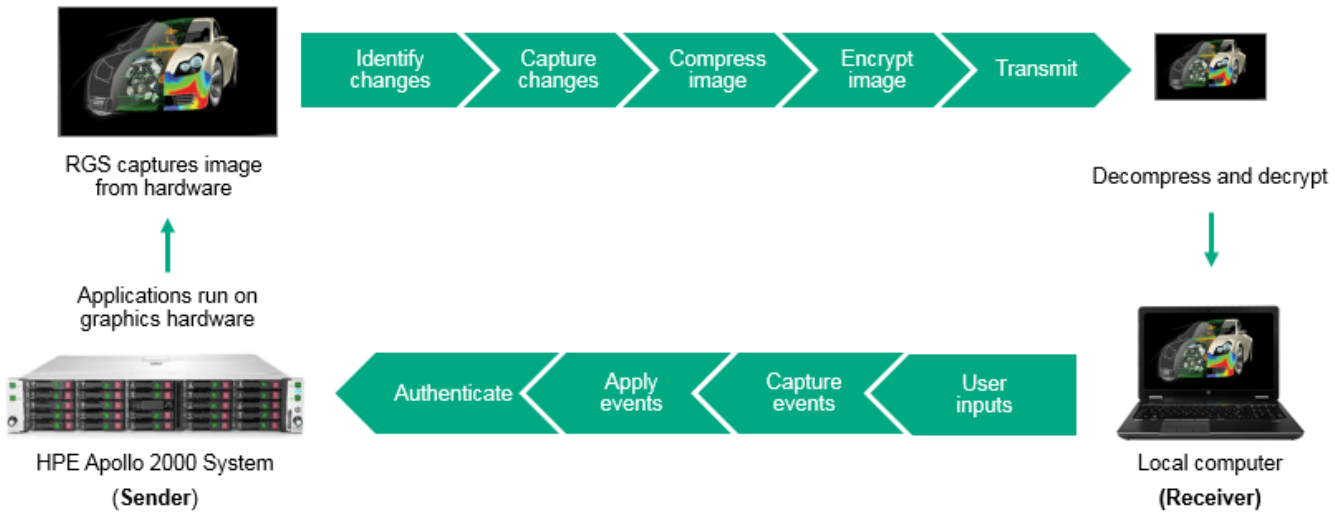

Figure 3. An overview of HP Remote Graphics Software

Because it requires no application modifications, HP RGS may be used in conjunction with any application software, and works seamlessly with Windows 7, Windows 8.1, Windows 10, Red Hat® 6, Red Hat 7, and SUSE 12 Linux® applications. Virtualization is enabled through support for hypervisors from Citrix, VMware, and Red Hat.

In addition to HP RGS, the HPE High Performance Remote Visualization Solution supports HDX 3D Pro and VMware® Horizon™ Blast Extreme or VMware Horizon PCoIP as remoting protocols. While Horizon and HDX 3D Pro simply transport 3D graphics for viewing purposes, HP RGS holds the performance advantage on poor networks—making it the ideal choice for corporations with engineers located in low bandwidth, high latency network locations.

Table 1. Feature benefits of different remoting protocols

	Multi-OS (Linux/Windows)	vGPU support	Hardware encoding on NVIDIA GPU	WAN Optimizer
HP RGS	✓	✓	✓	✓
NICE DCV	✓	✓	✓	
HDX 3D Pro		✓	✓	
VMware Horizon	✓	✓	✓	

HPE Apollo systems

When massive data demands extreme performance and scale, customers are choosing HPE Apollo. This high-density server family delivers breakthrough performance with efficient rack-scale compute, storage, networking, power, and cooling for the most demanding HPC, massive data analytics, and object storage workloads. HPE Apollo combines a modular design with innovative power distribution and air-cooling techniques to provide up to four times more performance per square foot than standard rack servers. With HPE Apollo high-density servers, the data-driven organization accelerates time to value.

HPE Apollo 2000 System

The HPE Apollo 2000 System is a high-density solution that packs performance and workload capacity into a small amount of data center space. In addition to the traditional data center, the Apollo 2000 is also ideal for the space-constrained data center or remote site. Apollo Systems are designed to exploit a modular and efficient infrastructure to provide storage density and operating efficiency.

Flexible configuration options make the Apollo 2000 System a great fit for a variety of workloads, including HPC clusters and remote visualization. The ability to mix and match both 1U and 2U servers in the same chassis and the unique drive mapping flexibility can help to support optimized configurations for many applications. The Apollo 2000 System Chassis, or groups of chassis, can be custom-configured to act as affordable, modular, 2U building blocks for specific implementations at scale—and for future growth.

There are two server models available for the Apollo 2000. The XL170r is for compute intense deployments not requiring support for accelerators or GPUs, and offers extreme server density. The XL190r provides the 2U height to support GPU cards, but reduces the server's width by half when compared to traditional rack servers, giving customers the density of 4 GPUs (2 GPUs per HPE ProLiant XL190r Server) in a 2U Apollo 2000 System Chassis.

HPE BladeSystem

HPE BladeSystem is a modular infrastructure platform that converges servers, storage, and network fabric to accelerate operations and speed delivery of applications and services running in physical, virtual, and cloud-computing environments. Because the core infrastructure is shared, capital costs can be significantly lower. Blades share power, cooling, network, and storage infrastructure at the BladeSystem enclosure level.

NVIDIA GPUs

For more than two decades, NVIDIA has pioneered visual computing, and the art and science of computer graphics. With a singular focus on proliferating the application of the GPU to the most important technology transitions taking place, NVIDIA offers specialized platforms for the gaming, professional visualization, data center, and automotive markets. Beginning as a visual computing company, NVIDIA has transformed into a specialized platform company that targets artificial intelligence, deep learning, and professional visualization—where the power of the GPU is essential and uniquely beneficial.

A range of NVIDIA cards are offered for the HPE Apollo 2000 System high-performance remote visualization solution (Table 2). These NVIDIA GPUs are optimized for engineering virtual desktop infrastructures (eVDI) using NVIDIA GRID Virtual GPU Manager. Please note that the M10 card is currently under design investigation and may be supported on the HPE Apollo 2000 System at a later date.

Table 2. NVIDIA Tesla GPU cards offered for the HPE Apollo 2000 System remote visualization solution

	User Density optimized	**Blade optimized**	**Performance optimized**
	M10	**M6**	**M60**
GPU	Quad Mid-Level Maxwell	Single High-end Maxwell	Dual High-end Maxwell
CUDA cores	2560 (640 per GPU)	1536	4096 (2048 per GPU)
Memory size	32 GB GDDR5 (8 GB per GPU)	8 GB GDDR5	16 GB GDDR5 (8 GB per GPU)
Max. vGPU instances	64	16	32
Form factor	PCIe 3.0 Dual Slot (rack servers)	MXM (blade servers)	PCIe 3.0 Dual Slot (rack servers)
Power	225 W	100 W (75 W opt.)	200 W/300 W (225 W opt.)
Thermal	Passive	Bare board	Active/passive

Supporting software options

The key software elements complementing HPE Remove Visualization are a connection broker, virtualization software, and cluster tools. The connection broker simplifies overall system management by managing permissions and making it easier to add users or move users in case of a failure. The virtualization software uses hypervisors to create and manage the virtual machines (VMs).

Cluster managers, such as HPE Insight Cluster Management Utility, support installation of images across the cluster nodes and monitoring of servers—including components such as GPUs, and provide a management interface for administrators.

Connection brokers

In desktop virtualization, a connection broker is a software program that allows the end user to connect to an available desktop by managing a pool of connections to resources such as databases or remote desktops. At a minimum, connection brokers:

- Validate users' credentials and establish connections
- Monitor activity levels of virtual machines
- Handle reassignment of a virtual machine after user disconnect

Popular connection brokers include those provided by Leostream, Citrix, NICE, and VMware.

Hypervisors

Virtualization software is required to deploy the pass-through GPU and virtualized GPU solutions. Enterprise hypervisors supporting VMs on GPUs and qualified with the HPE visualization technology include:

- Citrix XenServer
- VMware® ESXi™

HPE Insight Cluster Management Utility

The HPE GPU ecosystem includes HPE Cluster Platform specification and qualification, HPE-supported GPU-aware cluster software, and also third-party GPU-aware cluster software for NVIDIA Tesla, Quadro, and GRID modules on HPE ProLiant servers. In particular, the HPE Insight Cluster Management Utility (CMU) monitors and displays GPU health sensors such as temperature. HPE Insight CMU also installs and provisions the GPU drivers and the CUDA software. HPE Insight CMU is integrated with popular schedulers such as Adaptive Moab, Altair PBS Professional, and IBM Platform LSF, all of which have the capability of scheduling jobs based on GPU requirements.

NVIDIA GRID

NVIDIA GRID is the enterprise-grade platform for graphics virtualization, leveraging NVIDIA Tesla data center GPU cards (M10, M6, M60) in combination with NVIDIA GRID software. For virtualized environments, NVIDIA GRID allows sharing of virtual GPUs (vGPUs) across multiple virtual desktop and applications instances—without compromising the graphics experience. Application features and compatibility are exactly the same as they would be at the desk. With GRID technology, the graphics commands of each virtual machine are passed from the NVIDIA driver in the guest OS, to the GPU in the host via the NVIDIA GRID Virtual GPU Manager, without translation by the hypervisor. This allows the GPU hardware to be time-sliced to deliver the ultimate in shared virtualized graphics performance (Figure 4). NVIDIA GRID gives IT professionals and data center managers the power to centralize apps and data, providing virtual workspaces with improved security, productivity, and simplified manageability.

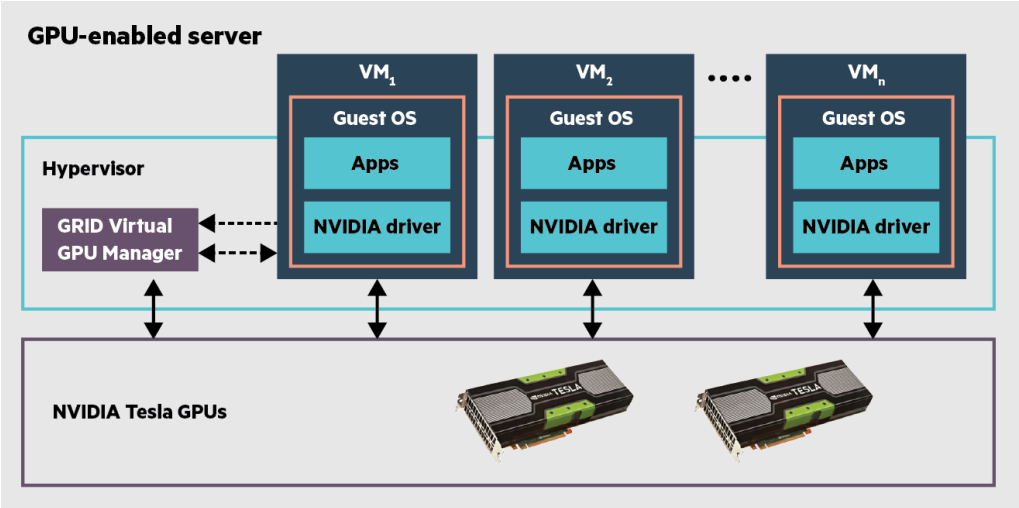

Figure 4. NVIDIA GRID allows sharing of virtual GPUs (vGPUs)

Solution implementation guidelines

This section provides examples of remote visualization implementations utilizing the HPE Apollo 2000 System as the visualization server for accelerated graphics, and the HPE BladeSystem as the platform for dedicated hardware implementations. For applications demanding high-density plus fast and large IOPS and storage capacity, the Apollo 2000 Systems offers the highest IOPS and HDD capacity per 1U of rack space. Also, the Apollo 2000 System chassis has a 2U rack-server form factor, while delivering the cost and power efficiency of a shared cooling and power infrastructure. The design of any remote visualization solution needs to be assessed for each deployment in terms of the quality of user experience required, application workloads and problem size, resolution (pixel) needs at the desktop, and constraints such as network bandwidth and latency.

Remote visualization solutions

A foundation of the Apollo 2000 System with two ProLiant XL190r Servers is used for both the pass-through GPU and virtualized GPU (virtual machine with virtualized GPU can be used for both), and listed in table 3. This method allows direct mapping of discrete GPU devices to a virtual machine (VM) through the hypervisor, providing a dedicated GPU for every VM. The dedicated GPU supports workstation-class graphics and GPU computing in a virtual environment.

The ProLiant XL190r Server can be configured with two GPU cards. Each GPU supports multiple users—for example, the NVIDIA Tesla M60 will support up to 32 simultaneous remote users, for a total of 64 users per ProLiant XL190r Server.[1] When creating a VM with workstation-class graphics performance, you will also need to give it workstation-class resources with at least multicore processing and large memory.

Table 3. Pass-through mode configuration example with HPE Apollo platform

Function	Recommendation
Local client	Workstation, mobile workstation, PC, thin client, Windows tablet, Mac (coming soon)
Receiving agent	HP RGS Receiver
Connection broker (optional)	Leostream Connection Broker
Remoting protocol	HP RGS
Operating system	Windows 7, 8.1, and 10 Enterprise 64-bit Red Hat Enterprise Linux SUSE Linux Enterprise Server
vGPU Hypervisor Support (if Virtualizing GPUs)	NVIDIA GRID
Hypervisor	Citrix XenServer VMware ESXi
Remote hardware	HPE Apollo 2000 System Chassis with 2 x ProLiant XL190r Server nodes **Base node configuration** 2 x Intel® Xeon® processor 2 x high-performance GPUs 512 GB Memory 12 x 1 TB SATA HDD (with array controller if needed for I/O performance)

[1] Note that this is the design maximum; the actual number of users is highly dependent on application and user requirements

Dedicated hardware solution

In this implementation, the server is basically a remote high-performance workstation. There are no virtual machines on the remote system and there is only one user per system. The end user connects to the remote resource via the receiver software. In some cases, a connection broker, such as the Leostream Connection Broker, may be used to simplify overall system management by managing permissions and making it easier to add users or move users in case of a failure.

Table 4. Dedicated remote hardware configuration example with HPE BladeSystem WS460c platform

Function	Recommendation
Local client	Workstation, mobile workstation, thin client, Windows tablet, PC, Mac (coming soon)
Receiving agent	HP RGS Receiver
Connection broker (optional)	Leostream Connection Broker
Remoting protocol	HP RGS
Operating system	Microsoft® Windows 7 Professional (64-bit), Enterprise (64-bit) **Note:** Windows 7 is supported only with Legacy BIOS mode Microsoft Windows 8.1 Professional (64-bit), Enterprise (64-bit) Red Hat Enterprise Linux (RHEL) Desktop 6.5 or later (64-bit only)
Remote hardware	HPE BladeSystem c7000 chassis, with up to 8 HPE WS460c + extension blade: each blade configured with Intel Xeon processors, 128 GB high speed memory, 2 x 960 GB SSD MU Expansion blade allows installation of full-size high-end graphics cards, including 4 NVIDIA M6 cards

Visualization server integration with HPC clusters

The flexible, modular design of the HPE Apollo Systems supports deployment of HPC resources capable of high compute performance, large and fast storage that support data-intensive modeling and visualization. The scale-out compute part of the HPE Apollo System portfolio includes the Apollo 2000 System for hyperscale, remote visualization, and general-purpose scale-out computing; the Apollo 4000 System Family for analytics and object storage; and the Apollo 6000 and 8000 Systems for HPC and supercomputing. This complete range of offerings makes high-density server storage, management, and rack-scale efficiency available to organizations of all sizes, with a tiered approach that provides a logical and simplified starting point for companies seeking to utilize HPC.

The Apollo platform makes it simple for HPC clusters to include a broad range of different node types, all sharing a high-performance network (such as InfiniBand). Nodes can be configured to support specific roles such as visualization and computation. For example, simulations can be run on compute nodes, and data-intensive results transferred to visualization nodes for rapid rendering. HPE Insight Cluster Management Utility simplifies the management of multiple node images, as well as management and monitoring.

Hewlett Packard Enterprise offers in-factory and onsite cluster integration services, quick start options, and cluster training.

Implementation options for virtual desktops

This paper is focused on remote visualization to support high-performance scientific and engineering applications. These use cases are characterized by large data-intensive computation that drives real-time rendering to be co-located with the data set at the HPC center. However, desktop virtualization is deployed across many industries and situations that do not place the same demands on the GPU and server resources. The end user may not need the same maximum resolution or low latency, reducing the per display requirements.

In vGPU mode, HP RGS is fully supported. In this environment, and features such as OpenGL and DirectX offer full 3D graphics API support for workstation-class performance. Citrix XenServer and VMware ESXi currently support fully virtualized solutions.

The range of options and configurations to enable desktop virtualization is extensive and dynamic. HPE systems support a range of GPU cards for visualization as well as compute acceleration. Companies such as Citrix and VMware have ongoing programs to enable desktop virtualization on technologies as soon as new technologies are introduced.

Configuration guidance for HPE High Performance Remote Visualization

The following representative bill of materials can be used as a starting point for developing customer-specific configurations. The optimum configuration for a customer will be significantly influenced by several factors in the technology and business environment including, but not limited to, the applications being used, user requirements, data file sizes, virtualization philosophy, and price-performance considerations.

Each HPE Apollo 2000 System can contain two remote visualization nodes (HPE ProLiant XL190r Servers). The node and chassis configurations are listed in table 5 and table 6. Lists bill of materials associated with the HPE Apollo 2000 System Chassis infrastructure.

Table 5. Remote visualization node (ProLiant XL190r Server) configuration

Description	SKU/PN	Category	Quantity
NVIDIA Tesla M60 RAF Dual GPU Module	M3X67A	Graphics Card—2 GPUs per card	2
HPE XL1x0r Gen9 E5-2670v3 FIO Kit	793026-L21	Xeon Multi-Core (2.3 GHz/12 cores) Processor	1
HPE XL1x0r Gen9 E5-2670v3 Kit	793026-B21	Xeon Multi-Core (2.3 GHz/12 cores) Processor	1
HPE 16GB 2Rx4 PC4-2133P-R Kit	726719-B21	Memory (256 GB per node/64 GB per GPU)	16
HPE 1TB 6G SATA 7.2k 2.5in SC MDL HDD	655710-B21	Hard drives (3 HDDs/3 TB per GPU)	12
HPE Ethernet 10Gb 2P 546FLR-SFP+ Adapter	779799-B21	Networking	1
HPE ProLiant XL190r Gen9 Server	798156-B21	Server	2
HP RGS 7 Single E-LTU	K8Z51AAE	RGS Sender software	1 per end point device

Table 6. Apollo 2000 System Chassis configuration

Description	SKU/PN	Category	Quantity
HPE Apollo 2600 24SFF CTO Chassis	798153-B21	Enclosure	1
HPE 1400W FS Plat Pl Hot Plug Power Supply Kit	720620-B21	Power supply	2
HPE DL2000 Hardware Rail Kit 2000	611428-B21	BT others	1

Notes

1. Graphics solutions include NVIDIA Tesla M60 in combination with NVIDIA GRID software.

2. Processor, memory, and hard drive selection and configuration are highly dependent on use case.

Summary—High Performance Remote Visualization capabilities and expertise

Our comprehensive suite of high-performance remote visualization solutions comprised of high-performance server platforms, industry-leading graphics, software, and partnerships, deliver uncompromising performance. Hewlett Packard Enterprise has deep expertise in remote visualization gained through decades of experience with leading companies across industries. This depth and breadth of experience is leveraged to develop optimized systems for high-performance 3D remote graphics for the most demanding applications.

Hewlett Packard Enterprise develops and implements comprehensive testing programs to cover the hardware and software stack as a means to deliver high reliability systems and achieve industry-leading quality and reliability. Customers also have access to HPE's comprehensive suite of service capabilities, which assure optimum system performance and high availability. These capabilities provide a full range of value-added services that can be tailored to meet each customer's unique needs.

Resources

Implementing a proof of concept

As a matter of best practice for all deployments, Hewlett Packard Enterprise recommends implementing a proof of concept using a test environment that matches as closely as possible to the planned production environment. In this way, appropriate performance and scalability characterizations can be obtained. For help with a proof-of-concept, contact an HPE Services representative or your HPE partner.

Additional information

Contact your local HPE representative for questions about HPE high-performance remote visualization solutions and additional HPC solutions.

- HPC and Big Data Industry Solutions website: hpe.com/info/hpc-bigdata-industrysolutions
- HPE Apollo Systems website: hpe.com/us/en/servers/apollo.html
- HP RGS website: hp.com/go/rgs
- NVIDIA GRID website: nvidia.com/grid

Learn more at

hpe.com/info/hpc-bigdata-industrysolutions

Sign up for updates

4AA6-8108ENW, November 2016