



NVIDIA Virtual GPU Software Packaging, Pricing, and Licensing Guide

Application Note

Document History

DA-09924-001_v07

Version	Date	Authors	Description of Change
01	April 22, 2020	CH, SM	Initial Release
02	May 12, 2020	CH, SM	Updates to “Perpetual Concurrent User License” and “NVIDIA Education Pricing Program” sections
03	June 25, 2020	CH, RR, MR	Update NVIDIA Virtual Compute Server subscription price
04	July 20, 2020	NN	Update NVIDIA Virtual Compute Server abbreviated name
05	September 30, 2020	MR, RR	Add Ampere GPU support
06	February 17, 2021	MR, VK, NS, SM	Updated new Ampere GPU support, vGPU rebranding changes
07	April 14, 2021	NS, SM	Added A10 and A16

Table of Contents

- Terminology 1
- Overview 2
 - General Purchasing Information 5
- Virtual GPU Software Product Details 7
 - Virtual GPU Software Products 7
 - NVIDIA vApps 8
 - NVIDIA vPC 8
 - NVIDIA RTX vWS 8
 - NVIDIA Virtual Compute Server 8
 - vGPU Software Products and Entitlement 8
- vGPU Software Licensing and Pricing 11
 - Enterprise vGPU Software Pricing 11
 - Subscription Concurrent User License 11
 - Subscription per GPU User License 12
 - Perpetual Concurrent User License 12
 - NVIDIA Education Pricing Program 13
- Deciding the Right License Based on Capability and Entitlement 14
- NVIDIA vGPU Software License Server 17

List of Tables

Table 1.	Terminology	1
Table 2.	NVIDIA Virtual GPU Software Licensed Products.....	2
Table 3.	Supported NVIDIA Maxwell Graphics Cards	3
Table 4.	Supported NVIDIA Pascal Graphics Cards.....	3
Table 5.	Supported NVIDIA T4, NVIDIA RTX, and NVIDIA V100 Graphics Cards.....	4
Table 6.	Supported NVIDIA Ampere Graphics Cards.....	4
Table 7.	General Purchasing Information	5
Table 8.	Feature Entitlements	9
Table 9.	Concurrent User License Annual Subscription Pricing	12
Table 10.	GPU User License Annual Subscription Pricing.....	12
Table 11.	Enterprise Perpetual Licensing plus SUMS Pricing.....	12
Table 12.	NVIDIA Education Pricing.....	13
Table 13.	Common Use Cases and Solutions	16

Terminology

The following table lists the terms that will be used throughout this guide.

Table 1. Terminology

Term	Meaning
SUMS	Support, Upgrade and Maintenance program
Perpetual License	A non-expiring, permanent software license that can be used on a perpetual basis without the need to renew. SUMS is required and is available in four-year or five-year increments. One-year SUMS available only for renewals.
Annual Subscription	A software license that is active for a fixed period as defined by the terms of the subscription. An annual subscription includes SUMS for the duration of the license term.
License Server	An application that manages license allocation, installed on a physical or virtual server.
Concurrent User (CCU) Licensing	A method of allocating licenses based on the number of VMs that are concurrently being used. A CCU license allows only one concurrent VM to be hosted per license.
GPU	Graphics processing unit
Per GPU Licensing	A method of allocating licenses based on the number of GPUs that are being used. A per GPU license allows a certain number of concurrent VMs to be hosted per license.

Overview

NVIDIA virtual GPU (vGPU) solutions bring the power of NVIDIA GPUs to virtual desktops, applications, and workstations, accelerating graphics and compute to make virtualized workspaces accessible to creative and technical professionals working from home offices or anywhere.

This guide covers the entitlement, packaging, and licensing of the NVIDIA virtual GPU (vGPU) software family of products. It is intended to be a quick reference to understand the product portfolio at a high level, with the corresponding SKU information. It does not contain detailed product information, which can be obtained from the NVIDIA vGPU website at <http://www.nvidia.com/virtualgpu>. This guide is not intended to replace or contradict the End User License Agreement (EULA). Refer to the [EULA](#) for more detailed information.

Table 2. NVIDIA Virtual GPU Software Licensed Products

Product	Description
NVIDIA Virtual Applications (vApps)	For organizations deploying Citrix Virtual Apps and Desktops, RDSH or other app streaming or session-based solutions. Designed for PC level applications and server-based desktops.
NVIDIA Virtual PC (vPC)	For users who want a virtual desktop but need great user experience leveraging PC Windows applications, browsers, and high-definition video.
NVIDIA RTX™ Virtual Workstation (vWS)	For users who want to be able to use remote professional graphics applications with full performance on any device, anywhere.
NVIDIA Virtual Compute Server (vCS)	For compute-intensive server workloads, such as artificial intelligence (AI), deep learning, or high-performance computing (HPC).

NVIDIA vGPU software brings graphics and virtualization capabilities to NVIDIA data center deployments and is currently supported on the NVIDIA graphics cards listed in Table 3 through Table 6. Find certified servers that are supported by NVIDIA GPUs and vGPU software at <https://www.nvidia.com/en-us/data-center/resources/vgpu-certified-servers/>.

Table 3. Supported NVIDIA Maxwell Graphics Cards

	NVIDIA M10	NVIDIA M60	NVIDIA M6
Recommended Use case	NVIDIA vPC – User Density Optimized	Performance-Optimized	Blade-Optimized
Number of GPUs	4 NVIDIA Maxwell™ GPUs	2 NVIDIA Maxwell GPUs	1 NVIDIA Maxwell GPU
Total NVIDIA® CUDA® Cores	2,560 (640 per GPU)	4,096 (2,048 per GPU)	1,536
Total Memory Size	32 GB GDDR5	32 GB GDDR5	8 GB GDDR5
(8 GB per GPU)	16 GB GDDR5	16 GB GDDR5	
(8 GB per GPU)	8 GB GDDR5	8 GB GDDR5	
Max Power	225 W	300 W	100 W

Table 4. Supported NVIDIA Pascal Graphics Cards

	NVIDIA P4	NVIDIA P6	NVIDIA P40	NVIDIA P100
Recommended Use case	NVIDIA vWS (Entry to mid)-	NVIDIA vWS or NVIDIA vPC - Blade-Optimized	NVIDIA vWS - Performance Optimized (Mid to high)	NVIDIA vWS (Mid to high)
Number of GPUs	1 NVIDIA Pascal GP104	1 NVIDIA Pascal GP104	1 NVIDIA Pascal GP102	1 NVIDIA Pascal GP100
Total CUDA Cores	2,560	2,048	3,840	3,584
Total Memory Size	8 GB GDDR5	16 GB GDDR5	24 GB GDDR5	12 GB HBM2/16 GB HBM2
Max Power	75 W	90 W	250 W	250 W/300 W
Form Factor	PCIe 3.0 Single Slot	MXM	PCIe 3.0 Dual Slot	SXM2/PCIe 3.0
Board Dimensions	2.7" × 6.6"	3.2" × 4.1"	10.5" × 4.4"	10.5" × 4.4"
Cooling Solution	Passive	Bare Board	Passive	Passive

Table 5. Supported NVIDIA T4, NVIDIA RTX, and NVIDIA V100 Graphics Cards

	NVIDIA T4	NVIDIA RTX 6000	NVIDIA RTX 8000	NVIDIA V100S/V100 SXM2
Recommended Use case	NVIDIA vWS - Performance Optimized (Entry); vPC - Density Optimized; vCS - Compute Optimized	NVIDIA vWS - Performance Optimized (High end); vCS - Compute Optimized	NVIDIA vWS - Performance Optimized (High end); vCS - Compute Optimized	NVIDIA vWS - Performance-Optimized; vCS - Compute Optimized
Number of GPUs	1 NVIDIA Turing™ TU104	1 NVIDIA Turing TU102	1 NVIDIA Turing TU102	1 NVIDIA Volta™ GPU
Total CUDA Cores	2,560	4,608	4,608	5,120
Tensor Cores	320	576	576	640
NVIDIA® TensorRT™ Cores	40	72	72	--
Total Memory Size	16 GB GDDR6	24 GB GDDR6	48 GB GDDR6	32 GB HBM2
Max GPU Power	70 W	260 W / 250 W	260 W / 250 W	250 W / 300 W
Form Factor	PCIe 3.0 Single Slot	PCIe 3.0 Dual Slot	PCIe 3.0 Dual Slot	PCIe, SXM2, Full Height/Length
Board Dimensions	2.7" × 6.6"	10.5" × 4.4"	10.5" × 4.4"	10.5" × 4.4" 5.5" × 3.1" × .5"
Cooling Solution	Passive	Active or Passive	Active or Passive	Passive

Table 6. Supported NVIDIA Ampere Graphics Cards

	NVIDIA HGX A100	NVIDIA A100	NVIDIA A40	NVIDIA A10
Recommended Use case	vCS - Compute Optimized	vCS - Compute Optimized	NVIDIA vWS - Performance Optimized (High end); vCS - Compute Optimized	NVIDIA vWS - Performance Optimized (midrange) vCS - Compute Optimized
Number of GPUs	4 NVIDIA A100/ 8 NVIDIA A100	1 NVIDIA A100	1 NVIDIA A40	1 NVIDIA A10
FP32 Cores / GPU	6,912	6,912	10,572	9,216
Tensor Cores / GPU	432	432	336	288
TensorRT Cores	-	-	84	72
Total Memory Size / GPU	40 GB HBM2/80GB HBM2	40 GB HBM2	48 GB GDDR6	24 GB GDDR6
Max GPU Power / GPU	400 W	250 W	300 W	150 W

	NVIDIA HGX A100	NVIDIA A100	NVIDIA A40	NVIDIA A10
Form Factor	4x SXM4 GPUs/8x SXM4	PCIe 4.0 Dual-Slot FHFL	PCIe 4.0 Dual-Slot FHFL	PCIe 4.0 Single Slot FHFL
Board Dimensions		10.5" × 4.4"	10.5" × 4.4"	10.5" × 4.4"
Cooling Solution	Passive	Passive	Passive	Passive



Note: NVIDIA RTX A6000 and NVIDIA RTX A5000 are also supported with NVIDIA virtual GPU software.

General Purchasing Information

NVIDIA vGPU software products can be purchased through NVIDIA Preferred Partners and select server OEMs. A list of these Preferred Partners and OEMs can be obtained from: <http://www.nvidia.com/buygrid>.

NVIDIA vGPU software products can be purchased as either a perpetual license with a Support Updates and Maintenance Subscription (SUMS), or as an annual subscription. The perpetual license gives the user the right to use the software indefinitely, with no expiration. All NVIDIA vGPU software products with perpetual licenses must be purchased in conjunction with a four-year, or five-year SUMS. A one-year SUMS is available only for renewals.

The annual subscription offering is a more affordable option to allow IT departments to better manage the flexibility of license volumes. NVIDIA vGPU software products with an annual subscription are bundled with SUMS for the duration of the software's subscription license.

Table 7. General Purchasing Information

Entitlement	NVIDIA vGPU Production SUMS
Maintenance	Access to all maintenance releases, defect resolutions, and security patches for flexibility in upgrading as per the NVIDIA Virtual GPU Software Lifecycle Policy
Upgrades	Access to all new major version releases including feature enhancements and new hardware support
Long-term branch maintenance	Available for up to 3 years from general availability as per the NVIDIA Virtual GPU Software Lifecycle Policy
Direct support	Direct access to NVIDIA support engineering for timely resolution of customer-specific issues
Support availability	Customer support available during standard business hours Cases accepted 24 × 7
Knowledgebase access	✓

Entitlement	NVIDIA vGPU Production SUMS
Web support	✓
E-mail support	✓
Phone support	✓

Virtual GPU Software Product Details

NVIDIA vGPU software is the industry's most advanced technology for sharing true virtual GPU hardware acceleration between multiple users—without compromising the graphics experience. This virtualization technology ensures complete application compatibility, which means features and experience are the same as they would be on a physical device.

Virtual GPU Software Products

NVIDIA vGPU desktop and application virtualization solutions are designed to bring the power of virtualization to the users who need to be their most productive. vGPU technology ensures application compatibility, meaning any application that can run in a physical desktop environment can run in a virtual environment. Organizations can now expand their virtualization footprint without compromise.

The following NVIDIA vGPU software products are available:

- ▶ NVIDIA Virtual Applications (vApps)
- ▶ NVIDIA Virtual PC (vPC)
- ▶ NVIDIA RTX Virtual Workstation (vWS)
- ▶ NVIDIA Virtual Compute Server (vCS)

NVIDIA vApps

This product is for organizations deploying Citrix Virtual Apps and Desktops, RDSH or other app streaming or session-based solutions. Designed to deliver PC Windows applications at full performance, NVIDIA vApps allows users to access any Windows application at full performance on any device, anywhere.

Windows Server hosted RDSH desktops are also supported by NVIDIA vApps.

NVIDIA vPC

This product is ideal for users who want a virtual desktop but need great user experience leveraging PC Windows applications, browsers, and high-definition video. NVIDIA vPC delivers a native experience to users in a virtual environment, allowing them to run all their PC applications at full performance.

NVIDIA RTX vWS

This product is ideal for mainstream and high-end designers who use powerful 3D content creation applications such as Dassault CATIA, SOLIDWORKS, 3DExcite, Siemens NX, PTC Creo, Schlumberger Petrel, or Autodesk Maya. NVIDIA vWS allows users to access their professional graphics applications with full features and performance, anywhere, on any device.

NVIDIA Virtual Compute Server

This product is for organizations running compute-intensive server workloads such as Artificial Intelligence (AI), Deep Learning (DL), and High-Performance Computing (HPC). NVIDIA Virtual Compute Server (vCS) is software that enables the NVIDIA GPU to be virtualized to accelerate compute-intensive server workloads with features such as error-correcting code (ECC) memory, page retirement, peer-to-peer CUDA transfers over NVIDIA® NVLink®, and multiple vGPUs assigned to a single VM.

vGPU Software Products and Entitlement

Each NVIDIA vGPU software product includes the feature entitlements listed in Table 8.

Table 8. Feature Entitlements

Feature	NVIDIA vApps	NVIDIA vPC	NVIDIA vWS	NVIDIA vCS
License Entitlement				
Concurrent User (CCU)	✓	✓	✓	
Per GPU				✓ ¹
Capability Entitlement				
Desktop Virtualization		✓	✓	
RDSH App Hosting	✓	✓	✓ ²	
RDSH Desktop Hosting	✓	✓	✓ ²	
Compute Virtualization			✓	✓
Windows Guest OS Support	✓	✓	✓	
Linux Guest OS Support		✓	✓	✓
Maximum Displays	1 ³	4	4	1
Maximum Resolution ⁴	1280 × 1024	5120 × 2880 (5K)	7680 × 4320 (8K)	4096 × 2160 (4K)
NVIDIA RTX Enterprise Software Features			✓	
OpenGL, DirectX and Vulkan	✓	✓	✓	In-situ Graphics only
CUDA and OpenCL Support			✓ ⁵	✓
ECC and Page Retirement			✓	✓ ⁶
Multi-vGPU			✓	✓
NVLink			✓	✓
GPU Pass Through Support ⁷	✓		✓	✓
Bare Metal Support ⁸	✓		✓	
vGPU Profile Sizes Supported⁹				
512 MB		✓	✓	
1 GB	✓	✓	✓	
2 GB	✓	✓	✓	
3 GB	✓		✓	
4 GB	✓		✓	✓
5 GB				✓
6 GB	✓		✓	✓
8 GB	✓		✓	✓
10 GB				✓
12 GB	✓		✓	✓

Feature	NVIDIA vApps	NVIDIA vPC	NVIDIA vWS	NVIDIA vCS
16 GB	✓		✓	✓
20 GB				✓
24 GB	✓		✓	✓
32 GB	✓		✓	✓
40 GB				✓
48 GB			✓	✓
80 GB				✓

Notes:

¹Maximum 8 concurrent VMs per GPU

²With packaged NVIDIA vApps license

³Applies only to the console display in remote application environments. For details, see [Supported GPUs](#)

⁴Review the [Virtual GPU Software User Guide](#) for the supported display configurations for each profile

⁵Supported on 8GB 1:1 profile on NVIDIA Maxwell and all profiles on Pascal

⁶ECC support begins with Pascal

⁷Only supported on 1:1 profiles

⁸Only NVIDIA M6 hardware supported as primary display device

⁹Review the [Virtual GPU Software User Guide](#) for the vGPU profiles supported on your GPU

vGPU Software Licensing and Pricing

Enterprise vGPU Software Pricing

NVIDIA vPC, NVIDIA RTX vWS, and NVIDIA vApps are available on a per Concurrent User (CCU) model. A CCU license is required for every user who is accessing or using the software at any given time, whether an active connection to the virtualized desktop or session is maintained.

NVIDIA vCS is available on a per GPU model. A GPU license is required for every GPU that will host vCS-enabled VMs. A single vCS license enables a maximum of 20 concurrent VMs.

NVIDIA vGPU software products can be purchased by enterprises as either perpetual licenses with annual Support Updates and Maintenance Subscription (SUMS), or as an annual subscription. All NVIDIA vGPU software products with perpetual licenses must be purchased in conjunction with a four-year or five-year SUMS. A one-year SUMS is available only for renewals. For annual licenses, SUMS is bundled into the annual license cost.

Enterprise vGPU software licensing and pricing are described in the following subsections. [Find the full SKU list here](#). Pricing is suggested pricing only. Contact your authorized NVIDIA partner for final pricing. If you are looking to run or host a service by using NVIDIA vGPU software, you need to join the [NPN CSP partner program](#).

Subscription Concurrent User License

An annual Enterprise subscription is active for a fixed period as defined by the terms of the subscription license. To be kept active, the license will need to be renewed at the end of the subscription period. The subscription license includes the software license and production level SUMS for the duration of the license subscription period.

Table 9. Concurrent User License Annual Subscription Pricing

License	Annual Subscription Pricing
NVIDIA Virtual Applications	\$10 per CCU subscription
NVIDIA Virtual PC	\$50 per CCU subscription
	4 year subscription - \$43.75 per CCU
	5 year subscription - \$40 per CCU
NVIDIA RTX Virtual Workstation	\$250 per CCU subscription

Subscription per GPU User License

An annual enterprise subscription is active for a fixed period as defined by the terms of the subscription license. To be kept active, the license will need to be renewed at the end of the subscription period. The subscription license includes the software license and production-level SUMS for the duration of the license subscription period.

Table 10. GPU User License Annual Subscription Pricing

License	Annual Subscription Pricing
NVIDIA Virtual Compute Server	\$150 per GPU subscription

Perpetual Concurrent User License

A perpetual enterprise license allows for use of the licensed software indefinitely. Users that opt to license using this model are required to subscribe to SUMS for four or five years. The SUMS subscription can be renewed on a yearly basis after the expiration of the initial subscription.

Table 11. Enterprise Perpetual Licensing plus SUMS Pricing

License	Annual Licensing plus SUMS Pricing
NVIDIA Virtual Applications	\$20 per CCU perpetual license
	\$5 SUMS per year
NVIDIA Virtual PC	\$100 per CCU perpetual license
	\$25 SUMS per year
NVIDIA RTX Virtual Workstation	\$450 per CCU perpetual license
	\$100 SUMS per year

NVIDIA Education Pricing Program

The NVIDIA Education Pricing Program supports the use of visual computing in teaching and research institutions. The program makes it easy to procure and administer NVIDIA solutions, software licensing, and services for qualified educational institutions and helps reduce their total cost. For more information on eligibility, review the [NVIDIA Education Pricing Program](#) documentation.

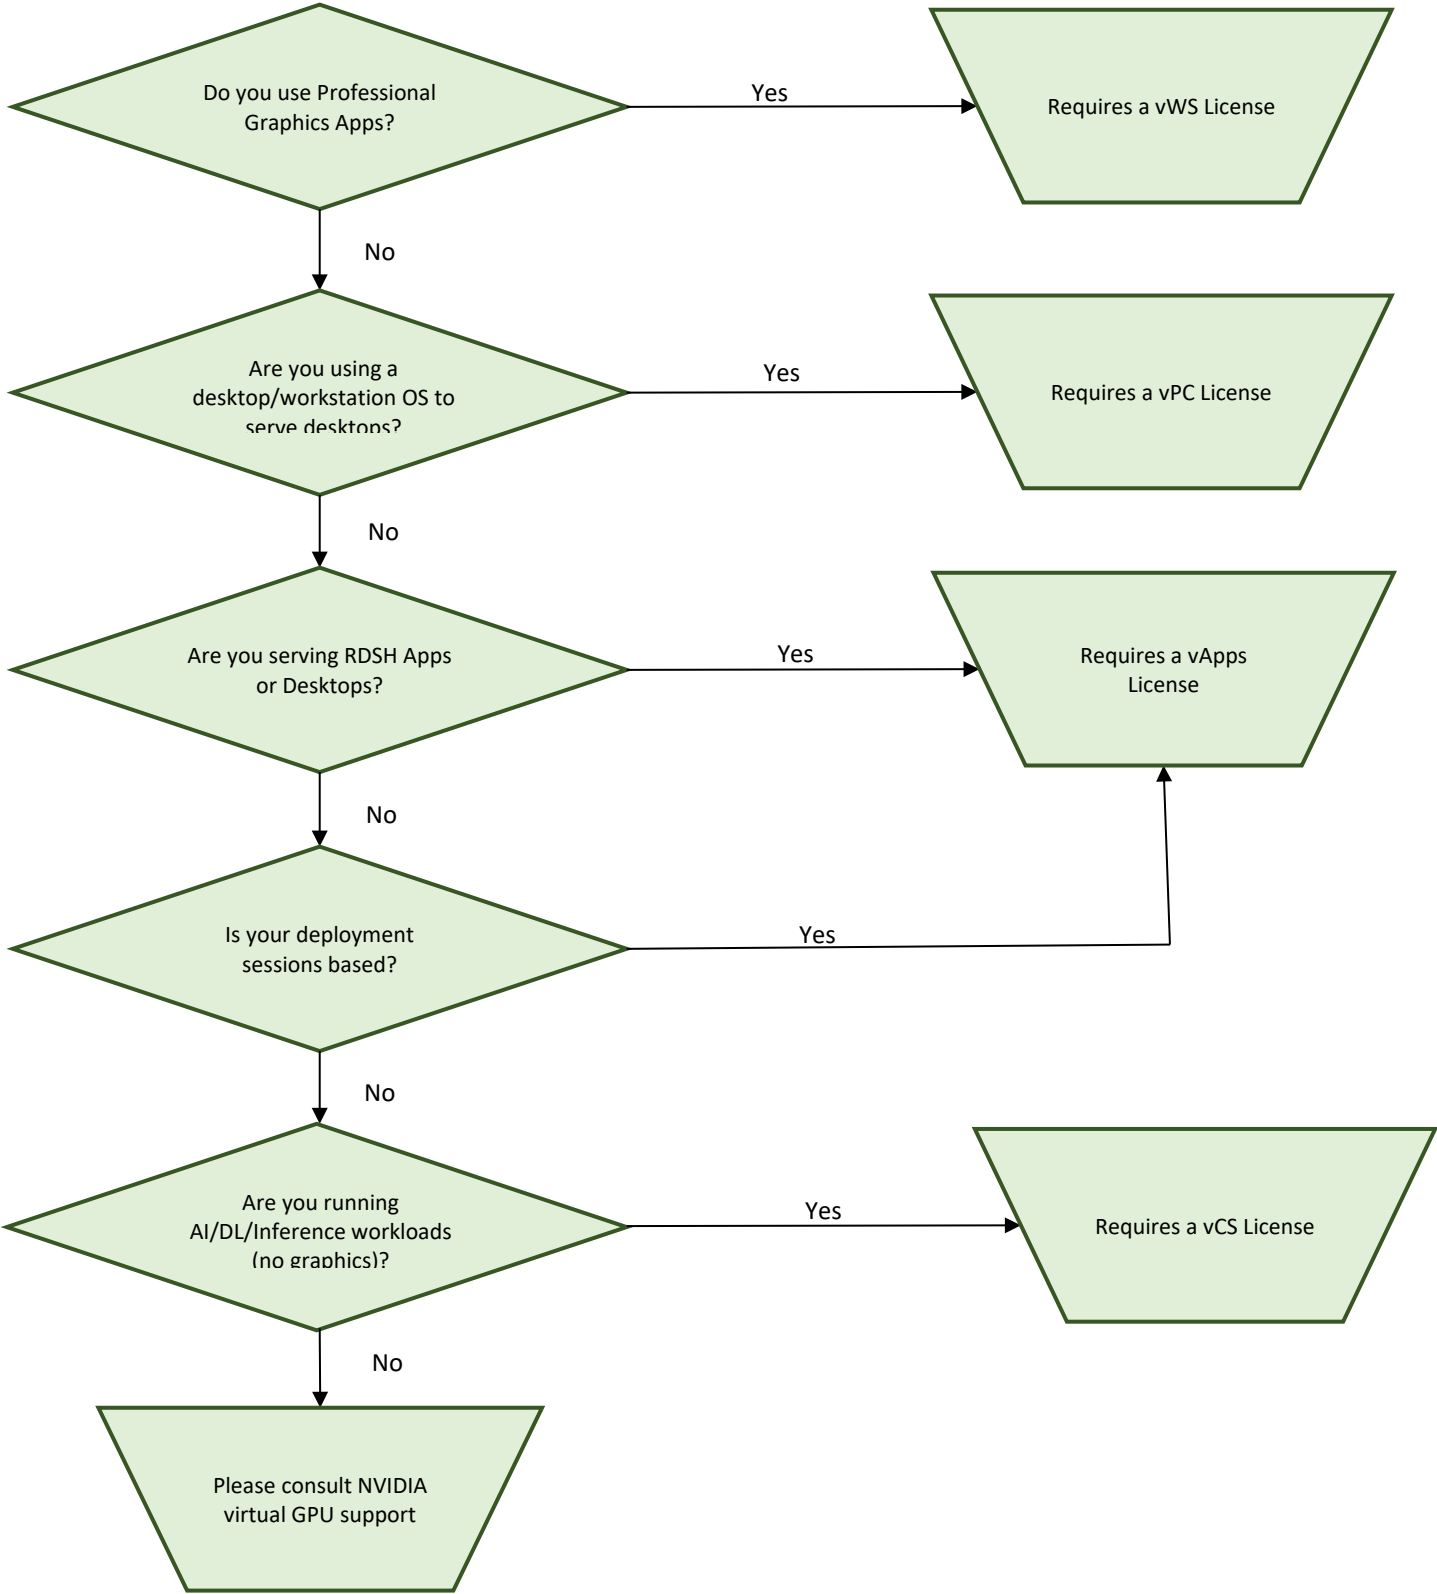
The Annual Subscription Model is available for NVIDIA vWS and vCS. The Perpetual Licensing Model is available only for NVIDIA vWS. Users that opt for the Perpetual Licensing Model are required to subscribe to SUMS for four or five years. The SUMS subscription can be renewed on a yearly basis after the expiration of the initial subscription.

Table 12. NVIDIA Education Pricing

License	Pricing Model	Price
NVIDIA vWS	Perpetual Licensing plus SUMS	\$99 per CCU perpetual license
		\$25 SUMS per year
	Annual Subscription	\$50 per CCU subscription
NVIDIA Virtual Compute Server	Annual Subscription	\$50 per GPU subscription

Deciding the Right License Based on Capability and Entitlement

The following flowchart provides a simple decision tree to help decide which license is required based on the desired entitlement and capability. If you have further questions or are unable to decide based on the decision tree, contact NVIDIA vGPU Support at <https://www.nvidia.com/en-us/support/enterprise/>.



The following table summarizes some common use cases of different solutions. This is not an all-inclusive list of possible solutions. If you have questions, contact NVIDIA vGPU Support.

Table 13. Common Use Cases and Solutions

I am using...	I need this license...
Citrix Virtual Desktops	NVIDIA vPC – for PC level applications NVIDIA vWS – for workstation/professional 3D use cases
VMware Horizon (View)	NVIDIA vPC – for PC level applications NVIDIA vWS – for each session utilizing a workstation/professional 3D use case
Citrix Virtual Apps and Desktops	NVIDIA vApps NVIDIA vWS – for each session utilizing a workstation/professional 3D use case
VMware Horizon RDSH	NVIDIA vApps NVIDIA vWS – for each session utilizing a workstation/professional 3D use case
Other RDSH / Session-based	NVIDIA vApps NVIDIA vWS – for each session utilizing a workstation/professional 3D use case
Microsoft RemoteFX	NVIDIA vPC – for PC level applications
VMware Horizon vSGA	NVIDIA vPC – for PC level applications
Microsoft Hyper-V (DDA)	NVIDIA vWS
Red Hat KVM	vCS – for AI/DL/Inference workloads NVIDIA vWS – for workstation/professional 3D use cases

NVIDIA vGPU Software License Server

The NVIDIA vGPU software license server provides monitoring and reporting on license usage for capacity planning. It is available for download with the vGPU software packages from the NVIDIA Licensing Portal. This license server can be installed on either a physical server or, more likely, a dedicated virtual machine.

To ensure that the end user experience is not impaired by license overages or connection issues, the vGPU software runs with or without a valid license server connection. The license server allows the IT administrator to track license usage and determine correct sizing for their environments. In compliance with the EULA, IT administrators may also use any other method that reliably tracks the software usage to ensure they have enough licenses for their CCU or per GPU usage.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. 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 document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information 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 ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

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 this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform 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 document. NVIDIA accepts no 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 document 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 document. 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 document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND 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. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. 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 Terms of Sale for the product.

OpenCL

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

Trademarks

NVIDIA, the NVIDIA logo, CUDA, NVIDIA HGX, NVIDIA Maxwell, NVIDIA Turing, NVIDIA RTX, NVIDIA Volta, NVLink, and TensorRT 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

© 2020, 2021 NVIDIA Corporation. All rights reserved.

