



The bridge to possible

White Paper
Public

Size Dassault SOLIDWORKS on Cisco UCS and Cisco HyperFlex for Virtualized Workstations

Contents

Executive summary	3
Solution overview	3
Features and benefits	4
Solution architecture and components	5
NVIDIA GPUs and software for virtualization	9
Sizing considerations	13
SOLIDWORKS sizing recommendations	17
Ordering information	20
Conclusion	20
Cisco Customer Experience Services	20
Cisco Capital flexible payment solutions to help you achieve your objectives	21
For more information	21

Executive summary

In many organizations, Dassault Systèmes SOLIDWORKS practitioners use physical workstation computers that are expensive to purchase and maintain. Recent advances in the Cisco Unified Computing System™ (Cisco UCS®) and Cisco HyperFlex™ virtualization platforms and important developments in NVIDIA graphics processing units (GPUs) and software make it possible to virtualize professional graphic workstations running SOLIDWORKS. This document discusses the organizational benefits of virtualization and provides initial guidance for sizing SOLIDWORKS deployments.

Solution overview

Dassault SOLIDWORKS delivers a powerful suite of tools that provides a completely integrated design to manufacturing solutions. To support SOLIDWORKS in a virtual environment, different user types in an organization need to be classified based on their roles.

Each user type's virtual workstation requirements will vary. Cisco recommends that organizations run a proof-of-concept study using their SOLIDWORKS implementation to better define each of their user types.

The foundation on which a virtual graphic workstation is built consists of the server and storage system that support it. Two server and storage architectures can be used to support virtual workstations: converged infrastructure and hyperconverged infrastructure (HCI). Cisco supports both foundations, giving customers several choices. Customers can choose from our industry-leading FlashStack or FlexPod converged infrastructure joint offerings with Pure Storage or NetApp respectively, or from our innovative hyperconverged platform: the Cisco HyperFlex system.

All of our solution platform offerings incorporate NVIDIA GPUs and NVIDIA RTX Virtual Workstation (vWS) software, which enables each virtual workstation to access one or more full GPUs as native PCI devices.

The solution is virtualized on the VMware ESXi hypervisor to use its industry-leading support for the capabilities of virtualized GPUs.

Citrix Virtual Apps and Desktops and VMware Horizon desktop brokers are supported for this solution.

The SOLIDWORKS application runs on each virtual workstation built on a Microsoft Windows 10 64-bit desktop operating system.

Figure 1 illustrates the interaction of the components that deliver high-performance virtual solutions.

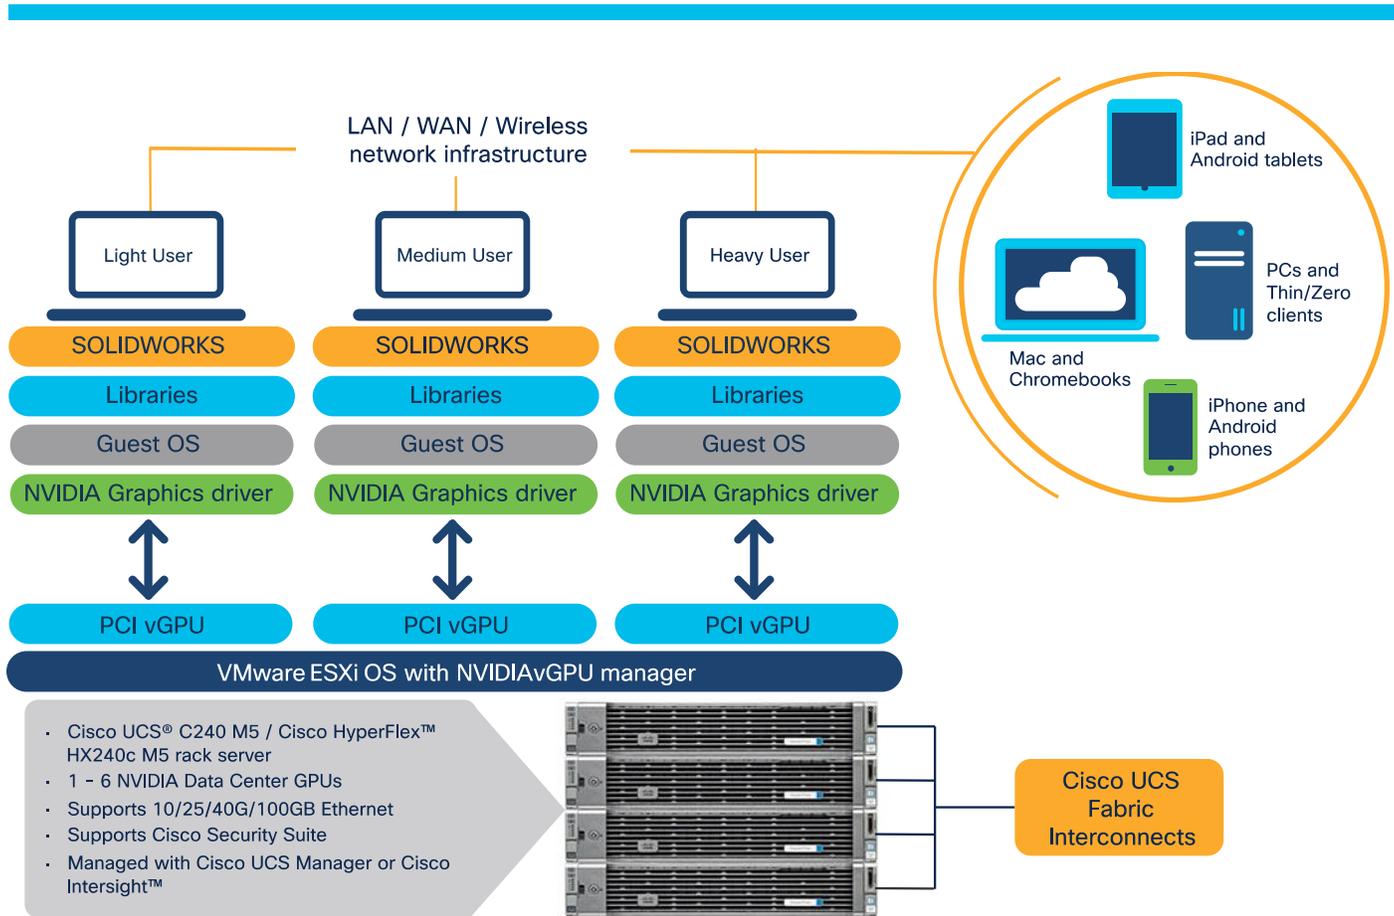


Figure 1.
Dassault SOLIDWORKS virtual workstation solution logical diagram
LAN, WAN, and wireless network infrastructure

Features and benefits

Table 1 summarizes the main advantages of running SOLIDWORKS on a virtual workstation.

Table 1. Features and benefits of deploying Dassault SOLIDWORKS on virtual graphic workstations

Feature	Benefit
Global collaboration	<ul style="list-style-type: none"> Eliminate geographic restrictions that prevent physical graphic workstation users from working together on large files. Enable organizations to deploy their best workers to solve their biggest challenges in real time regardless of a worker's location. Enable teams to work on projects at any time to accelerate completion. Enable workers to work from anywhere on any supported device if security policy permits.
Enhanced productivity	<ul style="list-style-type: none"> Essentially eliminate wait times for large file downloads because the data and the virtual workstation are co-located in the data center. Reduce file download and upload wait times for engineering professionals by an hour or more per day. Achieve faster time to productivity. Enable bring-your-own-device (BYOD) operation. Enable engineering professionals to access their work from multiple device types. Provide private cloud access from anywhere.
More secure intellectual	<ul style="list-style-type: none"> Intellectual property does not leave the corporate data center.

Feature	Benefit
property	<ul style="list-style-type: none"> • Critical files cannot be damaged during download or upload. • Latest versions of designs are always stored in the data center. • Intellectual property backups can be as detailed as by the minute. • Policies can be set to both authenticate and authorize various levels of file access by users. • Policies can be set to prevent downloading, copying, or deletion of any file in the data center.
Reduced capital expenses	<ul style="list-style-type: none"> • Eliminate the need to purchase expensive physical graphic workstations. • Consolidate and enhance design computing resources in fewer powerful graphic workstation server hosts. • Extend the life of the design infrastructure by maintaining it in a controlled, secured data center environment. • Deploy much lower-cost thin or zero clients to replace physical graphic workstations.
Reduced operating expenses	<ul style="list-style-type: none"> • Eliminate expensive maintenance contracts for physical graphic workstations. • Eliminate the need to install, upgrade, and maintain desktop operating systems on physical graphic workstations. • Eliminate installation of SOLIDWORKS and other required applications on physical graphic workstations. • Maintain master images for key user types and deploy updates to applications and operating systems at next login. • Deploy patches, updates, security fixes, and antivirus updates with little or no user disruption.
Excellent application performance	<ul style="list-style-type: none"> • Engineering resources experience physical workstation-like end user experience. • Engineering resources are not tethered to physical devices. • SOLIDWORKS 3D models are displayed at the highest resolution supported by endpoints.

Solution architecture and components

Cisco brings a robust, secure, hybrid cloud approach to virtual desktop infrastructure (VDI) and virtual graphic workstations, as shown in Figure 2. This wholistic approach allows customers to enable remote work for all functions in its workforce. In addition, if security policy allows, the solution enables work from anywhere on any device.

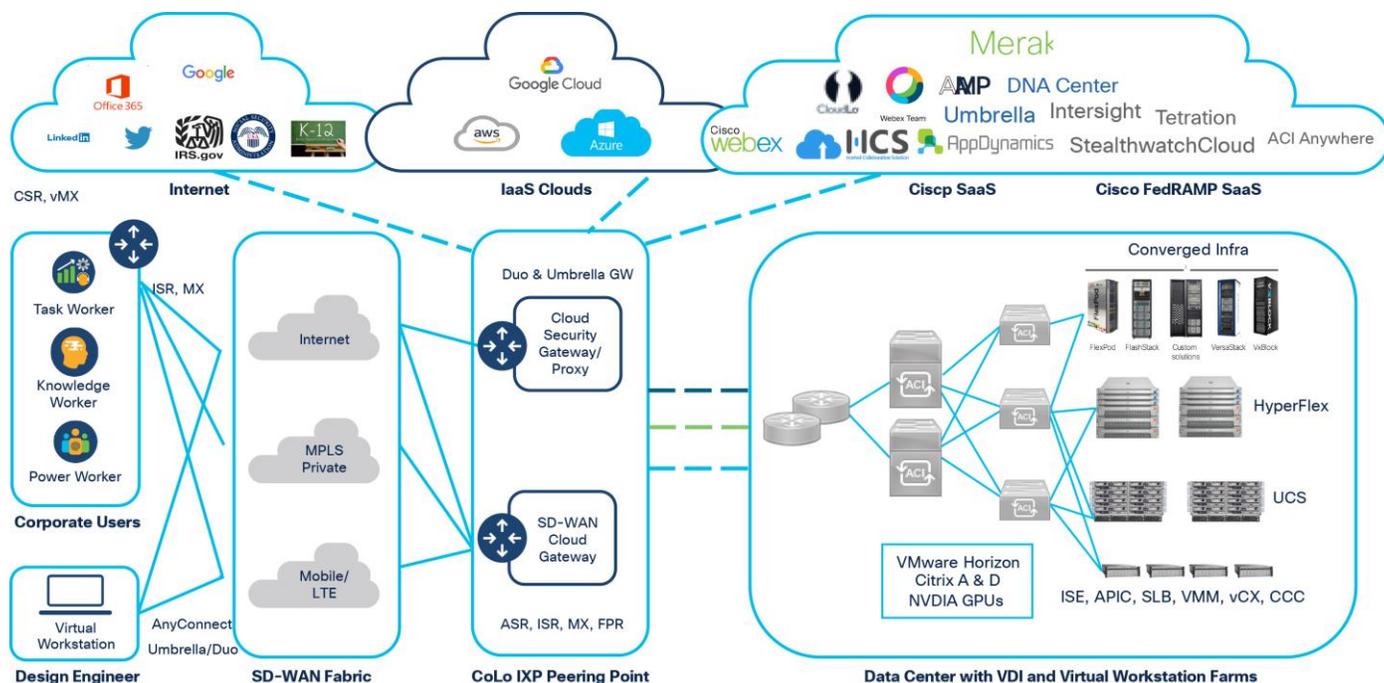


Figure 2.
Cisco solution architecture

SOLIDWORKS designs can be hosted on either of two types of infrastructure that support virtual workstations: converged infrastructure or hyperconverged infrastructure. Each has its unique sets of features and benefits.

Converged infrastructure

Built on Cisco UCS, converged infrastructure relies on Cisco® innovations introduced in 2009 that dramatically changed the data center server market. With Cisco UCS, we created a completely stateless software-defined storage network and computing platform that is defined by subject-matter experts (SMEs) in each field. It has evolved significantly over the past decade and now extends to the cloud with the Cisco Intersight™ platform.

Cisco UCS Manager software runs on specialized Cisco switches called fabric interconnects. The fabric interconnects are deployed in pairs for fault tolerance. Each fabric interconnect pair defines a Cisco UCS domain, which is a management boundary. Multiple Cisco UCS domains can be managed by the central Cisco UCS Manager, which deploys and enforces policy and character globally.

The customer organization’s SMEs define policy for their areas of expertise. Policies defined by storage, network, and computing SMEs are combined into a Cisco UCS service profile. When a Cisco UCS server is connected to the fabric interconnects, it is inventoried automatically. Based on policy in the service profiles that exist in the Cisco UCS domain, the connected server may have its service profile assigned automatically with no administrator intervention.

For example, if a server is connected to a Cisco UCS domain with SOLIDWORKS virtual workstation host service profiles with appropriate identifiers, such as a certain processor, a certain memory configuration, and a particular NVIDIA GPU, then the server can be associated with the SOLIDWORKS service profile automatically.

Alternatively, an administrator can select the service profile for recently discovered servers from the Cisco UCS Manager HTML 5 user interface.

In a converged infrastructure, the service profile contains information about how the server connects to the partner storage array to access boot logical unit numbers (LUNs) and file storage repositories. The service profile also provides network connectivity definitions for Fibre Channel host bus adapters (HBAs), Small Computer System Interface over IP (iSCSI) initiators, and 10/25/40/100 Gigabit Ethernet interfaces through Cisco UCS virtual interface cards (VICs), which provide single-cable, multiprotocol connectivity.

The converged infrastructure is built using the building blocks listed in Table 2.

Table 2. Converged infrastructure components

Product family	Platforms supported	NVIDIA GPU support
Cisco UCS blade servers	FlashStack and FlexPod	NVIDIA P6 front and NVIDIA P6 rear
Cisco UCS 5108 Blade Server Chassis	FlashStack and FlexPod	
Cisco UCS rack-mount servers	FlashStack and FlexPod	NVIDIA T4, NVIDIA RTX 6000, and NVIDIA RTX 8000
Cisco UCS 6154 and 64108 Fabric Interconnects	FlashStack and FlexPod	
Cisco UCS Manager, Cisco UCS Central Software, and Cisco Intersight	FlashStack and FlexPod	All NVIDIA Pascal, Turing, and RTX server GPUs
Cisco Nexus® 9300 platform Layer 2 switches (Cisco Application Centric Infrastructure [Cisco ACI®] supported)	FlashStack and FlexPod	
NetApp All Flash and All NVME A200, A400, and A800 arrays	FlexPod	
Pure Storage FlashArray//m20, //m50, and //m70 arrays	FlashStack	
Cisco security products (Cisco Duo, Cisco Umbrella®, Cisco Adaptive Security Appliance [ASA], Cisco Tetration Analytics™, etc.)	FlashStack and FlexPod	
Cisco performance products (Cisco Workload Optimization Manager [CWOM], Cisco AppDynamics®, Cisco Stealthwatch® Cloud, etc.)	FlashStack and FlexPod	

Hyperconverged infrastructure

Cisco HyperFlex systems, built on Cisco UCS, combine computing, physical storage, and networking resources in a unified platform that is well suited to deliver both VDI and virtual workstations. To help you meet the challenge of deploying any application, in any cloud, anywhere, we introduced Cisco HyperFlex systems, an adaptive system to power any application anywhere with the simplicity of hyperconvergence. Cisco HyperFlex systems deliver flexibility to support multiple hypervisors, containerized environments, multicloud services, and edge deployment to efficiently and economically deploy, monitor, and manage today's applications.

We designed Cisco HyperFlex systems as a next-generation platform that can adapt to meet new information technology challenges as they evolve. Cisco HyperFlex clusters are built with a minimum of three hyperconverged nodes and can support up to 64 nodes in a single cluster. (Limited-function edge clusters can be built with as few as two hyperconverged nodes.) A significant differentiator for Cisco is our Cisco HyperFlex platform that support computing-only nodes, in a ratio of two computing-only nodes to one hyperconverged node. The computing-only nodes contribute CPU, memory, and graphics capabilities and share the storage presented by the Cisco HyperFlex hyperconverged nodes. The computing-only nodes do not require the license costs for the Cisco HyperFlex HX Data Platform software. This factor presents a significant economic advantage for VDI and SOLIDWORKS users, whose CPU and memory demands outstrip storage requirements.

The Cisco HyperFlex HX Data Platform is a purpose-built, high-performance, scale-out file system with a wide array of enterprise-class data management services. The data platform's innovations redefine distributed storage technology, giving you complete hyperconvergence with enterprise storage features. Here is how the solution works.

An HX Data Platform controller resides on each node and implements a distributed file system. The controller runs in user space within a virtual machine and intercepts and handles all I/O from guest virtual machines. Dedicated CPU cores and memory allow the controller to deliver consistent performance without affecting the performance of the other virtual machines in the cluster. When nodes are configured with self-encrypting drives, the controller negotiates with Cisco UCS Manager to receive the encryption keys that enable the drives to encrypt and decrypt data that flows to and from the various storage layers.

The data platform has modules to support the specific hypervisor or container platform in use. The controller accesses all of the node's disk storage through hypervisor bypass mechanisms for excellent performance. It uses the node's memory and dedicated solid-state disk (SSD) drives or Non-Volatile Memory Express (NVMe) storage as part of a distributed caching layer, and it uses the node's hard-disk drives (HDDs), SSD drives, or NVMe storage for distributed storage. The data platform controller interfaces with the hypervisor in two ways:

- IO Visor: The data platform controller intercepts all I/O requests and routes them to the nodes responsible for storing or retrieving the blocks. IO Visor makes the existence of the hyperconvergence layer transparent to the hypervisor.
- Hypervisor agent: A module uses the hypervisor APIs to support advanced storage system operations such as snapshots and cloning. These are accessed through the hypervisor so that the hyperconvergence layer appears as if it were enterprise shared storage. The controller accelerates operations by manipulating metadata rather than performing actual data copying, providing rapid response and thus rapid deployment of new application environments.

Data distribution

The HX Data Platform controller handles all read and write requests for volumes that the hypervisor accesses and thus intermediates all I/O from the virtual machines and containers. Recognizing the importance of data distribution, the HX Data Platform is designed to exploit low network latencies and parallelism, in contrast to other approaches that build on node-local affinity and can easily cause data hot spots.

With data distribution, the data platform stripes data evenly across all nodes, with the number of data replicas determined by the policies you set.

This approach helps prevent both network and storage hot spots and makes I/O performance the same regardless of virtual machine location. This feature gives you more flexibility in workload placement and

contrasts with other architectures in which a data locality approach does not fully utilize all available networking and I/O resources.

- **Data write operations:** For write operations, data is written to the local SSD or NVMe cache, and the replicas are written to remote caches in parallel before the write operation is acknowledged. Write operations are later synchronously flushed to the capacity layer HDDs (for hybrid nodes) or SSD drives (for all-flash nodes) or NVMe storage (for NVMe nodes).
- **Data read operations:** For read operations in all-flash nodes, local and remote data is read directly from storage in the distributed capacity layer. For read operations in hybrid configurations, data that is local usually is read directly from the cache. This process allows the platform to use all solid-state storage for read operations, reducing bottlenecks and delivering excellent performance.

In addition, when a virtual machine migrates to a new location, the data platform does not require data movement because any virtual machine can read its data from any location. Thus, moving virtual machines has no performance impact or cost.

The hyperconverged infrastructure is built using the building blocks listed in Table 3.

Table 3. Hyperconverged infrastructure components

Product family	Platforms supported	NVIDIA GPU support
Cisco HyperFlex hyperconverged servers (hybrid, all flash, and all NVMe)	1 rack unit (1RU) and 2RU	NVIDIA T4, NVIDIA RTX 6000, and NVIDIA RTX 8000
Cisco HyperFlex computing-only servers	1RU, 2RU, and blade	NVIDIA P6, NVIDIA T4, NVIDIA RTX 6000, and NVIDIA RTX 8000
Cisco UCS 6154 and 64108 Fabric Interconnects	HCI and computing only	
Cisco UCS Manager, Cisco UCS Central Software, and Cisco Intersight platform	HCI and computing only	All NVIDIA Pascal, Turing, and RTX server GPUs
Cisco Nexus 9300 platform Layer 2 switches (Cisco ACI supported)	HCI and computing only	
Cisco security products (Cisco Duo, Cisco Umbrella, Cisco ASA, Cisco Tetration Analytics, etc.)	HCI and computing only	
Cisco performance products (Cisco CWOM, Cisco AppDynamics, Cisco Stealthwatch Cloud, etc.)	HCI and computing only	

NVIDIA GPUs and software for virtualization

NVIDIA virtual GPU (vGPU) software enables delivery of graphics-rich virtual desktops and workstations accelerated by NVIDIA GPUs, the most powerful data center GPUs on the market today. With NVIDIA vGPU software, GPU resources can be divided so that the GPUs are shared across multiple virtual machines, or multiple GPUs can be allocated to a single virtual machine to power the most demanding workflows. NVIDIA vGPU software runs on GPUs based on NVIDIA Turing, Volta, Pascal, Maxwell, and Ampere™ architectures.

For SOLIDWORKS, four NVIDIA GPUs are best suited:

- NVIDIA P6 GPU (Figure 3)
 - Designed for blade servers and best suited for light users
 - Delivers high graphics performance, improved energy efficiency, and increased user density with 16 GB of frame buffer
 - Mobile PCI Express (PCIe) module (MXM) form factor that runs on less than 90 watts (W) of power for high-density data centers
 - Up to two NVIDIA P6 cards per Cisco UCS B200 M5 blade server.

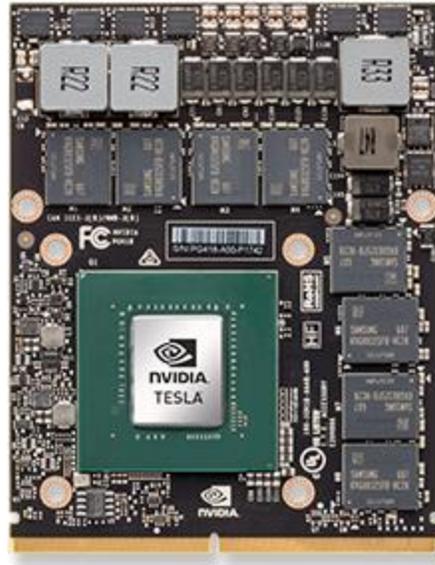


Figure 3.
NVIDIA P6 GPU

- NVIDIA T4 GPU (Figure 4)
 - With vWS software, it provides an excellent solution for light and medium 3D design and engineering workflows
 - With a single slot, low-profile form factor, and just 70 W of power consumption, achieves maximum GPU density per server
 - Up to six NVIDIA T4 cards per Cisco UCS C240 M5 or Cisco HyperFlex HX240c M5 rack server
 - Up to two NVIDIA T4 cards per Cisco UCS C220 M5 or Cisco HyperFlex HX220c M5 rack server

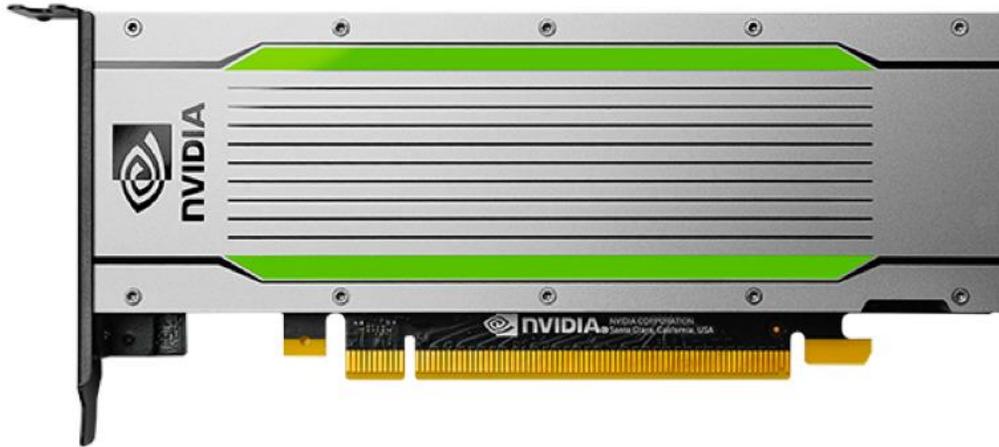


Figure 4.
NVIDIA T4 GPU

- NVIDIA RTX 6000 GPU (Figure 5)
 - Combined with vWS software, enables design engineers to work from high-powered virtual design workstations and render nodes to accelerate design workflows and arrive at their best creations faster
 - With ray-tracing (RT) cores, a larger frame buffer, and multiple profile sizes, gives engineers and designers the flexibility to run demanding workloads from the data center
 - Up to two NVIDIA RTX 6000 cards per Cisco UCS C240 M5 (July 2020) or Cisco HyperFlex HX240c M5 (Q4 2020) rack server



Figure 5.
NVIDIA RTX 6000 GPU

- NVIDIA RTX 8000 GPU (Figure 6)
 - Includes 48 GB of memory for the largest models—twice the frame buffer of the RTX 6000
 - Enables designers to work with the largest and most complex ray-tracing and visual-computing workloads

- Delivers the ultimate flexibility with vWS software, powering virtual design workstations and render nodes to propel creative workflows
- Up to two NVIDIA RTX 8000 cards per Cisco UCS C240 M5 (July 2020) or Cisco HyperFlex HX240c M5 (Q4 2020) rack server



Figure 6.
NVIDIA RTX 8000 GPU

NVIDIA RTX Virtual Workstation (vWS) software

From stunning industrial design to advanced special effects to complex scientific visualization, NVIDIA RTX is the world’s preeminent visual computing platform. And with NVIDIA RTX Virtual Workstation (vWS) software, you can now deliver the most powerful virtual workstation from the data center or cloud to any device, anywhere. Millions of creative and technical professionals can access the most demanding applications from any device, work from anywhere, and tackle larger data sets, all while meeting the need for greater security.

Virtual workstations free users from the confines of physical location, delivering resources from the data center and providing secure access on any device, anywhere. NVIDIA vWS extends the trusted benefits of NVIDIA RTX to deliver a true GPU-accelerated data center. This capability lets IT virtualize any application from the data center with a workstation-class user experience. Now your business can eliminate constrained workflows that inhibit agility, and users can securely collaborate in real time without borders or limits. You can efficiently centralize all your applications and data for dramatically lower IT operating expenses.

Table 4 lists the main features of vWS software. These features allow excellent flexibility for supporting important capabilities and selecting the virtual workstation operating system.

Table 4. NVIDIA vWS features

	RTX Virtual Workstation	Virtual PC	Virtual Compute Server
Desktop Virtualization	✓	✓	
Server Virtualization			✓
Windows OS Support	✓	✓	
Linux OS Support	✓	✓	✓
NVIDIA Graphics Driver	✓	✓	
NVIDIA RTX Enterprise Driver	✓		
NVIDIA Compute Driver			✓
Multi-vGPU/NVLink	✓		✓
ECC Reporting and Handling	✓		✓
Page Retirement	✓		✓

Display	RTX Virtual Workstation	Virtual PC	Virtual Compute Server
Maximum Hardware Rendered Display	Four 5K, Two 8K	Four QHD, Two 4K, One 5K	One 4K
Maximum Resolution	7680x4302	5120x2880	4096x2160

Advanced Professional Features	RTX Virtual Workstation	Virtual PC	Virtual Compute Server
ISV Certifications	✓		
NVIDIA CUDA/OpenCL	✓		✓

Graphics Features and APIs	RTX Virtual Workstation	Virtual PC	Virtual Compute Server
NVENC	✓	✓	✓
OpenGL Extensions (WebGL)	✓	✓	
Insitu Graphics/GL Support			✓
RTX Enterprise Driver Optimizations	✓		
DirectX	✓	✓	
Vulkan Support	✓		✓

Profiles	RTX Virtual Workstation	Virtual PC	Virtual Compute Server
Max Frame Buffer Supported	48GB	2GB	48GB
Available Profiles	0Q, 1Q, 2Q, 3Q, 4Q, 6Q, 8Q, 12Q, 16Q, 24Q, 32Q, 48Q	0B, 1B, 2B	4C, 6C, 8C, 12C, 16C, 24C, 32C, 48C

Sizing considerations

You need to consider four main factors when sizing a SOLIDWORKS deployment. Each is discussed here in detail.

- User types and requirements
- Display resolution and monitors
- CPU and memory selection
- GPU selection
- Scheduling engine
- Frame buffer and frame rates

User types and requirements

Grouping users by their SOLIDWORKS use pattern is the first step in a successful virtual graphics workstation deployment. Identify five main criteria for each group. Figure 7 shows the criteria.

User requirements

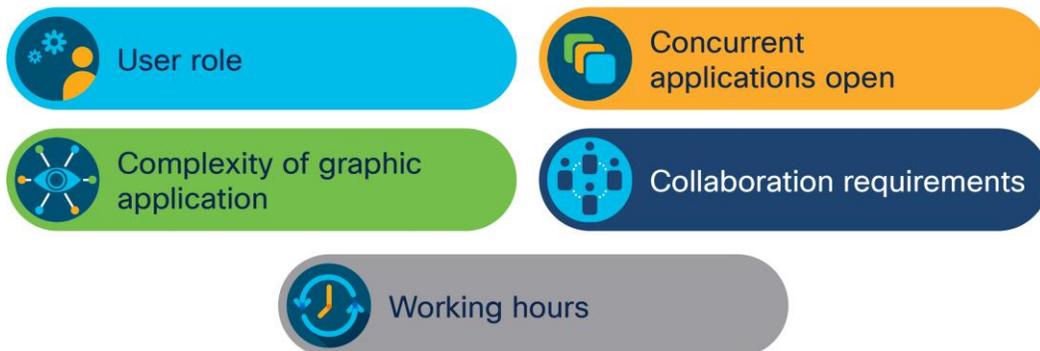


Figure 7.

User requirements

- User role: Identify the primary functions for SOLIDWORKS. Identify roles in terms of the intensity and purpose of use.
- Complexity: Think in terms of read-only users, editors, designers, and the size of the system being created or manipulated.
- Concurrency: How many applications does the user group member use at the same time? Are other graphics-intensive applications besides SOLIDWORKS being used?
- Collaboration: Does the user group have a collaboration requirement that spans geographic locations? Is the collaboration cross-functional? Across user roles?
- Working hours: Are your SOLIDWORKS users working across time zones? Does SOLIDWORKS user activity follow the sun—that is, does user activity occur at all times?

As a starting point, Cisco uses the classifications listed in Table 5.

Table 5. Starting-point recommendations for user classifications

User classification	Characteristics
Light	<ul style="list-style-type: none">• Read-only, for review and documentation• Project management• Small parts or subassembly
Medium	<ul style="list-style-type: none">• Read-only and full application• Medium assemblies
Heavy	<ul style="list-style-type: none">• Full application• Large assemblies or full product

For this document, virtual workstations will be constructed with the three user classifications listed in the table. As stated previously, Cisco recommends running a proof-of-concept study with your users to develop user classifications based on your specific use case.

Display resolution and monitor selection

Different SOLIDWORKS user types generally have different requirements for display resolution and the number and types of monitors. This information is crucial in sizing the virtual graphic workstations that will support each user type.

- Resolution: HDMI, 4K, 5K, or higher
- Size: 24, 32, or 40 inches, or larger
- Number: 1, 2, or 4

For this document, SPECviewperf 13 was used, with HDMI resolution, to measure performance and determine starting-point sizing recommendations for SOLIDWORKS.

CPU and memory selection

NVIDIA GPUs work synergistically with the Intel® Xeon® Scalable family of processors and high-frequency memory installed on the virtual graphics workstation host server. It is important to pair the proper CPU and memory with the chosen NVIDIA GPU for the user type and virtual machine density expected for the virtual graphic workstation hosts.

- CPUs: High-frequency (3.0 GHz or faster) processors such as like the Intel Xeon Scalable 6248R, 6246R, or 6242R are recommended.
- Memory: High-frequency (2933 MHz) memory in a balanced configuration (12 or 24 DIMMs) provide optimal performance.

GPU selection

Each of the user classifications in Table 5 may be served by a different NVIDIA graphics card. Careful understanding of each user class’s use of the graphics card should be built by observation and by measurement using tools such as the [GPUProfiler](#). In general, for SOLIDWORKS, Cisco recommends the NVIDIA T4, RTX 6000 and RTX 8000 (Table 6).

Table 6. NVIDIA GPUs for SOLIDWORKS specifications

Specifications	NVIDIA T4	NVIDIA RTX 6000	NVIDIA RTX 8000
GPU and board (architecture)	1 (Turing)	1 (Turing)	1 (Turing)
Compute Unified Device Architecture (CUDA) cores	2560	4608	4608
Tensor cores	320	576	576
RT cores	40	72	72
Memory size	16 GB GDDR6	24 GB GDDR6	48 GB GDDR6
vGPU profiles	1, 2, 4, 8, and 16 GB	1, 2, 3, 4, 6, 8, 12, and 24 GB	1, 2, 3, 4, 6, 8, 12, 16, 24, and 48 GB
Form factor	PCIe 3.0 single slot	PCIe 3.0 dual slot	PCIe 3.0 dual slot
Power	70W	250W	250W
Thermal	Passive	Passive	Passive

As you can see in the table, all three cards are built on NVIDIA’s Turing architecture. Each card contains CUDA cores, Tensor cores, and ray-tracing cores. From a computing perspective, the T4 card is slightly more than 50 percent of the RTX cards. The RTX cards have a larger processor and significantly more memory, making them excellent candidates for hosting very large models and rendering activities.

Keep in mind that you can assign more than one NVIDIA GPU to a virtual workstation for demanding multitaskers.

The Cisco UCS C240 M5 server platform can accommodate up to six NVIDIA T4 cards or up to two NVIDIA RTX cards. The Cisco UCS C220 M5 server can accommodate up to two NVIDIA T4 cards.

Scheduling engine selection

NVIDIA offers a choice of three scheduling engines for use with SOLIDWORKS. Figure 8 provides details about the choices available.

Multi-user graphics cards scheduling engines



NVIDIA supports 3 models

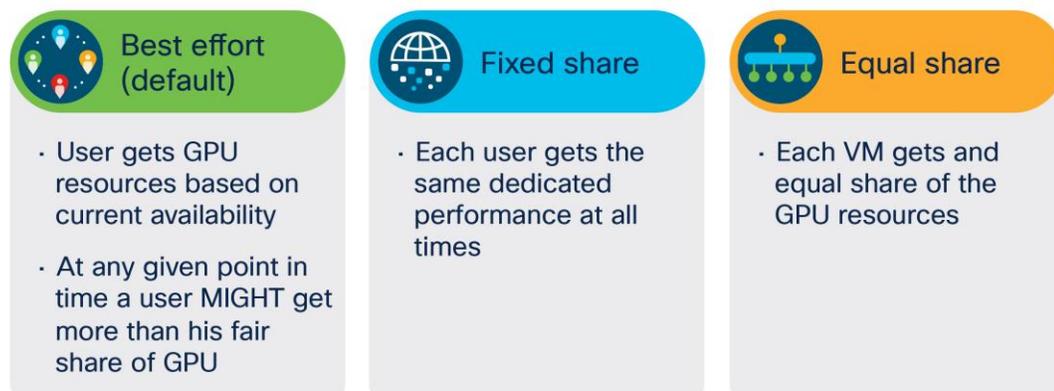


Figure 8.
NVIDIA scheduling engine options

As shown in the figure, the best-effort scheduler is enabled by default. The idea behind this scheduler is that all users assigned to a graphics card will seldom all use its capabilities simultaneously. This scheduler potentially provides users with more GPU resources than their frame buffer allocation would suggest if used strictly to allocate resources. This approach could result in inconsistent behavior if all users assigned to a card are actively using it. However, NVIDIA recommends using this method to schedule the GPU.

Frame buffer selection and frame rates

Frame buffer selection is important for two reasons. First, it allocates the portion of the full GPU that will be allocated to the virtual machine that is assigned to it when all the virtual machines assigned to the card are powered on and being actively used. The frame buffer cannot be oversubscribed. Second, the frame buffer can be used to allocate resources supported by GDDR6 memory. The larger the amount of frame buffer allocated, the more of a model that can be loaded into GPU memory for fast access and processing.

The frame buffer is assigned through the NVIDIA profile selected when the virtual machine is created and the NVIDIA PCI vGPU is assigned. There are several choices for the vWS profiles listed in Table 4.

Frame rates are important for end users, particularly the heavy users identified in Figure 9.

Frame rates



The great equalizer for performance

- For computer video displays,
 - frame rate = #frames or images displayed per second
- For a given application
 - Provides a mechanism to compare systems performance
 - Describes a mechanism by which system requirements can be stated

Virtual Graphics Workstation insights

- Frame rate can be controlled – or not
- Frame rate can be set in the NVIDIA and Desktop Broker software policy

Figure 9.
Frame rates

Target frame rates for smooth manipulation of SOLIDWORKS objects are 24 to 30 frames per second. Higher is better. If everything has been evaluated and provisioned correctly, all user types should achieve this performance level or better for their workloads.

Keep in mind that for virtual graphic workstations, the maximum frame rate allowed by policy is 60 frames per second. The default policy for both Citrix Virtual Apps and Desktops and VMware Horizon is 30 frames per second.

SOLIDWORKS sizing recommendations

Cisco uses a tool from SPEC.org called SPECviewperf 13 to evaluate various GPUs, GPU profiles, and CPU combinations to tune to starting-point recommendations for each individual. This is a publicly available tool that you can use to conduct your own comparative testing. The tool can be found [here](#).

The SPECviewperf 13 benchmark is the worldwide standard for measuring graphics performance based on professional applications. The benchmark measures the 3D graphics performance of systems running under the OpenGL and Direct X APIs. The benchmark's workloads, called viewsets, represent graphics content and behavior from actual applications.

The latest version of the SPECviewperf benchmark is the SPECviewperf 13 benchmark, released on May 23, 2018.

The SPECviewperf 13 benchmark is a comprehensive upgrade of previous versions of the benchmark. Medical and energy viewsets incorporate new models and ray casting for volume visualization; the Maya viewset features new models based on the [SPECapc for Maya 2017](#) benchmark; and the Creo viewset has been updated with fresh application traces. All other viewsets have been recompiled with minor changes. Results from the SPECviewperf 13 benchmark are not comparable to those from earlier versions.

A SPECviewperf 13 benchmark license covers both Microsoft Windows and Linux versions; current paid license holders of the Windows version can receive the Linux edition free of charge.

To evaluate SOLIDWORKS, SPECviewperf 13 was used in nonbenchmark mode, and that application only was evaluated with a series of graphic card, profile, and CPU and memory combinations.

NVIDIA RTX 6000 and RTX 8000 sample results

Figures 10 through 13 show sample data collected from SPECviewperf 13 during the testing process for reference.



Figure 10.
8Q profile (RTX 6000 only)



Figure 11.
12Q profile (RTX 8000 and RTX 6000)

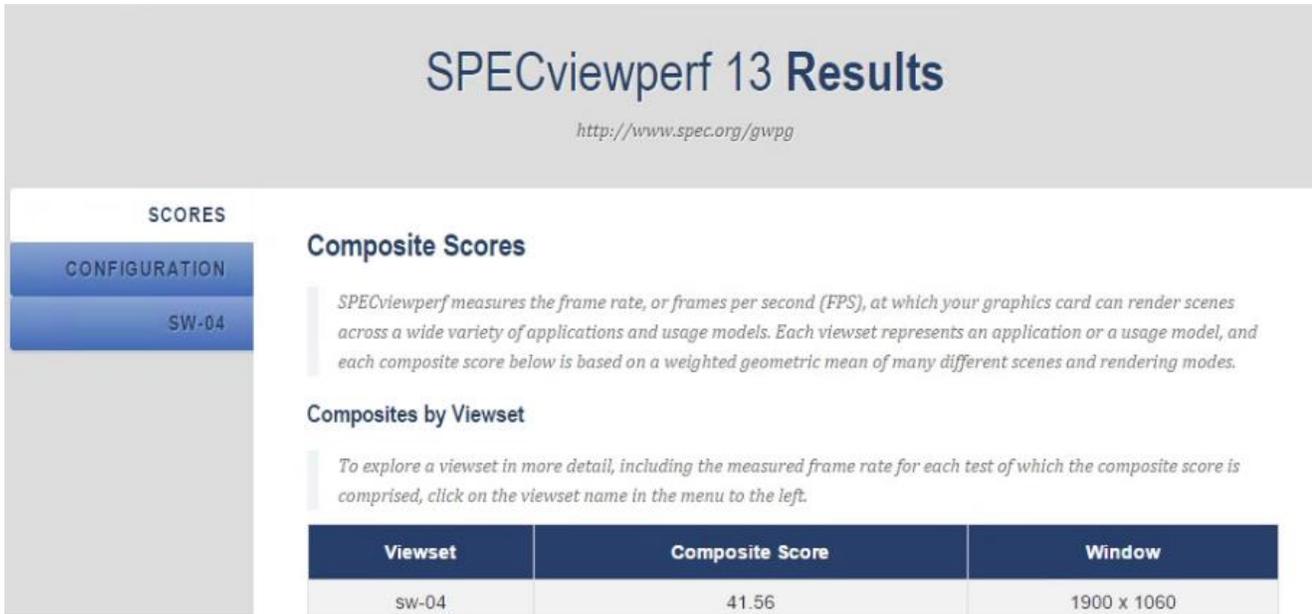


Figure 12.
24Q profile (RTX 8000)



Figure 13.
48Q profile (RTX 8000 only)

The test results led to the recommendations listed in Table 7 for SOLIDWORKS as a starting point for the three user classifications detailed earlier in this document.

Table 7. SOLIDWORKS starting-point virtual graphic workstation configuration by user class

User type	Equivalent performance	Users/server	vCPU/user	Memory/user	Server CPU	Server memory	NVIDIA GPU	Quadro profile	Storage type	Network
Light	Quadro P1000	48	4	12-16	Intel Xeon 6248R	768	NVIDIA T4 (6)	T4-2Q	Flash	10Gb+

User type	Equivalent performance	Users/server	vCPU/user	Memory/user	Server CPU	Server memory	NVIDIA GPU	Quadro profile	Storage type	Network
Medium	Quadro P2000	24	4-6	16-32	Intel Xeon 6246R	768	NVIDIA T4 (6)	T4-4Q	Flash	10Gb+
Heavy	Quadro P5000	4	8-12	96+	Intel Xeon 6246R	768	RTX (2)	RTX6-24Q RTX8-24Q RTX8-48Q	Flash	10Gb+

Note: The recommendations in Table 7 reflect starting points. Customers should run proof-of-concept implementations to determine optimal configurations for their specific environments. Cisco can help.

Ordering information

Contact your Cisco account manager or partner account manager for preconfigured virtual graphic workstation bundles for SOLIDWORKS and other high-performance graphics applications.

Conclusion

With the proper analysis and planning, Dassault Systèmes SOLIDWORKS can be virtualized for use by three standard classes of users, as discussed in this document. The main advantages of virtualizing SOLIDWORKS are summarized in Figure 14.

5 Key Benefits VDI Bring to SOLIDWORKS



Figure 14. Benefits of virtualizing SOLIDWORKS on Cisco UCS servers with NVIDIA graphics cards

Cisco Customer Experience Services

Access the collaborative intelligence of Cisco Customer Experience (CX) Services, which combine world-class expertise and best practices with insights and automation derived from data and artificial intelligence (AI) and machine learning to solve your unique VDI and graphical workstation challenges.

With our CX Success Portfolio, we help you at every step of the lifecycle to achieve optimal value from your technology investments to accelerate your success. [Learn more.](#)

Cisco Capital flexible payment solutions to help you achieve your objectives

Cisco Capital® flexible payment solutions make it easier to get the right technology to help you achieve your objectives, enable business transformation, and stay competitive. We can help you reduce the total cost of ownership (TCO), conserve capital, and accelerate growth. In more than 100 countries, our flexible payment solutions can help you acquire hardware, software, services, and complementary third-party equipment in easy, predictable payments. [Learn more.](#)

For more information

For additional information, see the following: <https://www.cisco.com/c/en/us/solutions/data-center-virtualization/desktop-virtualization/index.html>

Document history

New or revised topic	Described in	Date
SOLDIWORKS Sizing First Release		October 15, 2020

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)