



RTX VIRTUAL WORKSTATION CLOUD ON ORACLE CLOUD INFRASTRUCTURE

DU-09942-001 _v10 | October 2022

Quick Start Guide



TABLE OF CONTENTS

Chapter 1. Creating and Using an NVIDIA RTX Virtual Workstation Instance from the Oracle Cloud Infrastructure Marketplace.....	1
1.1. Creating a GPU-Accelerated Virtual Workstation.....	1
1.1.1. Creating a Virtual Cloud Network (VCN).....	2
1.1.2. Launching an NVIDIA RTX Virtual Workstation Instance.....	2
1.2. Connecting to a GPU-Accelerated Virtual Workstation.....	3
1.3. Verifying the Creation of your GPU-Accelerated Virtual Workstation.....	4
Appendix A. Resources for NVIDIA RTX Virtual Workstation on Oracle Cloud Infrastructure... 	6

Chapter 1.

CREATING AND USING AN NVIDIA RTX VIRTUAL WORKSTATION INSTANCE FROM THE ORACLE CLOUD INFRASTRUCTURE 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 Oracle ensures that the image is always up to date with the latest 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.

1.1. Creating a GPU-Accelerated Virtual Workstation

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

- ▶ Ensure that you have a Oracle Cloud Infrastructure account and an active subscription.
- ▶ Decide the Oracle Cloud Infrastructure instance shape on which you want to deploy your GPU-accelerated virtual workstation.

For information about the Oracle Cloud Infrastructure instance shapes that support NVIDIA RTX Virtual Workstation, see [NVIDIA RTX Virtual Workstation on Oracle Cloud Infrastructure Release Notes](#).

Creating a GPU-accelerated virtual workstation involves following this sequence of instructions:

1. [Creating a Virtual Cloud Network \(VCN\)](#)

2. Launching an NVIDIA RTX Virtual Workstation Instance

1.1.1. Creating a Virtual Cloud Network (VCN)

1. In a web browser, sign in to the Oracle Cloud Infrastructure Console as explained in [Signing In to the Console](#) in the Oracle Cloud Infrastructure documentation.
2. Select a region where GPUs are available, for example, **Germany Central (Frankfurt)**.
3. Under **Quick Actions**, click **Set up a network with a wizard**.
4. Select **VCN with Internet Connectivity** and click **Start VCN Wizard**.
5. Provide a name for the VCN and select a compartment.
6. Review the default settings to see if they meet your requirements, change any settings as required, and click **Next**.
7. Review your choices and click **Create**.

1.1.2. Launching an NVIDIA RTX Virtual Workstation Instance

1. Go to the NVIDIA RTX Virtual Workstation listing for the guest OS that you want to use:
 - ▶ [NVIDIA RTX Virtual Workstation - Windows Server 2016 VM](#)
2. In the upper right corner of the page, click **Get App**.
The **Install Application** page opens.
3. From the **Select OCI Region** list, select the region where you created your VCN and click **Sign In**.
4. Select the latest image version, for example, **2020.04.03-1-vdws-10.1**, select your compartment, confirm that you have reviewed the terms of use, and click **Launch Instance**.
5. Specify the details required to create the instance.

Image Source

Select **NVIDIA GPU Cloud Machine Image**.

Availability Domain

Select an availability domain where GPUs are available, for example, **EU-FRANKFURT-1-AD-3**.

Instance Type

Select **Virtual Machine** or **Bare Metal**.

Instance Shape

Select one of the available instance shapes listed in [NVIDIA RTX Virtual Workstation on Oracle Cloud Infrastructure Release Notes](#).

Configure Networking

Select the correct compartment and the VCN and subnets that you created.

For details, see [Connecting to a Windows Instance](#) in the Oracle Cloud Infrastructure documentation.

6. If you want to connect to the instance through RDP, add a stateful ingress security rule for TCP traffic on destination port 3389 from source 0.0.0.0/0 and any source port.

You can add this security rule either to a network security group to which the Windows instance belongs or to a security list that is used by the instance's subnet.

For more information, see the following topics in the Oracle Cloud Infrastructure documentation:

- ▶ [Security Rules](#)
- ▶ [Network Security Groups](#)
- ▶ [Security Lists](#)

7. Optional: If necessary, add block storage to provide additional space for your workloads.

For details, see [Adding a Block Volume](#) in the Oracle Cloud Infrastructure documentation.

8. In preparation for connecting to your instance, note the public IP address of the instance from the Oracle Cloud Infrastructure console.

Connect to your instance as explained in [Connecting to a GPU-Accelerated Virtual Workstation](#).

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

1. On your local computer, start a Remote Desktop client, for example, **Remote Desktop Connection** on Windows.
2. In the **Computer** field, type the public IP address of the instance
If you don't already have the public IP address of the instance, you can get it from the Oracle Cloud Infrastructure console.
3. In the **User name** field, type the user name **opc**.
Some Remote Desktop clients require you to connect to the instance before you can provide the user name.
4. Click **Connect**.
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**.
6. When prompted, type your password and press **Enter**.

- ▶ If you are connecting to the instance for the first time, type the initial password that was provided to you by Oracle Cloud Infrastructure when you launched the instance.

You are prompted to change the password as soon as you log in. Your new password must be at least 12 characters long and meet the requirements described in [Password must meet complexity requirements](#) on the Microsoft documentation site.

- ▶ If you have previously connected to the instance and created a password, type the password that you created.
- ▶ If you are using a custom image, type the appropriate password, which might be the password for the instance from which the image was created.

For details about Windows custom images, see [Creating Windows Custom Images](#) in the Oracle Cloud Infrastructure documentation.

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.

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

```
C:\Users\opc>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 one NVIDIA P100 GPU.

```
C:\Program Files\NVIDIA Corporation\NVSMI>nvidia-smi
Tue Apr 21 00:50:57 2020
+-----+
| NVIDIA-SMI 432.33 Driver Version: 432.33 CUDA Version: 10.1 |
+-----+-----+
| 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 P100-SXM2... WDDM | 00000000:00:04.0 Off | 0 |
| N/A 34C P0 31W / 300W | 340MiB / 16384MiB | 2% Default |
+-----+-----+-----+
+-----+
| Processes: GPU Memory | | GPU PID Type Process name Usage |
+-----+-----+-----+
| 0 452 C+G Insufficient Permissions N/A |
| 0 1116 C+G ...t_cw5n1h2txyewy\ShellExperienceHost.exe N/A |
+-----+-----+-----+
```


Appendix A.

RESOURCES FOR NVIDIA RTX VIRTUAL WORKSTATION ON ORACLE CLOUD INFRASTRUCTURE

- ▶ *NVIDIA RTX Virtual Workstation on Oracle Cloud Infrastructure Release Notes*
- ▶ NVIDIA RTX Virtual Workstation listings on Oracle Cloud Infrastructure marketplace:
 - ▶ NVIDIA RTX Virtual Workstation - Windows Server 2016 VM
- ▶ Oracle Cloud Infrastructure documentation:
 - ▶ Signing In to the Console
 - ▶ Connecting to a Windows Instance
 - ▶ Security Rules
 - ▶ Network Security Groups
 - ▶ Security Lists
 - ▶ Adding a Block Volume
 - ▶ Creating Windows Custom Images

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.