



RTX Virtual Workstation Cloud on Amazon Web Services Elastic Compute Cloud Quick Start Guide

NVIDIA Corporation

Oct 02, 2026

Contents

1	Creating and Using an NVIDIA RTX Virtual Workstation Instance from the Amazon Web Services Elastic Compute Cloud 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	6
2	Resources for NVIDIA RTX Virtual Workstation on Amazon Web Services Elastic Compute Cloud	9
3	Notices	11

Chapter 1. Creating and Using an NVIDIA RTX Virtual Workstation Instance from the Amazon Web Services Elastic Compute Cloud 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 Amazon Web Services Elastic Compute Cloud marketplace as follows:

- ▶ Ensure that you have a Amazon Web Services Elastic Compute Cloud account and an active subscription.
- ▶ Decide the EC2 instance on which you want to deploy your GPU-accelerated virtual workstation. Instances are grouped according to the NVIDIA GPU with which the instances are configured. For details, refer to the following table, which links to information about the instances on the Amazon Web Services Elastic Compute Cloud website.

NVIDIA GPU	Amazon Web Services Elastic Compute Cloud Instances
NVIDIA L4	Amazon EC2 G6 Instances
NVIDIA A10G	Amazon EC2 G5 Instances
NVIDIA T4	Amazon EC2 G4 Instances

1. List the available [NVIDIA RTX Virtual Workstation AMIs](#).
 1. In a web browser, sign in to [AWS Marketplace \(https://aws.amazon.com/marketplace/\)](https://aws.amazon.com/marketplace/).
 2. In the search box, type `rtx`.
 3. Refine the results by following the **Business Applications** link and selecting **NVIDIA** as the publisher.
2. Go to the listing NVIDIA RTX Virtual Workstation for the AMI that you want to use, for example, **NVIDIA RTX Virtual Workstation - WinServer 2025**.
3. Click **View purchase options**.
4. On the Subscribe to NVIDIA RTX Virtual Workstation page that opens, review the purchase options and click **Subscribe**.
5. When the subscription is complete, click **Launch your software**.
6. On the Launch NVIDIA RTX Virtual Workstation page that opens, select the **Launch method** option that you want.

The only **Service** option available is **Amazon EC2**.

- If you want full control over the configuration of your instance, select **Launch from EC2 Console**.

If you select this option, the **Launch from EC2** button is added to this page.

- If you want a quick deployment with minimal configuration, select **One-click launch from AWS Marketplace**.

If you select this option, the options to enable you to configure the instance are added to this page.

7. From the **Region** drop-down list, select the region in which you want the instance to be launched.
8. Configure and launch the instance.
 - If you selected **Launch from EC2 Console**, click **Launch from EC2** and in the Launch an Instance page that opens, configure and launch the instance.
For more information, refer to [Launch an Amazon EC2 instance from an AWS Marketplace AMI](#) in the Amazon Elastic Compute Cloud documentation.
 - If you selected **One-click launch from AWS Marketplace**, use the options that were added to this page to configure the instance and click **Launch**.

When prompted to create a key pair or add a key pair to the set of keys authorized for this instance, ensure that the following conditions are met:

- The key pair type is RSA.
- The private key file is in `.pem` format.
- You have a copy of the private key file.

Otherwise, you will not be able to log in to the instance.

When the instance is launched, a confirmation page is displayed with a link to the instance on the EC2 Console.

9. In the launch confirmation page, click **View instance on EC2** to confirm that the instance is running and that all status checks are completed.
10. If you are creating a Windows instance and want to connect to the instance through RDP, edit the instance's security group to add an inbound RDP security rule to open port 3389 for TCP connections.

Follow the instructions for your chosen OS to connect to the instance:

- [Connecting to a GPU-Accelerated Virtual Workstation on a Windows VM Instance](#)
- [Connect to your Linux instance](#) in *Amazon EC2 User Guide for Linux Instances*

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

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

Note

For instructions for a Linux VM, refer to [Connect to your Linux instance](#) in *Amazon EC2 User Guide for Linux Instances*.

1. In the list of instances on the EC2 Dashboard, select the instance and click **Connect**.
2. In the **Connect** window that opens, download the RDP file and get the password for connecting to your instance.
 1. Click the **RDP client** tab.
 2. Set the **Connection Type** option to **Connect using RDP client**.
 3. Click **Download remote desktop file**.
 4. Click **Get password**.
 5. In the **Get Windows password** window, provide the private key file and click **Decrypt password**.
3. When the remote desktop file is downloaded, double-click the file to start a Remote Desktop Connection session on the VM.
4. If you are warned that the publisher of the remote connection cannot be identified and are asked about whether to connect anyway, click **Connect**.
5. When you are prompted, log in to the VM with the user name Administrator and the password that you got earlier.

Note

If you are logging in from a computer that is joined to a domain, specify the user name as `~.\administrator`. In this way, you ensure that you log in as Administrator on the VM where the Remote Desktop Connection session is running. Otherwise, Windows automatically prefixes the user name with the name of the domain to which your computer is joined and attempts to log in there.

6. If you are warned that the publisher of the remote connection cannot be identified and are asked about whether to connect anyway, click **Yes**.
7. Close the **Connect to Your Instance** window.

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:\Users\Administrator>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 L4 GPU.

```
C:\Program Files\NVIDIA Corporation\NVSMI> nvidia-smi
Fri Jan  9 22:41:58 2026
```

```
+-----+
| NVIDIA-SMI 581.42                Driver Version: 581.42                CUDA Version: 13.0
```

(continues on next page)

(continued from previous page)

```

→ |
|-----+-----+-----+
→--+
| GPU Name Driver-Model | Bus-Id Disp.A | Volatile Uncorr.
→ECC |
| Fan Temp Perf Pwr:Usage/Cap| Memory-Usage | GPU-Util Compute M.
→ |
| | | | | | |
→ | | | | | | |
|-----+-----+-----+
| 0 NVIDIA L4 WDDM |00000000:00:1E.0 Off |
→0 |
| N/A 29C P8 14W / 70W | 0MiB / 23028MiB | 2%
→Default |
| | | | | | |
→A |
|-----+-----+-----+
→--+

+-----+-----+-----+
→--+
| Processes:
→ |
| GPU GI CI PID Type Process name GPU
→Memory | | | | | | |
| ID ID | | | | | |
→ | | | | | | |
|-----+-----+-----+
| 0 N/A N/A 660 C+G C:\Windows\System32\WUDFHost.exe N/A
→ |
| 0 N/A N/A 860 C+G C:\Windows\System32\dwm.exe N/A
→ |
| 0 N/A N/A 964 C+G ...indows\System32\ShellHost.exe N/A
→ |
| 0 N/A N/A 1220 C+G C:\Windows\System32\dwm.exe N/A
→ |
| 0 N/A N/A 5272 C+G ..._cw5n1h2txyewy\SearchHost.exe N/A
→ |
| 0 N/A N/A 5800 C+G ...2txyewy\CrossDeviceResume.exe N/A
→ |
| 0 N/A N/A 5844 C+G C:\Windows\explorer.exe N/A
→ |
| 0 N/A N/A 5932 C+G C:\Windows\System32\LogonUI.exe N/A
→ |
| 0 N/A N/A 7020 C+G ....0.3650.96\msedgewebview2.exe N/A
→ |
| 0 N/A N/A 7672 C+G ...y\StartMenuExperienceHost.exe N/A
→ |
| 0 N/A N/A 9872 C+G ...yb3d8bbwe\WindowsTerminal.exe N/A
→ |
|-----+-----+-----+
→--+

```


(continues on next page)

(continued from previous page)

```
C:\Program Files\NVIDIA Corporation\NVSMI>
```

3. Start **NVIDIA Control Panel** to verify that the NVIDIA driver is running.
 1. Right-click on the desktop.
 2. 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 A10G GPU.

```
~$ nvidia-smi
Fri Jan  9 23:21:53 2026
+-----+
| NVIDIA-SMI 580.95.05      Driver Version: 580.95.05      CUDA Version: 13.0
|
+-----+-----+-----+
| GPU   Name                Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC
| Fan   Temp   Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M.
|
|                               |                      |              MIG M.
+-----+-----+-----+
```

(continues on next page)

(continued from previous page)

```

→ |
|=====+=====+=====|
| 0 NVIDIA A10G          On | 00000000:00:1E.0 Off | 0
→ |
| 0% 24C P8 11W / 300W | 0MiB / 23028MiB | 0% Default
→ |
| | | | |
→ |
+-----+-----+-----+
→+
+-----+-----+-----+
→+
| Processes:
→ |
| GPU  GI  CI          PID  Type  Process name          GPU Memory
→ |
|      ID  ID                                   Usage
→ |
|=====+=====+=====|
| No running processes found
→ |
+-----+-----+-----+
→+
:~$

```

Chapter 2. Resources for NVIDIA RTX Virtual Workstation on Amazon Web Services Elastic Compute Cloud

- ▶ [NVIDIA RTX Virtual Workstation on Amazon Web Services Elastic Compute Cloud Release Notes](#)
- ▶ [NVIDIA RTX Virtual Workstation AMIs](#)
- ▶ Amazon Web Services Elastic Compute Cloud documentation:
 - ▶ [Launch an Amazon EC2 instance from an AWS Marketplace AMI](#)
 - ▶ [Connect to your Windows instance](#)
 - ▶ [Connect to your Linux instance](#)

Chapter 3. Notices

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 Turing, NVIDIA Volta, GPUDirect, Quadro, 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

©2019-2026, NVIDIA Corporation