



NVIDIA GRID DASSAULT SYSTÈMES CATIA V5/V6 SCALABILITY GUIDE

NVIDIA Performance Engineering Labs

PerfEngDoc-SG-DSC01v1

March 2016



HOW MANY USERS CAN I GET ON A SERVER?

The purpose of this guide is to give a detailed analysis of how many users organizations can expect to get per server based on the application being used and the types of users they have in their environment. The data for this analysis is sourced from a combination of application recommendations from the independent software vendor (ISV), along with customer data from actual deployments. As with any deployment, your usage will depend on your models and equipment, so this Scalability Guide is best used as a starting point for your implementation. This testing was completed by NVIDIA, Dassault Systèmes does not certify GRID vGPU configuration. Any software issue should be reproduced on an officially certified workstation. Please refer to Dassault Systèmes website for details.

DASSAULT SYSTÈMES CATIA

Dassault Systèmes CATIA is one of the most globally used CAD (Computer Aided Design) applications, being used for almost all aircraft designs in the world and over 80% of the automotive market. As a result, the use of CATIA is pervasive throughout the supply chains supporting these industries. CATIA leverages both CPU and GPU to deliver a high quality user experience so there are several things that need to be considered in architecting your NVIDIA GRID Virtual Workstation (vWS) solution. These include; the size of your data, the concurrency of your users, and the level of interaction with 3D data. Catia has been very successful as a legacy standalone application and has now become a brand as part of the 3DEXPERIENCE (aka V6) platform. This guides gives recommendation on deployment for Catia V5 standalone application.

Recommended CATIA Virtual System Requirements

Working with virtualization partners like Citrix and VMware, and alongside our shared customers with tested and/or production environments, we are recommending the following system requirements for deploying CATIA in a virtual environment. Again, this is a starting point and your requirements will depend on your users and data sets, testing with those users will provide the best guidance:

Recommended Virtualization Platforms			
Virtualization Software	Citrix XenServer v6.5 SP1 or later Citrix XenDesktop v7.8 or later VMware vSphere 6 or later VMware Horizon 7 or later		
Virtual Machine Operating System	Microsoft Windows 7 SP1 64-bit: Enterprise, Ultimate, or Professional Microsoft Windows 8.1 64-bit: Enterprise, or Professional		
Host Server Recommendation	Minimum	Recommended	Advanced
CPU	2.3 GHz+ Intel® Xeon	2.6 GHz+ Intel® Xeon	2.6 GHz+ Intel® Xeon E5

	E5 v3 or greater	E5 v3 or greater	v3 or greater
Memory	196 GB	320 GB	320 GB
Networking	1 GB minimum 10 G recommended	10 G	10 G or greater
Storage	~500+ IOPS Per User	~750+ IOPS Per User	~1000+ IOPS Per User
GPU	NVIDIA Tesla M60/M6 or greater NVIDIA GRID K2 or greater	NVIDIA Tesla M60/M6 or greater NVIDIA GRID K2 or greater	NVIDIA Tesla M60/M6 or greater NVIDIA GRID K2 or greater
Virtual Machine Settings	Minimum	Recommended	3D & Large Datasets
Memory	8 GB RAM	16 GB RAM	24 GB RAM or greater
vCPUs	4 vCPUs	4+ vCPUs	6+ vCPUs
Disk Space	100 GB free disk space	100 GB free disk space	100 GB free disk space
Virtual Graphics License & Profile	NVIDIA GRID Quadro vWorkstation License K240Q / M60-1Q profile or greater	NVIDIA GRID Quadro vWorkstation License K260Q / M60-2Q profile or greater	NVIDIA GRID Quadro vWorkstation License K280Q / M60-4Q profile or greater

CATIA SCALABILITY BASED ON CUSTOMER DATA

Based on testing and production results from several customers in the manufacturing space, here are the expected user densities for the different user types on NVIDIA GRID environments with GRID software and NVIDIA Tesla® M60 or GRID K2 cards. Please note that these results are what these customers determined were appropriate for their needs and provide insight into what one might expect. Results vary based on many variables and doing your own testing will determine your specific scalability.

TECH TIP! It is important to note that mixed user types have produced the best density for some customers, for example placing 1 large and 7 mediums user types per host.

CATIA User Classifications	Application(s)	Description	Users per server
Small/Standard	Dassault Systèmes CATIA	View-Only or Full application Accessing individual parts or small assemblies Manufacturing shop floor	8-16
Medium	Dassault Systèmes CATIA	View-Only or Full application Accessing medium assemblies Manufacturing shop floor	8-16

Large/Power User	Dassault Systèmes CATIA	Full Application Accessing large assemblies or full model	4-8
Analyst	Dassault Systèmes CATIA	Full Application CUDA/OpenCL Analysis of/or accessing large assemblies or full model	4

TYPICAL CATIA VIRTUAL WORKSTATION BUILDS

NVIDIA is developing recommendations for architecting your virtualized CATIA deployment. The following are actual builds from current customers, in production or POC use today, and can provide a starting point for building your own environment. Your own tests with your own models will determine if these recommendations meet your specific needs.

Automotive Manufacturing:

Three automotive manufacturers provided data for CATIA on GRID software with K2 and M6/M60 GPUs. This solution is being used to centralize data and compute into global regions, improving productivity, collaboration, supply chain (see below) integration, and of course security.

Two automotive parts manufacturers provided builds and user profiles that they architect for. These users are either located on a metro area campus setting, or in satellite offices. Having to work with multiple automotive manufacturers requires parts suppliers to work with multiple Dassault Systèmes CATIA versions, and the easiest way to achieve this is with virtual workstations. Virtualization also allows the parts suppliers to centralize their data, collaborate with ease, distribute their workforce, and retain intellectual property.

Resulting User Builds:

User Classification Matrix			
Customer User Classifications	STD CAD Desktop	ADV CAD Desktop	ADV+ CAD Desktop
CATIA Build Spec	4-6 vCPUs 16GB RAM 1GB Frame Buffer (K240Q / M60-1Q)	6 vCPUs 16GB RAM 2GB Frame Buffer (K260Q / M60-2Q)	8-16 vCPUs 16GB RAM 4GB Frame Buffer (K280Q / M60-4Q)
Exported model size	150-200 MB		500-600 MB
Total triangles	~100,000		~1,600,00
Visualization time (CPU vs vGPU)	122ms/52ms		937ms/386ms

Users per host (dual CPU, 2 GRID K2's per host)	16	8	4
NVIDIA GRID Software	Virtual Workstation (vWS) Edition		
Users per host (dual CPU, 2 Tesla M60's per host)	32	16	4+ (CUDA / OpenCL req.)

Aerospace manufacturing:

This aerospace manufacturer has four different user profiles that they architect for. They have run production users of CATIA workloads on both GRID K2 GPUs as well as GRID software and with Tesla M60 GPUs. These users were located in satellite offices, connected via enterprise class WAN links where a 70% productivity improvement was reported thanks to file proximity/transfer time savings alone, to those based on the LAN where even then a 30% productivity gain was reported. With Frame Rate Limiting (FRL) turned on, typical and appropriate for production environments, they saw no increase in Frame Per Second (FPS) but were able to double the number of users per host with no reported degradation of end user performance.

Resulting User Builds:

User Classification Matrix				
Customer User Classifications	Small/Standard	Medium/Moderate	Large/Power User	Jumbo/Analyst
CATIA Build Spec	2vCPU 16GB RAM 1GB Frame Buffer (K240Q / M60-1Q)	4vCPU 32GB RAM 1GB Frame Buffer (K240Q / M60-1Q)	4vCPU 64GB RAM 1GB Frame Buffer (K240Q / M60-1Q)	4vCPU 128GB RAM 1GB Frame Buffer (K240Q / M60-1Q)
Users per host (dual CPU, 2 GRID K2 per host)	16	16	8	4
NVIDIA GRID Software	Virtual Workstation (vWS) Edition			
Users per host (dual CPU, 2 GRID M60 per host)	32	32	16	8

We will continue to collect field data from customers and partners to further refine this data. If you would like to share your data, please reach out to your local NVIDIA team and let us know.

ANALYSIS SUMMARY

CATIA requires significant CPU resources, so investing in higher CPU speeds and more cores will pay off on performance and scalability. Smaller workloads can leverage the 1GB vGPU profile, but for medium to large models, 2GB or greater vGPU profile performance generally produce better results. Since each situation is different, you must test with your own models to ensure the most accurate results.

To test NVIDIA GRID in your environment you can choose to get started with a certified NVIDIA partner or start a POC yourself with a certified server and our 90-day evaluation licenses.

TESTING YOUR NVIDIA GRID ENVIRONMENT

As with all scalability testing, user workloads and environment must be taken into account for scalability analysis. In order to test NVIDIA GRID in your environment you can choose to get started with a [certified NVIDIA partner](#) or start a proof of concept (POC) with a certified server and our [90-day evaluation licenses](#).

Important things to remember during your POC

1. **Define “acceptable” user experience:** Defining user experience (UX) requires careful examination of user and application interaction. This can be obvious, like the rendering time for an image to appear, or smoothly panning across that image. It can also be less obvious, like the ability to smoothly scroll down a page or the “snappy” reaction for a menu to appear after a right click. Ask users to report metrics, and to judge specific activities or functions using finite scales (e.g. 1-5, 5 being best), to avoid generic feedback.
2. **Compare real world workloads:** In virtual environments, time-slicing of resources allows users to get the same level of performance even when sharing resources. This is due to user “think time” which includes any time the user is not actually interacting with the application, or when not using the application or even sitting at their desk. Add up all the time away from the application (meetings, lunch, out of office, etc.) and one could expect to get even more benefits from shared resources. These benefits equate to more resources for the user’s session and typically a more responsive application, thus a better-perceived experience by the end user, as opposed to peak workload benchmarks with inhuman like uninterrupted work.
3. **Test with real users:** It’s important to actually look at the application running be sure that the experience is enjoyable for users. While idle systems or benchmarks can be used as a starting point, real user workflows may require more or less graphics resources. Numbers of monitors and resolution will impact graphics requirements so BYO devices or users that logging in from different locations may impact user experience. Add in the point above about real world workloads and you see why real users are the most accurate means of testing.

Notice

The information provided in this specification is believed to be accurate and reliable as of the date provided. However, NVIDIA Corporation (“NVIDIA”) does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This publication supersedes and replaces all other specifications for the product that may have been previously supplied.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and other changes to this specification, at any time and/or to discontinue any product or service without notice. Customer should obtain the latest relevant specification before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer. NVIDIA hereby expressly objects to applying any customer general terms and conditions with regard to the purchase of the NVIDIA product referenced in this specification.

NVIDIA products are not designed, authorized or warranted to be suitable for use in medical, military, aircraft, space or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on these specifications will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this specification. NVIDIA does not accept any liability related to any default, damage, costs or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this specification, or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this specification. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA. Reproduction of information in this specification is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the NVIDIA terms and conditions of sale for the product.

VESA DisplayPort

DisplayPort and DisplayPort Compliance Logo, DisplayPort Compliance Logo for Dual-mode Sources, and DisplayPort Compliance Logo for Active Cables are trademarks owned by the Video Electronics Standards Association in the United States and other countries.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

ROVI Compliance Statement

NVIDIA Products that support Rovi Corporation’s Revision 7.1.L1 Anti-Copy Process (ACP) encoding technology can only be sold or distributed to buyers with a valid and existing authorization from ROVI to purchase and incorporate the device into buyer’s products.

This device is protected by U.S. patent numbers 6,516,132; 5,583,936; 6,836,549; 7,050,698; and 7,492,896 and other intellectual property rights. The use of ROVI Corporation’s copy protection technology in the device must be authorized by ROVI Corporation and is intended for home and other limited pay-per-view uses only, unless otherwise authorized in writing by ROVI Corporation. Reverse engineering or disassembly is prohibited.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Trademarks

NVIDIA, the NVIDIA logo, NVIDIA GRID, and NVIDIA GRID vGPU are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2016 NVIDIA Corporation. All rights reserved.