



Enhancing the performance of the 3DEXPERIENCE throughout the product lifecycle

Unparalleled productivity, performance, and flexibility with
HPE systems and NVIDIA GPUs



Contents

Executive summary.....3

Product design in the interconnected era of experience.....3

GPU-accelerated VDI for high-performance CAD.....4

Solution architecture.....6

Solution components.....7

Designing the solution for 3DEXPERIENCE CATIA apps.....9

3DEXPERIENCE CATIA apps implementation: Customer example.....12

HPE advantage for 3DEXPERIENCE CATIA apps in manufacturing.....13

HPE and NVIDIA resources.....13



Recent advances in GPU technology, virtualization software, and remote protocols are enabling remote visualization through VDI solutions that deliver more performance, robustness, and improved TCO than ever before.

Executive summary

Manufacturers are under severe competitive pressures to develop, deliver, and support high quality and more custom products and services faster and at lower costs. To address these challenges, manufacturers are embracing digital manufacturing or smart manufacturing. The **3DEXPERIENCE®** from Dassault Systèmes is a leading digital manufacturing platform used extensively in product design and development.

Designers and engineers in today's environments are working with increasingly larger and complex models. Recent advances in GPU technology, virtualization software, and remote protocols are enabling remote visualization through virtualized desktop infrastructure (VDI) solutions that deliver more performance, robustness, and improved total cost of ownership (TCO) than ever before.

Working jointly with NVIDIA®, Hewlett Packard Enterprise provides a comprehensive portfolio of enterprise-grade, high-performance systems and software, and high-value services to help manufacturers implement the **3DEXPERIENCE** and VDI throughout value chain and product lifecycle. Many of these systems are already certified by Dassault Systèmes and deliver unparalleled productivity, performance, and flexibility.

In an environment of constant enhancement, product modifications and changes are expected. Recently, NVIDIA announced additional features and functions that will enhance the performance of the **3DEXPERIENCE** in the future:

- **Run multi-GPU workloads with NVIDIA Quadro Virtual Data Center Workstation (Quadro vDWS)**—Improvement in virtual GPU performance by aggregating the power of up to four NVIDIA Tesla GPUs in a single virtual machine (VM) for the most graphics- and compute-intensive rendering, simulation and design workflows.
- **Live migration VMware vSphere® vMotion®**—The IT group can migrate live NVIDIA GPU-accelerated VMs without impacting users or requiring scheduled downtime. This saves valuable time and your IT resources.
- **Support for NVIDIA Tesla T4 GPUs**—2X the framebuffer in the same low-profile, single-slot form factor as the previous generation Tesla P4.

The purpose of this white paper is to give guidance in sizing HPE Apollo, HPE ProLiant, and HPE Synergy Gen10 servers combined with NVIDIA GPUs for the **3DEXPERIENCE**. Data for this paper is sourced from a combination of recommendations from HPE, NVIDIA, and Dassault Systèmes engineering experts, along with customer data from actual deployments. As with any deployment, the information provided in this white paper is best used as a starting point.

Product design in the interconnected era of experience

To stay ahead of competitors and win market share in today's dynamic environment, manufacturers are becoming nimbler and embracing change and digital transformation. This profound digital transformation in the manufacturing industry is due to information technology (IT) and operational technology (OT) convergence and integration. Digital manufacturing, also known as Industry 4.0 or smart manufacturing, is the convergence of data and information throughout product lifecycle to improve the overall quality and effectiveness of a company's products, manufacturing processes, supply chain, and customer engagement.

This **3DEXPERIENCE** platform enables the convergence of information, operations, and physical production systems so that manufacturers and their suppliers can quickly respond to quality issues or design changes. This platform also supports collaboration in a way that was previously not possible, both internally between engineering and production and externally with suppliers to effectively overcome competitive pressures to deliver high-quality products faster and at lower costs.

All 3DEXPERIENCE native apps are supported on certified configurations: native apps from 3DEXPERIENCE platform, CATIA®, DELMIA®, ENOVIA®, SIMULIA®, SOLIDWORKS®, GEOVIA® and BIOVIA®. However, for the purposes of this guide, reference is being made to 3DEXPERIENCE CATIA apps only.



HPE and NVIDIA customers have several options for virtualizing and enhancing their computer-aided design (CAD) environments through remote visualization.

GPU-accelerated VDI for high-performance CAD

Designers and engineers in today's environments are working with increasingly complex models and simulations to build high-quality products. This requires the CAD applications to be more responsive. Recent advances in NVIDIA GPU technology, virtualization software from NVIDIA and others, and remote protocols support this need for high-quality and responsiveness, enabling virtualization solutions to deliver professional workstation performance robustness and improved total cost of ownership (TCO) more than ever before.

Remote visualization transmits complex 2D and 3D images from a sender system across standard computer networks. 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 secure in the data center in real time.

With GPU virtualization, multiple engineers and designers can have a high-quality 3D experience viewing PLM data on any device, supporting the need for mobility while keeping data and IP protected in the data center. GPU virtualization is also helping manufacturers to cost-effectively scale general-purpose VDI with a high-quality user experience.

Virtual workstations with virtual GPU performance bring value and better TCO to high-performance computing workflows in a variety of ways:

- **Increases 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. All members of the ecosystem (suppliers, users, contractors) have secure access.
- **Improves productivity, collaboration, and mobility:** Remote visualization drives increased productivity and mobility by permitting anytime, anywhere secure 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.
- **Improves manageability:** IT managers who manage tens or hundreds of applications with perhaps thousands of customers 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.
- **Optimizes 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. However, using remote visualization, the servers are assigned licenses, not the remote users. This consolidation reduces license and other operating costs.
- **Drives consistency:** Data version control of PLM data is enforced in the data center for more consistency.

Many options are available for remote visualization implementation. The straightest forward is the dedicated hardware/bare-metal option where a visualization server is dedicated to one customer at a time.

The other option is the virtual GPU option.

Virtual GPU options

Three models of remote visualization are commonly deployed and are discussed (Figure 1):

- **Fractional GPU**—Allowing several customers to simultaneously share GPU resources
- **Virtual whole GPU**—Allocating each GPU to single customer use
- **Multi-GPU**—Aggregate up to four virtual GPUs per virtual machine (**currently not a supported environment for 3DEXPERIENCE CATIA apps**)



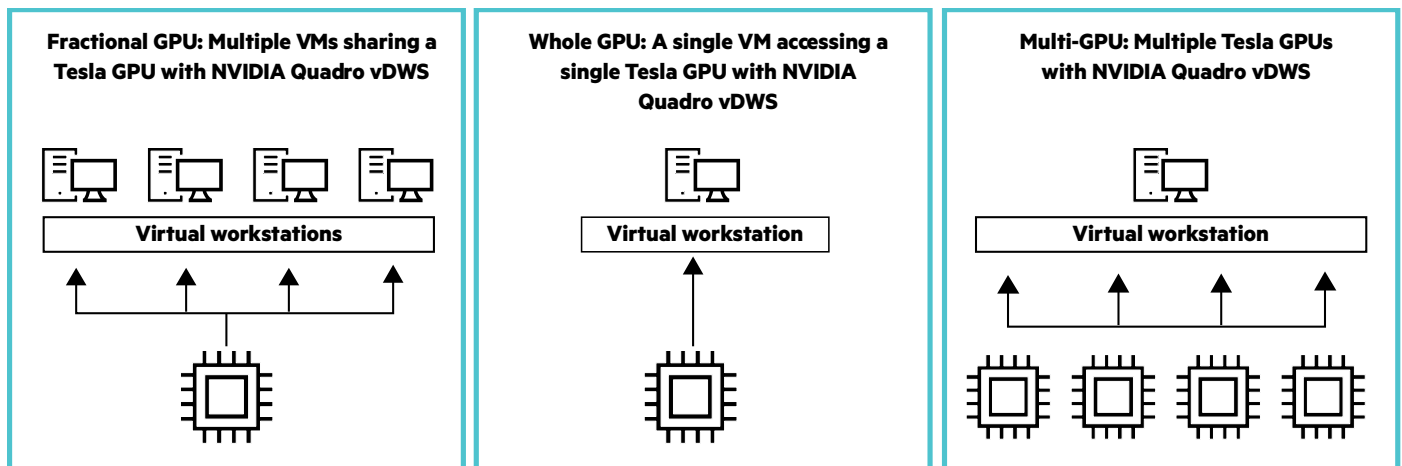


Figure 1. Virtualized GPU options

Each of these remote visualization methods are detailed in the following:

- **Fractional GPU-Virtual machine (VM) with Fractional GPU:** The virtualized GPU takes advantage of specialized GPU hardware and software, such as the NVIDIA Tesla P40, P4, P6, and T4 combined with NVIDIA Quadro vDWS software, which is designed to support fully virtualized environments. In this method, each VM can be assigned a fixed portion of the graphics card resources, giving each virtual desktop 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 remote protocol and an optional connection broker. This model enables a high-performance immersive virtual workspace with a user experience that matches that of a physical device, accessible from anywhere, on any device. Virtualized GPU also provides improved economics with multiple users simultaneously sharing each GPU, resulting in lower per-user cost. This mode of GPU option is suitable for **light** or **medium** users of **3DEXPERIENCE CATIA** apps.
- **Whole GPU:** This is also referred to generically as direct-attached GPU or vendor-specific vDGA (VMware®) and GPU pass-through (Citrix®). This method allows discrete GPU devices like NVIDIA M3000 SE to be directly mapped to a virtual machine for dedicated 1:1 use by the VM. The VM has full and direct access to the GPU, including the native graphics driver, allowing for full workstation-class graphics and GPU compute performance.
- 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 VM. Dual GPU cards can support two simultaneous customers in this mode. For high-performance applications, pass-through GPU delivers the high-performance customer or end-user experience available in a virtualized implementation. This mode of GPU option is suitable for **power** users.
- **Multi-GPU:** Users can experience significant improvement in virtual GPU performance by aggregating the power of up to four NVIDIA Tesla GPUs in a single VM for the most graphics- and compute-intensive rendering, simulation, and design workflows. This mode of GPU option is suitable for **power** users who need to speed-up their computation or whose models exceed the memory available to a single GPU.

Solution architecture

A high-level solution architecture—HPE servers with NVIDIA GPUs for **3DEXPERIENCE CATIA** apps—is shown in the following figure.

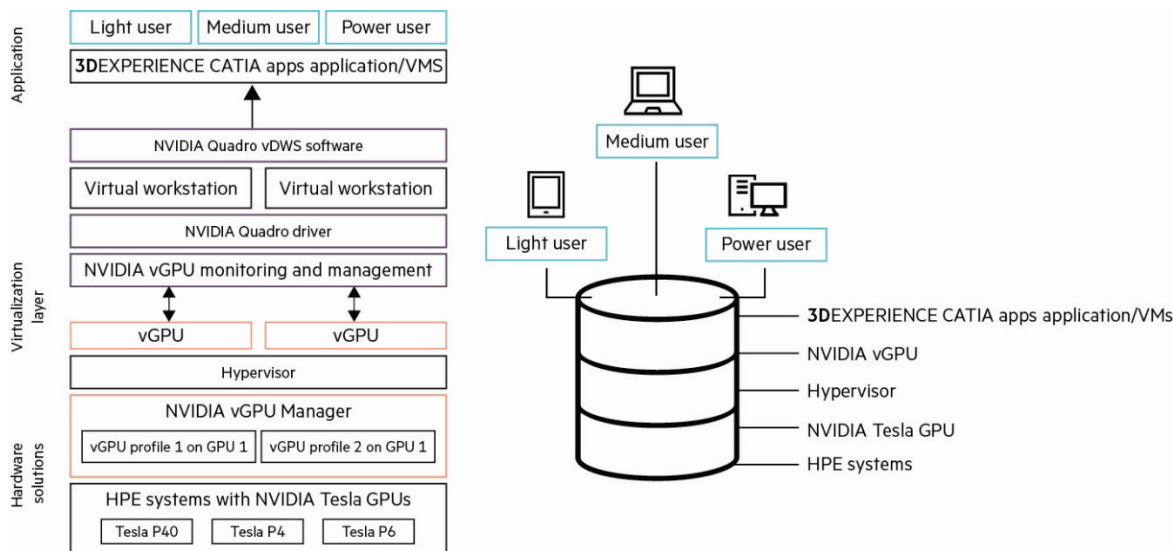


Figure 2. 3DEXPERIENCE CATIA apps, HPE, and NVIDIA solution components

The architecture consists of the application layer (**3DEXPERIENCE CATIA** apps), virtualization layer (GPU monitoring and management software, hypervisor), and systems layer (HPE systems with GPUs and associated systems software). A short description of elements of each of these layers is as follow:

User types

The user population is a combination of light, medium, and power users. The following table details the characteristics of each of these users.

Table 1. 3DEXPERIENCE CATIA apps user classification

User	Characteristics
Light user/knowledge worker	Predominantly uses professional applications (for example, Word, Excel); uses 3DEXPERIENCE CATIA apps for view-only or accessing individual parts or small assemblies and manufacturing shop floor.
Medium user/junior designer	3DEXPERIENCE CATIA apps user for view-only or accessing medium to large assemblies and manufacturing shop floor.
Power user/senior designer	3DEXPERIENCE CATIA apps user working with large assemblies or full model. They are real-time interactive simulations and iterative design and renders, requiring high performance.

NVIDIA vGPU layer

In a VDI environment powered by NVIDIA virtual GPUs, the Quadro vDWS¹ software is installed at the virtualization layer along with the hypervisor. This software creates virtual GPUs that enable every VM to share the physical GPU installed on the server. The software can also provision the use of a whole GPU or aggregate up to four GPUs to be accessed by a single virtual machine. The NVIDIA virtualization software includes a graphics driver for every VM.

Benefits of NVIDIA virtual GPU for 3DEXPERIENCE CATIA apps

An integral part of the HPE and **3DEXPERIENCE CATIA** apps solution is NVIDIA virtual GPU (vGPU) technology. NVIDIA vGPU solutions:

- **Improve productivity:** Deliver superior graphics performance for virtual desktops with the same responsive experience as a physical workstation.
- **Enhance collaboration anywhere, on any device:** Speed the design process by liberating employees to use the device of their choice to access fully 3D-capable virtual workstations from anywhere.
- **Protect IP:** Mitigate business from lost or stolen IP. Onboard new contractors in minutes while ensuring the security of project files in the data center.

¹ [nvidia.com/en-us/design-visualization/quadro-vdws/](https://www.nvidia.com/en-us/design-visualization/quadro-vdws/)

- **Drive data consistency:** Centralize the data center for greater consistency and consolidation of data as well as control over design changes.
- **Enable data center agility:** Live migration of GPU-enabled virtual machines (VMs) without impacting users or requiring scheduled downtime frees up IT resources to focus on more strategic projects.

HPE systems layer

HPE systems seamlessly integrate GPU computing. The GPUs along with the GPU software and the virtualization software (for example, VMware vSphere®, Citrix XenServer) reside in the HPE system. HPE system management software and tools assist in managing, administering, and monitoring these systems. HPE provides several system options in racks and blades.

3DEXPERIENCE CATIA is certified on some HPE servers which are outlined later in this paper. HPE certified servers deliver:

- **High performance:** Offers scalable high-performance systems that can be used for larger and more complex models and systems.
- **Reliability:** Minimizes downtime and unplanned outages and improves system availability for engineers and time to market.
- **Flexibility:** Provides choice of systems with a range of price points, configurations, and options that can be deployed in an individual department or in the data center.
- **Agility:** Automates server management and reconfigures on the fly to handle other workloads such as computer-aided engineering (CAE) and Internet of Things (IoT) applications to improve system utilization.
- **Efficiency:** Provides more performance in a smaller footprint, reduces data center floor space requirements, and lowers energy costs.
- **Advisory services:** Customizes solutions tailored to customer requirements, end-to-end system integration, and faster implementation of 3DEXPERIENCE CATIA apps solutions.
- **Partnerships:** Strong partnerships with technology leaders such as Dassault Systèmes, NVIDIA, VMware, and Citrix with deep expertise in their products and applications.

Solution components

Various solution components shown in the architectural diagram are detailed in the following.

HPE systems

Three HPE systems shown in Figure 3 are in the scope of this white paper.



Figure 3. HPE systems for 3DEXPERIENCE CATIA apps solution

Racks: HPE Apollo 2000 Gen10 system² is a density-optimized, 2U-shared infrastructure chassis for up to two HPE ProLiant XL190r Gen10 independent, hot-plug servers with all the traditional data center attributes—standard racks and cabling and rear-aisle serviceability access. With HPE Apollo 2000 Gen10 system servers, there is flexibility to tailor the system to the precise needs of each 3DEXPERIENCE CATIA apps workload.

² HPE Apollo 2000 Gen10 Chassis QuickSpecs and 2U for GPU server HPE ProLiant XL190r Gen10 Server QuickSpecs



HPE ProLiant DL380 Gen10 server³ is the latest addition to the successful HPE ProLiant DL380 product line of servers that offer numerous processor and connectivity options and is designed to reduce costs and complexity. The family offers a 71% performance gain and 27% increase in cores over Intel®'s previous generation Xeon E5 v4 family of processors, plus the HPE 2666 MT/s DDR4 SmartMemory supporting 3 TB, offering a broad range of options for **3DEXPERIENCE CATIA** apps.⁴

Blades: HPE Synergy 480 Gen10 compute module⁵ delivers capacity, efficiency, and flexibility in a 2-socket, half-height form factor to support demanding workloads. HPE Synergy 480 Gen10 compute module is an excellent platform for **3DEXPERIENCE CATIA** apps and other enterprise-grade workloads.

NVIDIA GPUs

The following NVIDIA GPUs are appropriate for **3DEXPERIENCE CATIA** apps vGPU VDI solutions.

Table 2. NVIDIA GPU comparison

Features	Tesla P40	Tesla T4	Tesla P4	Tesla P6
GPUs/board (architecture)	1 (Pascal)	1 (Turing)	1 (Pascal)	1 (Pascal)
CUDA cores	3840	2560	2560	2048
Memory size	24 GB GDDR5	16 GB GDDR5	8 GB GDDR5	16 GB GDDR5
vGPU profiles	2 GB, 3 GB, 4 GB, 6 GB, 8 GB, 12 GB, 24 GB	2 GB, 4 GB, 8 GB, 16 GB	2 GB, 4 GB, 8 GB	2 GB, 4 GB, 8 GB, 16 GB
Form factor	PCIe 3.0 dual slot (rack servers)	PCIe 3.0 single slot (rack servers)	PCIe 3.0 single slot (rack servers)	MXM (blade servers)
Power	250W	70W	75W	90W
Thermal	Passive	Passive	Passive	Bare board
Characteristics		Performance optimized		Blade optimized

- Tesla P40 is a 24 GB video memory size that is a good fit for midlevel to high-end 4 GB memory video profile sessions for professional graphics users. It includes use for single-precision compute workloads (for example, rendering and creating complex designs). The optimum fit is six users per card. The card can be used for 12 light users with 2 GB video profiles.
- Tesla P4 is an 8 GB video memory size that is a good fit for 2 midlevel to high-end 4 GB memory profiles or four medium 2 GB profile graphics users.
- Tesla T4 is an ideal upgrade path for users of Tesla P4, with 2X the framebuffer and added Tensor Cores, in the same energy-efficient, low-profile form factor.
- Tesla P6 is 16 GB video profile that is a good fit for four midlevel to high-end 4 GB memory profiles or eight light users with 2 GB video profiles.
- Tesla vGPUs can function with 1 GB vGPU user profile; however, **3DEXPERIENCE CATIA** apps is not supported with this user profile.

The configuration details and maximum GPU cards per node⁶ for the three selected HPE servers are provided in the following table.

Table 3. Maximum GPU cards per node for HPE systems

HPE server	Rack units	Node/chassis	Tesla P40	Tesla T4	Tesla P4	Tesla P6
HPE Apollo 2000 Gen10	2U	2	2 per node (4 per chassis)	2 per node (4 per chassis)	2 per node (4 per chassis)	N/A
HPE ProLiant DL380 Gen10	2U	1	3	5	5	N/A
HPE Synergy 480 Gen10	10U	12 (1 per node without expansion) 6 (with expansion chassis)	2 (2 per node with PCI expansion chassis leads to 12 per chassis)	N/A	N/A	1 (1 per node without expansion) 7 (7 per node with MXM expansion chassis leads to 42 per chassis)

³ HPE ProLiant DL380 Gen10 Server QuickSpecs; hpe.com/us/en/product-catalog/servers/proliant-servers/pip.hpe-proliant-dl380-gen10-server.1010026818.html

⁴ HPE ProLiant DL380 Gen10 4114 1P 32GB-R P408i-a 8SFF 800W PS Performance Server (P06421-B21)

⁵ HPE Synergy 480 Gen10 Compute Module QuickSpecs; Graphics accelerator options: HPE Synergy Graphics Accelerator Options Gen10 QuickSpecs

⁶ nvidia.com/object/grid-certified-servers.html



Designing the solution for 3DEXPERIENCE CATIA apps

3DEXPERIENCE CATIA apps leverages both the CPU and GPU to deliver a high-quality user experience. There are several considerations while architecting the Quadro vDWS solution. These include the size of user data, the concurrency of users, and the level of interaction with 3D data. In this section, we address each of these decision points and the pros and cons of each alternative.

User characteristics

The characteristics and composition of users using 3DEXPERIENCE CATIA apps on HPE with NVIDIA GPUs have a big bearing on sizing the system.

Here, we have assumed the users to be creative technical professionals who are extensive users of 3DEXPERIENCE CATIA apps Solids and other heavy-duty graphics applications. These user types are different from knowledge workers who use office productivity applications, which are also becoming increasingly graphics intensive.

Within this 3DEXPERIENCE CATIA apps user set, users are classified as power, medium, and light users based on their usage characteristics. The intent is to provide power users with more computing resources compared to the other two categories.

HPE server type selection

Three types of HPE systems are considered: HPE Apollo 2000 Gen10, HPE ProLiant DL380 Gen10, and HPE Synergy blades. All the three systems lines are well established, proven, and have a significant install base. A high-level comparison in an HPE environment is shown in Table 4.

Table 4. HPE server comparison

Server type	HPE Apollo 2000 Gen10	HPE ProLiant DL380 Gen10	HPE Synergy blades
Characteristics	<ul style="list-style-type: none"> Dense rack mounted (2x DL380 but with less configurability) Low TCO Same form factor for HPC workloads 	<ul style="list-style-type: none"> Rack mounted Low TCO Industry standard World's bestselling server⁷ Lots of options and configurability 	<ul style="list-style-type: none"> Software-defined composable infrastructure Blade format Network integrated into the chassis Highest density for NVIDIA vGPU 4Q profiles in a 10U form factor Highest thermal design power (TDP) processor choice for maximum performance

Preferences of centralized/decentralized IT, legacy installations, and the comfort level of IT management also play a major role in the selection of the system.

Select the right GPU for the application

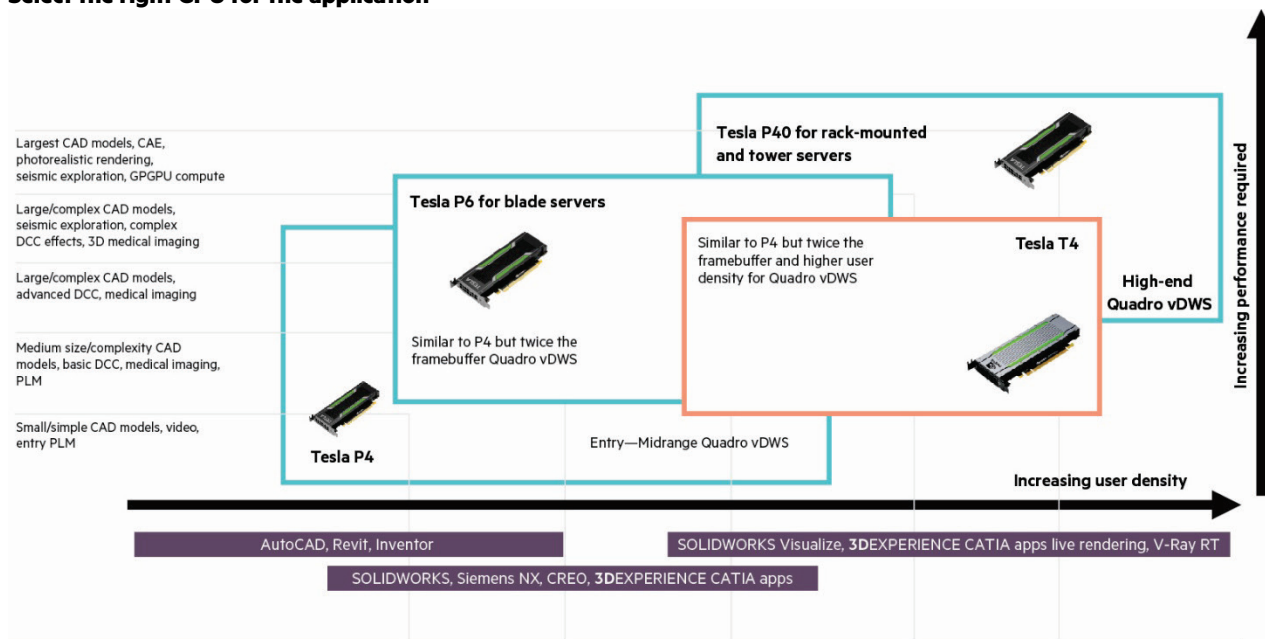


Figure 4. GPU selection guidance

⁷ With one HPE ProLiant server shipping every 14.6 seconds and more than 39.3 million sold, HPE ProLiant servers are the undisputed market share leader. CQ3'16 IDC Server Tracker; hpe.com/us/en/product-catalog/servers/proliant-servers/pip/hpe-proliant-dl380-gen10-server.1010026818.html

Guidance for GPU selection depending on the application, workload, and the type of server (that is, blade or rack) is provided in Figure 4.

NVIDIA has two types of quality of service (QoS) scheduling for server environments (non-cloud):⁸

- **Best Effort Scheduler:** Maximum utilization of GPU cycles, best potential performance, idle cycle is redistributed, no guarantee of QoS, no noisy neighbor protection
- **Equal Share Scheduler:** For compute workloads delivering guaranteed QoS, idle cycle is not redistributed, noisy neighbor protection

The choice of the scheduling depends on the use case. Customers should always start with Best Effort and if the POC reveals performance or QoS issues, equal share might be the answer.

Quadro vDWS is the ideal choice for professional graphics applications such as **3DEXPERIENCE CATIA** apps. The features of Quadro vDWS are shown in Table 5.

Table 5. Quadro vDWS features

Configuration and deployment	Supported
Desktop virtualization	✓
Remote desktop session host (RDSH) app hosting	✓
RDSH desktop hosting	✓
Windows® OS support	✓
Linux® OS support	✓
GPU pass-through support	✓
Bare-metal support	✓
Management and monitoring	✓
Live migration support	✓
Multi-GPU support	✓

Select the customer operating system

There are two choices VM OS: Windows 7 and Windows 10. The selection of the server OS depends on the customer's IT environment.

VDI environment for 3DEXPERIENCE CATIA apps on HPE servers

Citrix XenDesktop® or VMware Horizon View are the two preferred VDI environments for **3DEXPERIENCE CATIA** apps on HPE servers. The choice between the two depends on the customer's knowledge and ability to deploy and manage.

Virtualization hypervisor for 3DEXPERIENCE CATIA apps on HPE servers

VMware vSphere is the most prevalent virtualization hypervisor used in 75%⁹ of the implementations. Other customers are using Citrix XenServer.

IT can use VMware vMotion technology for live migration of NVIDIA vGPU-powered virtual desktops without end-user interruption or data loss. VMware vMotion is supported with NVIDIA virtual GPU software (vGPU 7.0) release.

Citrix XenMotion enables IT to migrate live NVIDIA GPU-accelerated VMs with no impact to the user and no loss of data. Citrix XenMotion is supported with NVIDIA virtual GPU software March 2018 release (aka vGPU 6.0) and identical NVIDIA Maxwell, Pascal, or Turing GPUs. Citrix XenServer 7.4 or later is also required.

⁸ "Optimizing NVIDIA virtual GPU for the best VDI user experience," NVIDIA, 2018

⁹ computerprofile.com/analytics-papers/vmware-far-largest-server-virtualisation-market/



3DEXPERIENCE CATIA apps certification on HPE servers

Dassault Systèmes certifies the configurations for 3DEXPERIENCE CATIA apps on a regular basis and the following table shows the recent certification status. Check with NVIDIA/HPE technical personnel to verify the certification status of a given configuration or check the Dassault Systèmes website with the list of Certified Configurations for VDI.¹⁰

Table 6. HPE systems and certified configurations for VDI

Certification status	Model	Processor and clock speed	Graphics card	Hypervisor	Desktop virtualization	Guest VM	vGPU driver
Certified since 10/2018	HPE WS460c Gen9	2 x Intel® Xeon® CPU E5-2690 v4 @2.60GHz	NVIDIA Quadro vDWS software, 4 x Tesla M6, M6-4Q	Citrix XenServer 7.1 Hypervisor driver: 410.68	Citrix XenDesktop 7.15 LTSR	Windows 7 64-bit, 8 VM, 4 vCPU GPU utilization 8:4	Driver: 411.81
Certified since 10/2018	HPE Synergy 480 Gen10	2 x Intel Xeon CPU Xeon Gold 6154 @3.00GHz	NVIDIA Quadro vDWS software, 4 x Tesla P6, P6-4Q	VMware ESXi™ 6.5 update 2 Hypervisor driver: 410.68	Citrix XenDesktop 7.15 LTSR	Windows 7 64-bit, 8 VM, 4 vCPU GPU utilization 8:4	Driver: 385.90
Certified since 09/2018	HPE WS460c Gen9	2 x Intel Xeon CPU E5-2699 v4 @3.20GHz	6x NVIDIA Quadro M3000 SE Pass-through	Citrix XenServer 7.1 Hypervisor driver: N/A	Citrix XenDesktop 7.15 LTSR	Windows 7 64-bit, 6 VM, 4 vCPU GPU utilization 6:6 Driver: 391.03	Driver: N/A
Certified since 05/2017	HPE Apollo 2000 Gen9	2 x Intel Xeon CPU E5-2667 v4 @3.20GHz	NVIDIA Quadro vDWS software, 1 x Tesla M60, M60-2Q	VMware ESXi 6.0 update 2 Hypervisor driver: 384.99	VMware Horizon® 7	Windows 7 64-bit, 8 VM, 2 vCPU GPU utilization 8:1	Driver: 385.90
Certified since 02/2019	HPE Apollo 2000 Gen10	2 x Intel Xeon CPU Xeon Gold 6136 @3.00GHz	NVIDIA Quadro vDWS software, 2 x P40, P40-4Q	Citrix XenServer 7.1 Hypervisor driver: 410.68	Citrix XenDesktop 7.15 LTSR	Windows 10 64-bit, 12 VM, 4 vCPU GPU utilization 12:2	Driver: 411.81
Certified since 03/2019	HPE Apollo 2000 Gen10	2 x Intel Xeon CPU Xeon Gold 6136 @3.00GHz	NVIDIA Quadro vDWS software, 1 x P40, P40-2Q	VMware ESXi 6.5 update 2 Hypervisor driver: 410.68	VMware Horizon 7	Windows 10 64-bit, 12 VM, 2 vCPU GPU utilization 12:1	Driver: 411.81
Certified since 03/2019	HPE DL380 Gen10	2 x Intel Xeon CPU Xeon Gold 6136 @3.00GHz	NVIDIA Quadro vDWS software, 2 x P40, P40-2Q	VMware ESXi 6.5 update 2 Hypervisor driver: 410.68	VMware Horizon 7	Windows 10 64-bit, 24 VM, 2 vCPU GPU utilization 24:2	Driver: 411.81

Sizing considerations

Sizing of systems for complex graphics applications such as 3DEXPERIENCE CATIA apps must be highly customized to a customer's needs and unique situation. Due to the varying nature of product complexity, sizing in one case may not apply to the other. As an example, an automobile may have 25,000 parts whereas an airliner may have over 2 million parts. As such, the system requirements for a junior designer in the aerospace industry will be much different for a similar designer in the automobile industry.

The following best practices are based on discussions with many 3DEXPERIENCE CATIA apps solution SMEs. **A common underlying message from all these discussions is that a proof of concept (POC) is an important element in sizing the system.**

Here are few suggestions for sizing the systems to be used in concert with earlier items discussed in this section.

- For a complex application such as 3DEXPERIENCE CATIA apps, the clock cycle of the CPU is very important: 3.0 GHz or above is preferred.
- Large framebuffer is essential for power users. In general:
 - 45% of users need 2 GB framebuffer.
 - 50% of users need 2 GB to 4 GB framebuffer.
 - 5% of users need greater than 8 GB framebuffer.

¹⁰ media.3ds.com/support/certified_hardware/Windows_VDI_implement.html



- GPU sizing
 - Increasing users per GPU will reduce performance.
 - Framebuffer size is **not** a good measure of expected performance.
 - Better performance per density achieved using more GPUs (Tesla P4 or T4) with a number of users versus fewer larger GPUs (P40) with similar number of users.
 - NVIDIA recommends maximum 4 to 6 users per GPU as a starting point for a POC.
 - For schedule management of GPUs **Best Effort** is preferable compared to **Equal Share**, as there will be less idle GPU compute cycles. NVIDIA recommends running a POC and measuring utilization of GPU, CPU, and others while gathering feedback from end users and then optimize/change based on the feedback. If customers need a specific performance at all times (QoS optimized), they should use the equal share scheduler and can run benchmarks like SPECviewperf.
- Sizing suggestions based on usage characteristics shown in the following table:

Table 7. Sizing based on user classification

Job description	Physical CPU cores	RAM	GPU framebuffer
Light user/knowledge worker	0.5	6 GB to 8 GB	2 GB
Medium user/junior designer	2	16 GB to 32 GB	3 or 6 GB
Power user/senior designer	4	32 GB to 64 GB	6 GB
Ultra-power	> 4	32 GB	Full GPU and up to 4 GPUs

3DEXPERIENCE CATIA apps implementation: Customer example

By providing a range of innovative platforms, certified solutions, and high-value services for 3DEXPERIENCE CATIA apps, HPE, and NVIDIA are driving and accelerating the digital transformation of manufacturing customers worldwide. Here is an example.

Honda R&D implementation: Global transportation equipment manufacturer

For more information, see the full [Honda R&D case study](#).

Typical manufacturing customer deployments of 3DEXPERIENCE CATIA apps in a VDI environment fit into the following three segments, depending on whether you have light, medium, or power users.

Table 8. Manufacturing customer deployments by user

Light users	Medium users	Power users
Users per server: 16–24	Users per server: 12–18	Users per server: 6–9
vGPU profile size: 2Q 4 vCPU	vGPU profile sizes: 2Q, 4Q 8 vCPU	vGPU profile sizes: 8Q, 12Q 12 vCPU
8 GB–16 GB RAM	16 GB–32 GB RAM	> 96 GB RAM

For more information, see the [NVIDIA Quadro vDWS Sizing Guide for CATIA](#).



HPE advantage for 3DEXPERIENCE CATIA apps in manufacturing

As more manufacturers embark on their VDI/3DEXPERIENCE CATIA apps journey, they need a reliable partner with deep design and manufacturing expertise along with knowledge of 3DEXPERIENCE CATIA apps. Working jointly with NVIDIA, HPE provides a comprehensive portfolio of enterprise-grade, high-performance systems and software, and high-value services to help manufacturing customers implement 3DEXPERIENCE CATIA apps and VDI throughout the value chain and product lifecycle.

Worldwide, many manufacturing companies are already using CAE, IoT, and other enterprise solutions from HPE. As 3DEXPERIENCE CATIA apps becomes an integral part of end-to-end operations in manufacturing, HPE is well positioned to help manufacturers innovate, improve product and service quality, and optimize operations to achieve new levels of profits and performance.

As a market leader in manufacturing solutions, HPE delivers enterprise-grade high-performance systems that are efficient, scalable, and cover a wide range of form factors satisfying the needs of manufacturing and design companies small and large. The HPE portfolio consisting of a wide array of rack mounted and blade servers work in concert with NVIDIA GPUs to exploit the full functionality of 3DEXPERIENCE CATIA apps. Many of these solutions are already certified by Dassault Systèmes. Both HPE and NVIDIA continue to invest in the highly skilled resources needed to expand this list of certified solutions so that manufacturing customers benefit from unparalleled productivity, performance, IP protection, and flexibility.

HPE and NVIDIA resources

[HPE Apollo 2000 Gen10 System](#)

[HPE ProLiant DL380 Gen10 Server](#)

[HPE Synergy 480 Gen10 Compute Module](#)


[NVIDIA and CATIA](#)

[NVIDIA vGPU Technology](#)

[NVIDIA vGPU evaluation license for up to 90 days](#)

Learn more at

hpe.com/software-solutions/hpe-nvidia-vgpu



 **Share now**

 **Get updates**

© Copyright 2019 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 paper was prepared by HPE and NVIDIA only. Neither Dassault Systèmes nor any of its subsidiaries makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or use of this paper, or of any information disclosed therein.

3DEXPERIENCE, the Compass logo and the 3DS logo, CATIA, SOLIDWORKS, ENOVIA, DELMIA, SIMULIA, GEOVIA, EXALEAD, 3D VIA, BIOVIA, NETVIBES, IFWE and 3DEXCITE are commercial trademarks or registered trademarks of Dassault Systèmes, a French "société européenne" (Versailles Commercial Register # B 322 306 440), or its subsidiaries in the U.S. and/or other countries. Use of any Dassault Systèmes or its subsidiaries trademarks is subject to their express written approval. Intel and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries. Windows is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. Citrix and XenDesktop are registered trademarks 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. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. VMware, VMware vSphere, VMware vSphere vMotion, VMware Horizon 7, and VMware ESXi are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. All other third-party marks are property of their respective owners.