



RTX VIRTUAL WORKSTATION CLOUD ON GOOGLE CLOUD PLATFORM

DU-09454-001 _v10 | October 2022

Quick Start Guide



TABLE OF CONTENTS

Chapter 1. Creating and Using an NVIDIA RTX Virtual Workstation Instance from the Google Cloud Platform Marketplace.....	1
1.1. Creating a GPU-Accelerated Virtual Workstation.....	1
1.2. Connecting to a GPU-Accelerated Virtual Workstation on a Windows VM Instance.....	3
1.3. Verifying the Creation of your GPU-Accelerated Virtual Workstation.....	4
1.3.1. Verifying the Creation of a GPU-Accelerated Virtual Workstation on a Windows VM Instance.....	4
1.3.2. Verifying the Creation of a GPU-Accelerated Virtual Workstation on a Linux VM Instance.....	5
1.4. Trying Out your GPU-Accelerated Virtual Workstation.....	6
1.4.1. Trying Out your GPU-Accelerated Virtual Workstation with Adobe Premiere Pro.....	6
1.4.2. Trying Out your GPU-Accelerated Virtual Workstation with Esri ArcGIS Pro.....	6
Appendix A. Resources for NVIDIA RTX Virtual Workstation on Google Cloud Platform.....	7

Chapter 1.

CREATING AND USING AN NVIDIA RTX VIRTUAL WORKSTATION INSTANCE FROM THE GOOGLE CLOUD PLATFORM MARKETPLACE

NVIDIA® RTX™ Virtual Workstation in the cloud is an NVIDIA Virtual Machine Image (VMI) preconfigured with NVIDIA RTX Virtual Workstation software and NVIDIA GPU hardware. The NVIDIA RTX Enterprise driver is preinstalled on the VMI and NVIDIA ensures that the image is always up to date with the latest NVIDIA RTX ISV certifications, patches, and upgrades. Support and technical information to help you get started are available on the [NVIDIA RTX Virtual Workstation \(vWS\) on CSP Marketplace](#) community forum and from additional resources.

1.1. Creating a GPU-Accelerated Virtual Workstation

Prepare for creating a GPU-accelerated virtual workstation from the Google Cloud Platform marketplace as follows:

- ▶ Ensure that you have a Google Cloud Platform account and an active subscription.
- ▶ Decide the machine type of the VM instance on which you want to deploy your GPU-accelerated virtual workstation.

For information about the machine types that support NVIDIA RTX Virtual Workstation, see [NVIDIA RTX Virtual Workstation on Google Cloud Platform Release Notes](#).

- ▶ If you are creating an Ubuntu VM instance and want to use SSH public key authentication for the administrator account, generate an SSH public key.
1. In a web browser, log on to [Google Cloud Platform portal \(https://console.cloud.google.com/\)](https://console.cloud.google.com/).

2. Go to the NVIDIA RTX Virtual Workstation listing for the guest OS that you want to use:
 - ▶ [NVIDIA RTX Virtual Workstation - Windows Server 2019](#)
 - ▶ [NVIDIA RTX Virtual Workstation - Windows Server 2016](#)
 - ▶ [NVIDIA RTX Virtual Workstation - Ubuntu 18.04](#)
3. Click **LAUNCH ON COMPUTE ENGINE** to go to the **New deployment** page for NVIDIA RTX Virtual Workstation on your chosen operating system.
4. On the **New deployment** page, provide the project details and instance details.
 - a) From the drop-down list at the top of the page, select your project.
 - b) If you don't want the default assigned deployment name, enter the name that you want for your NVIDIA RTX Virtual Workstation instance.
 - c) Select the zone for your NVIDIA RTX Virtual Workstation instance.

Ensure that the GPU type that you want is available in the selected zone. Not all GPU types that support NVIDIA RTX Virtual Workstation are available in all zones.
 - d) From the **Number of GPUs** drop-down list, select the number of GPUs that you want for your instance.

The number of number of GPUs selected limits the number of vCPUs allowed in your machine type. A maximum of 16 vCPUs is allowed for each GPU.
 - e) From the **GPU Type** drop-down list, select the GPU type that you want for your instance.

If the GPU type that you want to use is not listed, select a different zone until you find the GPU type that you want.
 - f) Specify the machine type with which you want to create your instance.
 - ▶ To create an instance with a predefined machine type, from the drop-down list under **Machine type**, select the machine type that you want.
 - ▶ To create an instance with custom virtualized hardware settings, follow the **Customize** link.

For guidance, see the following Google Cloud Platform documentation:

- ▶ [GPUs on Compute Engine](#)
 - ▶ [Regions and Zones](#)
 - ▶ [Machine Types](#)
 - ▶ [Creating a VM Instance with a Custom Machine Type](#)
5. In the **Boot Disk** and **Networking** sections, review the default settings to see if they meet your requirements and change any settings as required.

For guidance, see the following Google Cloud Platform documentation:

- ▶ [Storage Options](#)
 - ▶ [Virtual Private Cloud \(VPC\) Network Overview](#)
6. Click **Deploy**.

Deployment is complete in a few minutes.

- ▶ If you created a Windows VM instance, follow these instructions to create a password for the instance and connect to the VM instance:
 1. [Creating Passwords for Windows Instances](#) in the Google Cloud Platform documentation
 2. [Connecting to a GPU-Accelerated Virtual Workstation on a Windows VM Instance](#)
- ▶ If you created a Linux VM instance, follow these instructions to connect to the VM instance and verify that it was created properly:
 1. [Connecting to Linux instances](#) in the Google Cloud Platform documentation
 2. [Verifying the Creation of your GPU-Accelerated Virtual Workstation](#)

1.2. Connecting to a GPU-Accelerated Virtual Workstation on a Windows VM Instance



For instructions for an Ubuntu VM instance, see in [Connecting to Linux instances](#) in the Google Cloud Platform documentation.

Ensure that you have created a windows instance password as explain in [Creating Passwords for Windows Instances](#) in the Google Cloud Platform documentation.

1. On the Google Cloud Platform console **VM instances page**, from the **RDP** drop-down list, select **Download the RDP file**.
2. When the download is complete, double-click the RDP file to start a Remote Desktop Connection session on the VM.
3. If you are warned that the publisher of the remote connection cannot be identified and are asked about whether to connect anyway, click **Connect**.
4. When you are prompted, log in to the VM with the credentials for the account that you specified when you created the password for your Windows VM Instance.
5. If you are warned that the publisher of the remote connection cannot be identified and are asked about whether to connect anyway, click **Yes**.

After connecting to your GPU-accelerated virtual workstation, verify that it was created properly as explained in [Verifying the Creation of your GPU-Accelerated Virtual Workstation](#).

1.3. Verifying the Creation of your GPU-Accelerated Virtual Workstation

After connecting to your GPU-accelerated virtual workstation, verify that it was created properly by listing its GPUs. On a Windows VM instance you can also use **NVIDIA Control Panel** to verify that the NVIDIA driver is running.

For details, see:

- ▶ [Verifying the Creation of a GPU-Accelerated Virtual Workstation on a Windows VM Instance](#)
- ▶ [Verifying the Creation of a GPU-Accelerated Virtual Workstation on a Linux VM Instance](#)

After verifying the creation of your GPU-accelerated virtual workstation, you are now ready to run your design and engineering software.

1.3.1. Verifying the Creation of a GPU-Accelerated Virtual Workstation on a Windows VM Instance

1. Open a command prompt window and change to the folder that contains the `nvidia-smi` command.

```
C:\Program Files (x86)\Google\Cloud SDK>cd C:\Program Files\NVIDIA Corporation\NVSMI
```

2. List the GPUs in your GPU-accelerated virtual workstation by running the `nvidia-smi` command without any options.

The following example shows the output from `nvidia-smi` for a Windows VM instance configured with a single NVIDIA T4 GPU.

```
C:\Program Files\NVIDIA Corporation\NVSMI>nvidia-smi
Wed Apr 03 01:49:00 2019

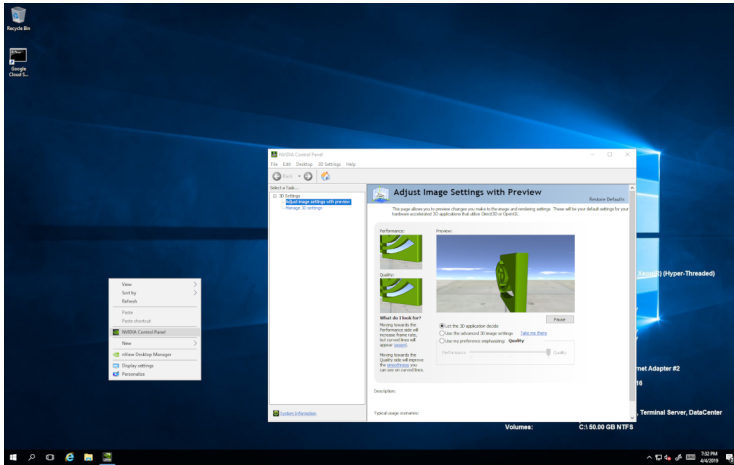
+-----+
| NVIDIA-SMI 412.16      Driver Version: 412.16      CUDA Version: 10.0      |
+-----+-----+-----+-----+-----+-----+
| GPU   Name           TCC/WDDM | Bus-Id      Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+-----+-----+-----+-----+
|    0   Tesla T4            WDDM | 00000000:00:04.0 Off  |            0         |
| N/A   36C    P8     12W / 70W |  310MiB / 15360MiB |    16%    Default   |
+-----+-----+-----+-----+-----+-----+

+-----+
| Processes:
| GPU      PID   Type   Process name                      GPU Memory |
+-----+-----+-----+-----+-----+-----+
|    0     404   C+G    Insufficient Permissions           N/A       |
|    0     896   C+G    Insufficient Permissions           N/A       |
|    0     908   C+G    Insufficient Permissions           N/A       |
|    0    2340   C+G    C:\Windows\explorer.exe           N/A       |
|    0    4304   C+G    ...ration\Control Panel Client\nvcp | N/A       |
|    0    4624   C+G    ...dows.Cortana_cw5n1h2txyewy\Search | N/A       |
+-----+-----+-----+-----+-----+-----+
```

```
| 0 5116 C+G ...t_cw5n1h2txyewy\ShellExperienceHost.exe N/A |
```

3. Start **NVIDIA Control Panel** to verify that the NVIDIA driver is running.
 - a) Right-click on the desktop.
 - b) From the menu that opens, choose **NVIDIA Control Panel**.
4. In the **NVIDIA Control Panel**, from the **Help** menu, choose **System Information** to get information about the GPU.

NVIDIA Control Panel reports the GPU that is being used, its capabilities, and the NVIDIA driver version that is loaded.



1.3.2. Verifying the Creation of a GPU-Accelerated Virtual Workstation on a Linux VM Instance

List the GPUs in your GPU-accelerated virtual workstation by running the `nvidia-smi` command without any options.

The following example shows the output from `nvidia-smi` for a Linux VM instance configured with a single NVIDIA P100 GPU.

```
qwwsuser@nvidia-quadro-virtual-workstation-with-ubuntu-1-vm:~$ nvidia-smi
Wed Apr  3 02:01:30 2019
+-----+
| NVIDIA-SMI 410.92          Driver Version: 410.92          CUDA Version: 10.0     |
+-----+
| GPU  Name                    Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+
|  0   Tesla P100-PCIE...    On          | 00000000:00:04.0 Off  |                0 |
| N/A   33C    P0       26W / 250W | 0MiB / 16280MiB |           0%      Default |
+-----+
+-----+
| Processes:
| GPU      PID   Type   Process name                               GPU Memory
+-----+
| No running processes found
+-----+
```

```
qvwsuser@nvidia-quadro-virtual-workstation-with-ubuntu-1-vm:~$
```

1.4. Trying Out your GPU-Accelerated Virtual Workstation

Some third-party application vendors have endorsed the use of some of their apps with NVIDIA RTX Virtual Workstation on Google Cloud Platform. If you can install an endorsed app in your VM, you can use the app to try out your cloud-based virtual workstation.

1.4.1. Trying Out your GPU-Accelerated Virtual Workstation with Adobe Premiere Pro

Adobe has endorsed the use of the **Adobe Premiere Pro** video editing app with NVIDIA RTX Virtual Workstation on Google Cloud Platform. If you can install the **Adobe Premiere Pro** app in your VM, you can use the app to try out your cloud-based virtual workstation.

1. Open the sample project in **Adobe Premiere Pro**.
2. Scrub through the video and see that the video continues to play smoothly and without stutter.
3. Adjust colors, contrast, and shadows and notice that the adjustments are applied instantaneously.
4. Play the video and see that playback is remarkably smooth with no dropped frames.

1.4.2. Trying Out your GPU-Accelerated Virtual Workstation with Esri ArcGIS Pro

Esri has endorsed the use of the **Esri ArcGIS Pro** app with NVIDIA RTX Virtual Workstation on Google Cloud Platform. If you can install the **Esri ArcGIS Pro** app in your VM, you can use the app to try out your cloud-based virtual workstation.

1. Open **Esri ArcGIS Pro**.
2. View geographic data and notice the responsiveness of the app.
3. Use **Detect Objects Using Deep Learning** to detect palm trees.
The geoprocessing feature provided by Esri deep learning tools uses inferencing to detect palm trees.

With NVIDIA RTX Virtual Workstation images from the Google Cloud Platform marketplace powered by NVIDIA GPUs, you can run complex graphics visualization applications with AI-enhanced features from anywhere.

Appendix A.

RESOURCES FOR NVIDIA RTX VIRTUAL WORKSTATION ON GOOGLE CLOUD PLATFORM

- ▶ *NVIDIA RTX Virtual Workstation on Google Cloud Platform Release Notes*
- ▶ NVIDIA RTX Virtual Workstation listings on Google Cloud Platform marketplace:
 - ▶ NVIDIA RTX Virtual Workstation - Windows Server 2019
 - ▶ NVIDIA RTX Virtual Workstation - Windows Server 2016
 - ▶ NVIDIA RTX Virtual Workstation - Ubuntu 18.04
- ▶ Google Cloud Platform documentation:
 - ▶ GPUs on Compute Engine
 - ▶ Regions and Zones
 - ▶ Machine Types
 - ▶ Creating a VM Instance with a Custom Machine Type
 - ▶ Storage Options
 - ▶ Virtual Private Cloud (VPC) Network Overview
 - ▶ Creating Passwords for Windows Instances
 - ▶ Connecting to Linux instances

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.

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.

OpenCL

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

Trademarks

NVIDIA, the NVIDIA logo, NVIDIA GRID, NVIDIA GRID vGPU, NVIDIA Maxwell, NVIDIA Pascal, NVIDIA RTX, NVIDIA Turing, NVIDIA Volta, and Tesla are trademarks 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

© 2022 NVIDIA Corporation. All rights reserved.