



VIRTUAL GPU CLIENT LICENSING

DU-07757-001 _v10.0 through 10.1 | February 2020

User Guide



TABLE OF CONTENTS

Chapter 1. Introduction to NVIDIA vGPU Software Licensing.....	1
1.1. NVIDIA vGPU Software Licensed Products.....	1
1.2. How NVIDIA vGPU Software Licensing Is Enforced.....	2
1.3. How this Guide Is Organized.....	4
1.4. Enabling License Management in NVIDIA X Server Settings.....	5
Chapter 2. Licensing an NVIDIA vGPU.....	6
2.1. NVIDIA vGPU License Requirements.....	6
2.2. Licensing an NVIDIA vGPU on Windows.....	7
2.3. Licensing an NVIDIA vGPU on Linux.....	9
Chapter 3. Licensing a Physical GPU.....	11
3.1. Licensing a Physical GPU for Quadro vDWS on Windows.....	11
3.2. Removing a Quadro vDWS License from a Physical GPU on Windows.....	13
3.3. Licensing a Physical GPU for Quadro vDWS or vComputeServer on Linux.....	13
Chapter 4. Advanced Topics.....	17
4.1. Setting the Client Host Identifier for License Checkouts.....	17
4.2. Licenses Obtained After Boot.....	18
4.3. Operating with Intermittent Connectivity to the License Server.....	18
4.4. Applying Windows License Settings Through the Windows Registry.....	18
Chapter 5. Troubleshooting.....	22
5.1. Known Issues.....	22
5.2. Licensing Event Logs.....	22
5.3. Troubleshooting Steps.....	24
Appendix A. Licensing NVIDIA vGPU Software on Linux by Using a Configuration File.....	25
A.1. Licensing an NVIDIA vGPU on Linux by Using a Configuration File.....	25
A.2. Licensing a Physical GPU on Linux by Using a Configuration File.....	27

LIST OF FIGURES

Figure 1	Software Enforcement of NVIDIA vGPU Software Licensing	4
Figure 2	Managing vGPU licensing in NVIDIA Control Panel	8
Figure 3	Managing Quadro vDWS Licensing in NVIDIA Control Panel	12
Figure 4	Managing Quadro vDWS Licensing in NVIDIA X Server Settings	15
Figure 5	Configuring vGPU Licensing Through Windows Registry Settings	21

LIST OF TABLES

Table 1	NVIDIA vGPU Software Licensed Products	1
Table 2	NVIDIA vGPU Software License Enforcement by Deployment Type	2
Table 3	Licensing Registry Settings	19

Chapter 1.

INTRODUCTION TO NVIDIA vGPU SOFTWARE LICENSING

The following products are available as licensed products on NVIDIA GPUs that support NVIDIA vGPU software:

- ▶ NVIDIA® Quadro® Virtual Data Center Workstation (Quadro vDWS)
- ▶ NVIDIA vComputeServer
- ▶ NVIDIA GRID™ Virtual PC
- ▶ GRID Virtual Applications

This guide describes these licensed products and how to enable and use them on supported hardware.

1.1. NVIDIA vGPU Software Licensed Products

The available NVIDIA vGPU software licensed products support different classes of NVIDIA vGPU software deployments.

Table 1 NVIDIA vGPU Software Licensed Products

NVIDIA vGPU Software Licensed Product	Target Users	Supported NVIDIA vGPU Software Deployments
GRID Virtual Applications	Users of PC-level applications and server-based desktops that use Citrix Virtual Apps and Desktops, VMware Horizon, RDSH, or other app streaming or session-based solutions	<ul style="list-style-type: none">▶ A-series NVIDIA vGPUs▶ GPU pass through▶ Microsoft DDA▶ VMware vDGA▶ Bare metal
GRID Virtual PC	Users of business virtual desktops who require a great user experience with PC applications for Windows, web browsers, and high-definition video	<ul style="list-style-type: none">▶ B-series NVIDIA vGPUs▶ Microsoft RemoteFX vGPU▶ VMware vSGA

NVIDIA vGPU Software Licensed Product	Target Users	Supported NVIDIA vGPU Software Deployments
vComputeServer	Users of compute-intensive virtual servers for artificial intelligence (AI), deep learning, or high-performance computing (HPC) workloads	<ul style="list-style-type: none"> ▶ C-series NVIDIA vGPUs ▶ GPU pass through ▶ Microsoft DDA ▶ VMware vDGA
Quadro vDWS	Users of mid-range and high-end workstations who require access to remote professional graphics applications with full performance on any device anywhere	<ul style="list-style-type: none"> ▶ Q-series NVIDIA vGPUs ▶ C-series NVIDIA vGPUs ▶ B-series NVIDIA vGPUs ▶ GPU pass through ▶ Microsoft DDA ▶ VMware vDGA ▶ Bare metal

- ▶ Microsoft DDA and VMware vDGA are functionally equivalent to GPU pass through.
- ▶ Microsoft RemoteFX vGPU and VMware vSGA are graphics virtualization technologies.
- ▶ Microsoft DDA requires Windows Server 2016 with the Hyper-V role.

1.2. How NVIDIA vGPU Software Licensing Is Enforced

Depending on the NVIDIA vGPU software deployment, licensing is enforced either through software or only through the end-user license agreement (EULA).

Table 2 NVIDIA vGPU Software License Enforcement by Deployment Type

NVIDIA vGPU Software Deployment	Required NVIDIA vGPU Software License	Enforcement
A-series NVIDIA vGPU	GRID Virtual Applications	Software See Note (1).
B-series NVIDIA vGPU	GRID Virtual PC or Quadro vDWS	Software
C-series NVIDIA vGPU	vComputeServer or Quadro vDWS	Software See Note (2).
Q-series NVIDIA vGPU	Quadro vDWS	Software See Note (3).
GPU pass through for workstation or professional 3D graphics	Quadro vDWS	Software
GPU pass through for compute-intensive virtual servers	vComputeServer	Software See Note (4).
GPU pass through for PC-level applications	GRID Virtual Applications	EULA only

NVIDIA vGPU Software Deployment	Required NVIDIA vGPU Software License	Enforcement
Microsoft DDA for workstation or professional 3D graphics	Quadro vDWS	Software
Microsoft DDA for compute-intensive virtual servers	vComputeServer	Software See Note (4).
Microsoft DDA for PC-level applications	GRID Virtual Applications	EULA only
VMware vDGA for workstation or professional 3D graphics	Quadro vDWS	Software
VMware vDGA for compute-intensive virtual servers	vComputeServer	Software See Note (4).
VMware vDGA for PC-level applications	GRID Virtual Applications	EULA only
Bare metal for workstation or professional 3D graphics	Quadro vDWS	Software
Bare metal for PC-level applications	GRID Virtual Applications	EULA only
Microsoft RemoteFX vGPU	GRID Virtual PC	EULA only
VMware vSGA	GRID Virtual PC	EULA only



1. For A-series NVIDIA vGPU deployments, one GRID Virtual Applications license is required for each concurrent user, but licensing for only one of these users is enforced through software. Licensing for the remaining users is enforced through the EULA.
2. For C-series NVIDIA vGPU deployments, one license per vGPU assigned to a VM is enforced through software. This license is valid for up to eight vGPU instances on a single GPU or for the assignment to a VM of one vGPU that is assigned all the physical GPU's frame buffer. When multiple C-series vGPUs are assigned to a single VM, one license for each vGPU assigned to the VM is required. One license is enforced through software. The remaining licenses are enforced through the EULA.
3. When multiple Q-series vGPUs are assigned to a single VM, only one license for the VM is required, **not** one license for each vGPU assigned to the VM.
4. For GPU pass through, Microsoft DDA, and VMware vDGA deployments for compute-intensive virtual servers, one license per VM is enforced through software. This license is valid for the assignment of one physical GPU to a VM. One additional license is required for each additional physical GPU assigned to a VM. Any additional licenses are enforced through the EULA.

Software Enforcement

When licensing is enforced through software, NVIDIA vGPU software runs at reduced capability until the Virtual Machine (VM) that is running the software obtains a license. A VM obtains a license over the network from an NVIDIA vGPU software license server. The license is “checked out” or “borrowed” when the VM is booted, and returned when the VM is shut down.

For NVIDIA vGPU deployments, the NVIDIA vGPU software automatically selects the correct type of license based on the vGPU type.



Information on setting up and running the NVIDIA vGPU software license server is provided in *Virtual GPU License Server User Guide*.

Figure 1 illustrates how NVIDIA vGPU software licensing is enforced through software.




Figure 1 Software Enforcement of NVIDIA vGPU Software Licensing

EULA-Only Enforcement

When licensing is enforced only through the EULA, no licenses are checked out from the license server. You **must** comply with the terms of the EULA by ensuring that **all** the licenses required for your NVIDIA vGPU software deployment have been purchased.

1.3. How this Guide Is Organized

The remainder of this guide is organized as follows:

- ▶ [Licensing an NVIDIA vGPU](#) describes licensing of NVIDIA vGPU.
- ▶ [Licensing a Physical GPU](#) describes licensing of Quadro vDWS features for a physical GPU in pass-through mode or in a bare-metal deployment.
- ▶ [Advanced Topics](#) discusses advanced licensing settings.
- ▶ [Troubleshooting](#) provides guidance on troubleshooting.

1.4. Enabling License Management in NVIDIA X Server Settings

By default, the **Manage License** option is not available in **NVIDIA X Server Settings**. Before using **NVIDIA X Server Settings** to license NVIDIA vGPU software, you must enable this option.



Do not enable the **Manage License** option with Red Hat Enterprise Linux 6.8 and 6.9 or CentOS 6.8 and 6.9. To prevent a segmentation fault in Dbus code from causing the `nvidia-gridd` service from exiting, the GUI for licensing must be disabled with these OS versions.

If you are licensing a physical GPU for vComputeServer, you **must** use the configuration file `/etc/nvidia/gridd.conf`.

This task requires `sudo` privileges.

1. If **NVIDIA X Server Settings** is running, shut it down.
2. If the `/etc/nvidia/gridd.conf` file does not already exist, create it by copying the supplied template file `/etc/nvidia/gridd.conf.template`.
3. As root, edit the `/etc/nvidia/gridd.conf` file to set the `EnableUI` option to `TRUE`.
4. Start the `nvidia-gridd` service.

```
# sudo service nvidia-gridd start
```

When **NVIDIA X Server Settings** is restarted, the **Manage License** option is now available.

Chapter 2.

LICENSING AN NVIDIA vGPU

NVIDIA vGPU is a licensed product. When booted on a supported GPU, a vGPU runs at reduced capability until a license is acquired.

The performance of an unlicensed vGPU is restricted as follows:

- ▶ Frame rate is capped at 3 frames per second.
- ▶ GPU resource allocations are limited, which will prevent some applications from running correctly.
- ▶ On vGPUs that support CUDA, CUDA is disabled.

These restrictions are removed when a license is acquired.

After you license NVIDIA vGPU, the VM that is set up to use NVIDIA vGPU is capable of running the full range of DirectX and OpenGL graphics applications.

If licensing is configured, the virtual machine (VM) obtains a license from the license server when a vGPU is booted on these GPUs. The VM retains the license until it is shut down. It then releases the license back to the license server. Licensing settings persist across reboots and need only be modified if the license server address changes, or the VM is switched to running GPU pass through.

2.1. NVIDIA vGPU License Requirements

The license edition that you require to use NVIDIA vGPU depends on the vGPU type.

The Quadro vDWS license edition includes the GRID Virtual PC license edition: virtual GPUs that require a GRID Virtual PC license can also be used with a Quadro vDWS license.



For more information on NVIDIA vGPU types, refer to *Virtual GPU Software User Guide*.

2.2. Licensing an NVIDIA vGPU on Windows

Perform this task from the guest VM to which the vGPU is assigned.

The **NVIDIA Control Panel** tool that you use to perform this task detects that a vGPU is assigned to the VM and, therefore, provides no options for selecting the license type. After you license the vGPU, NVIDIA vGPU software automatically selects the correct type of license based on the vGPU type.

1. Open **NVIDIA Control Panel**:
 - ▶ Right-click on the Windows desktop and select **NVIDIA Control Panel** from the menu.
 - ▶ Open **Windows Control Panel** and double-click the **NVIDIA Control Panel** icon.
2. In **NVIDIA Control Panel**, select the **Manage License** task in the **Licensing** section of the navigation pane.



If the **Licensing** section and **Manage License** task are not displayed in **NVIDIA Control Panel**, the system has been configured to hide licensing controls in **NVIDIA Control Panel**. For information about registry settings, see [Applying Windows License Settings Through the Windows Registry](#).

The **Manage License** task pane shows that NVIDIA vGPU is currently unlicensed.

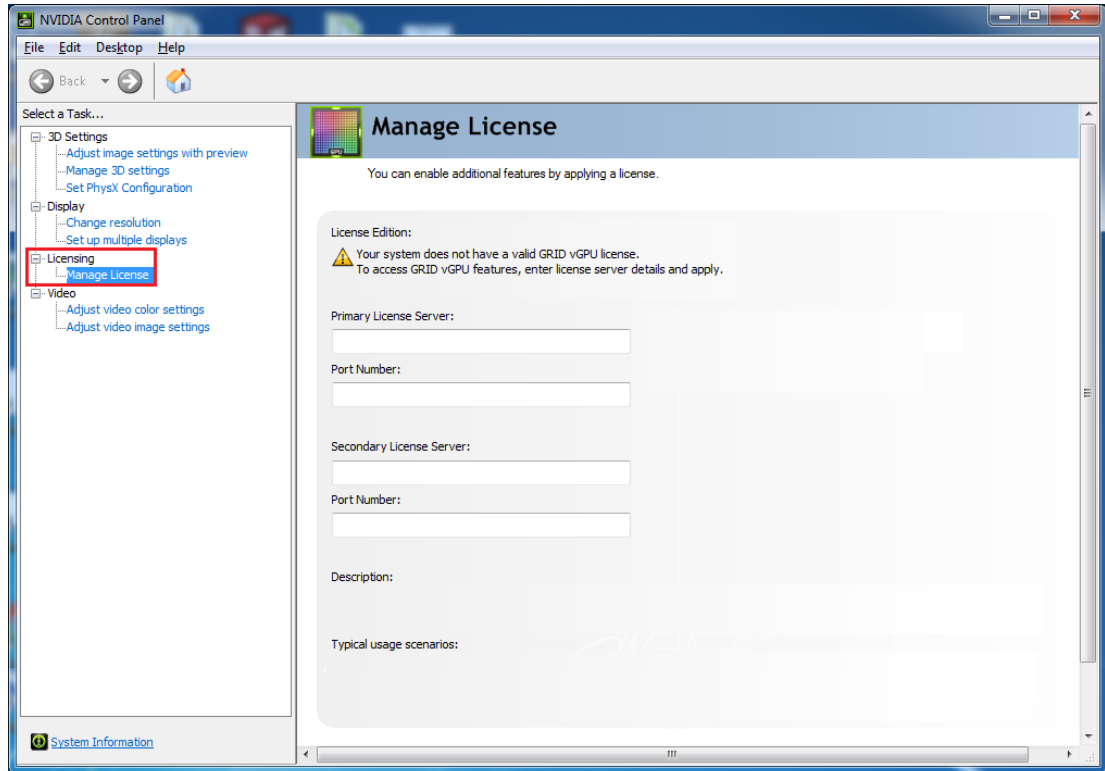


Figure 2 Managing vGPU licensing in NVIDIA Control Panel

3. In the **Primary License Server** field, enter the address of your primary NVIDIA vGPU software License Server.
 The address can be a fully-qualified domain name such as `gridlicense1.example.com`, or an IP address such as `10.31.20.45`.
 If you have only one license server configured, enter its address in this field.
4. Leave the **Port Number** field under the **Primary License Server** field unset.
 The port defaults to `7070`, which is the default port number used by NVIDIA vGPU software License Server.
5. In the **Secondary License Server** field, enter the address of your secondary NVIDIA vGPU software License Server.
 If you have only one license server configured, leave this field unset.
 The address can be a fully-qualified domain name such as `gridlicense2.example.com`, or an IP address such as `10.31.20.46`.
6. Leave the **Port Number** field under the **Secondary License Server** field unset.
 The port defaults to `7070`, which is the default port number used by NVIDIA vGPU software License Server.
7. Click **Apply** to assign the settings.
 The system requests the appropriate license for the current vGPU from the configured license server.

The vGPU within the VM should now exhibit full frame rate, resolution, and display output capabilities. The VM is now capable of running the full range of DirectX and OpenGL graphics applications.

If the system fails to obtain a license, see [Troubleshooting](#) for guidance on troubleshooting.

2.3. Licensing an NVIDIA vGPU on Linux

Perform this task from the guest VM to which the vGPU is assigned.

The **NVIDIA X Server Settings** tool that you use to perform this task detects that a vGPU is assigned to the VM and, therefore, provides no options for selecting the license type. After you license the vGPU, NVIDIA vGPU software automatically selects the correct type of license based on the vGPU type.

Ensure that the **Manage License** option is enabled as explained in [Enabling License Management in NVIDIA X Server Settings](#).



Do not enable the **Manage License** option with Red Hat Enterprise Linux 6.8 and 6.9 or CentOS 6.8 and 6.9. To prevent a segmentation fault in Dbus code from causing the `nvidia-gridd` service from exiting, the GUI for licensing must be disabled with these OS versions.

If you do not want to or cannot enable the **Manage License** option, use the the configuration file `/etc/nvidia/gridd.conf` as explained in [Licensing an NVIDIA vGPU on Linux by Using a Configuration File](#).

1. Start **NVIDIA X Server Settings** by using the method for launching applications provided by your Linux distribution.
For example, on Ubuntu Desktop, open the **Dash**, search for **NVIDIA X Server Settings**, and click the **NVIDIA X Server Settings** icon.
2. In the **NVIDIA X Server Settings** window that opens, click **Manage GRID License**.
The **License Edition** section of the **NVIDIA X Server Settings** window shows that NVIDIA vGPU is currently unlicensed.
3. In the **Primary Server** field, enter the address of your primary NVIDIA vGPU software License Server.
The address can be a fully-qualified domain name such as `gridlicense1.example.com`, or an IP address such as `10.31.20.45`.
If you have only one license server configured, enter its address in this field.
4. Leave the **Port Number** field under the **Primary Server** field unset.
The port defaults to `7070`, which is the default port number used by NVIDIA vGPU software License Server.
5. In the **Secondary Server** field, enter the address of your secondary NVIDIA vGPU software License Server.

If you have only one license server configured, leave this field unset.

The address can be a fully-qualified domain name such as `gridlicense2.example.com`, or an IP address such as `10.31.20.46`.

6. Leave the **Port Number** field under the **Secondary Server** field unset. The port defaults to 7070, which is the default port number used by NVIDIA vGPU software License Server.
7. Click **Apply** to assign the settings. The system requests the appropriate license for the current vGPU from the configured license server.

The vGPU within the VM should now exhibit full frame rate, resolution, and display output capabilities. The VM is now capable of running the full range of DirectX and OpenGL graphics applications.

If the system fails to obtain a license, see [Troubleshooting](#) for guidance on troubleshooting.

Chapter 3.

LICENSING A PHYSICAL GPU

To operate at full capability, a physical GPU running in pass-through mode or a bare-metal deployment requires a license. The type of license required depends on how the physical GPU is deployed.

In GPU pass-through mode on Windows, or in a bare-metal deployment on Windows or Linux, a physical GPU requires a Quadro vDWS license. In GPU pass-through mode on Linux, a physical GPU requires a Quadro vDWS license or a vComputeServer license.

You can also operate a physical GPU in pass-through mode or a bare-metal deployment at reduced capability with a GRID Virtual Apps license. GRID Virtual PC is not available on GPUs running in pass-through mode or bare-metal deployments.

A physical GPU with a Quadro vDWS license supports up to four virtual display heads at 4K resolution.

A physical GPU with a GRID Virtual Apps license supports only a single virtual display with a maximum resolution that depends on the GPU architecture:

- ▶ For GPUs based on the Pascal GPU architecture and later GPU architectures, the maximum resolution is 1280×1024.
- ▶ For GPUs based on the Maxwell GPU architecture, the maximum resolution is 2560×1600.


3.1. Licensing a Physical GPU for Quadro vDWS on Windows

Perform this task from the guest VM to which the GPU is passed through or from the bare-metal host.

The **NVIDIA Control Panel** tool that you use to perform this task detects that a GPU is passed through to the VM or deployed to the bare-metal host and, therefore, provides the options for selecting between the Quadro vDWS and GRID Virtual Apps license types.

1. Open **NVIDIA Control Panel**:

- ▶ Right-click on the Windows desktop and select **NVIDIA Control Panel** from the menu.
 - ▶ Open **Windows Control Panel** and double-click the **NVIDIA Control Panel** icon.
2. In **NVIDIA Control Panel**, select the **Manage License** task in the **Licensing** section of the navigation pane.

 If the **Licensing** section and **Manage License** task are not displayed in **NVIDIA Control Panel**, the GPU or driver you are using do not support NVIDIA vGPU software licensed features.

The **Manage License** task pane shows the current License Edition being used.

The default is GRID Virtual Apps.

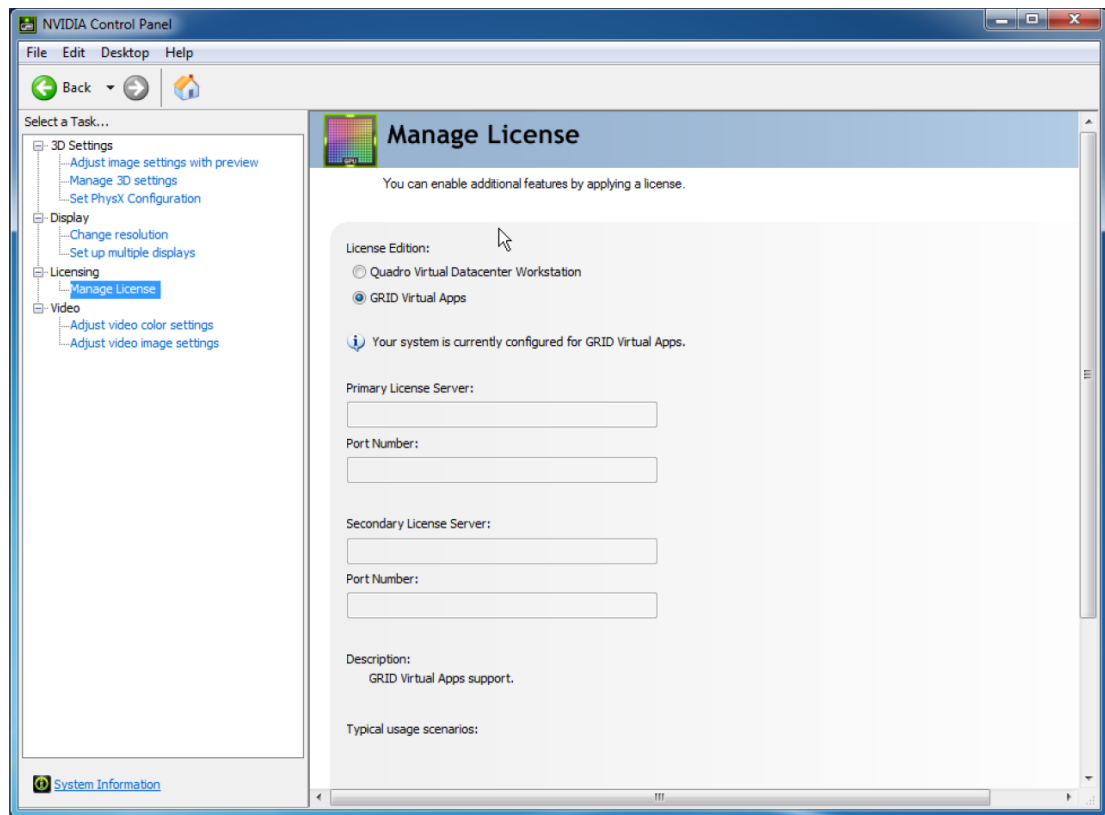


Figure 3 Managing Quadro vDWS Licensing in NVIDIA Control Panel

3. Select **Quadro vDWS**.
4. In the **Primary License Server** field, enter the address of your primary NVIDIA vGPU software License Server.

The address can be a fully-qualified domain name such as `gridlicense1.example.com`, or an IP address such as `10.31.20.45`.

If you have only one license server configured, enter its address in this field.

5. Leave the **Port Number** field under the **Primary License Server** field unset. It will default to 7070, which is the default port number used by NVIDIA vGPU software License Server.

6. In the **Secondary License Server** field, enter the address of your secondary NVIDIA vGPU software License Server.
If you have only one license server configured, leave this field unset.
The address can be a fully-qualified domain name such as `gridlicense2.example.com`, or an IP address such as `10.31.20.46`.
7. Leave the **Port Number** field under the **Secondary License Server** field unset.
The port defaults to 7070, which is the default port number used by NVIDIA vGPU software License Server.
8. Click **Apply** to assign the settings.
The **License Edition** section will indicate that your system is running in Quadro vDWS mode. The system will also request the appropriate license for the current GPU from the configured license server. If unsuccessful, the system warns you that it could not obtain a license. If necessary, see [Troubleshooting](#) for guidance on troubleshooting.

Once configured in **NVIDIA Control Panel**, licensing settings persist across reboots and need only be modified if the license server address changes, or the VM is switched to running NVIDIA vGPU.

3.2. Removing a Quadro vDWS License from a Physical GPU on Windows

1. Open **NVIDIA Control Panel**.
2. In the **Manage License** task, select **GRID Virtual Apps**.
3. Click **Apply**.

The setting does not take effect until the next time the system is shut down or rebooted. Quadro vDWS features remain available until then.

3.3. Licensing a Physical GPU for Quadro vDWS or vComputeServer on Linux

Perform this task from the guest VM to which the GPU is passed through or from the bare-metal host.

The **NVIDIA X Server Settings** tool that you use to perform this task detects that the GPU is passed through to the VM or deployed to the bare-metal host and, therefore, provides options for selecting between the following license types:

- ▶ Quadro vDWS
- ▶ NVIDIA vComputeServer
- ▶ GRID Virtual Applications

Ensure that the **Manage License** option is enabled as explained in [Enabling License Management in NVIDIA X Server Settings](#).



Do not enable the **Manage License** option with Red Hat Enterprise Linux 6.8 and 6.9 or CentOS 6.8 and 6.9. To prevent a segmentation fault in Dbus code from causing the `nvidia-gridd` service from exiting, the GUI for licensing must be disabled with these OS versions.

If you do not want to or cannot enable the **Manage License** option, use the configuration file `/etc/nvidia/gridd.conf`.

If you are licensing a physical GPU for vComputeServer, you **must** use the configuration file `/etc/nvidia/gridd.conf`.

For instructions, see [Licensing a Physical GPU on Linux by Using a Configuration File](#).

1. Start **NVIDIA X Server Settings** by using the method for launching applications provided by your Linux distribution.

For example, on Ubuntu Desktop, open the **Dash**, search for **NVIDIA X Server Settings**, and click the **NVIDIA X Server Settings** icon.

2. In the **NVIDIA X Server Settings** window that opens, click **Manage GRID License**.

The **License Edition** section of the **NVIDIA X Server Settings** window shows the current License Edition being used.

The default is GRID Virtual Applications.

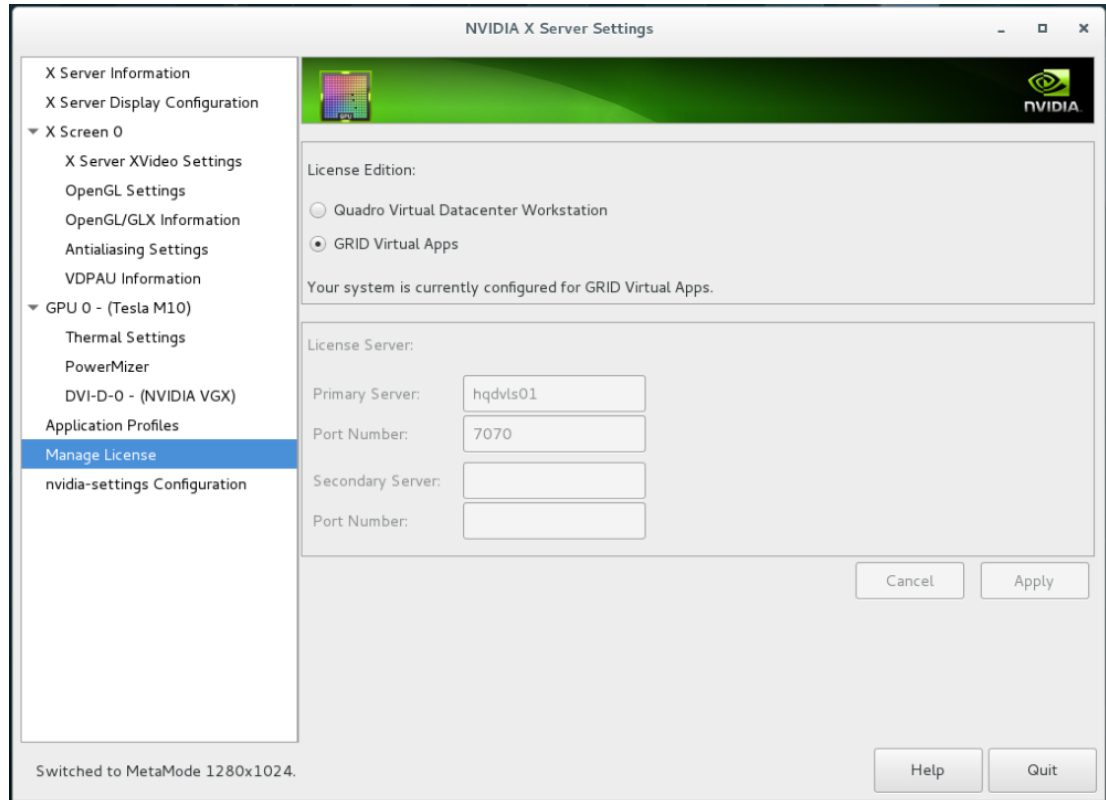


Figure 4 Managing Quadro vDWS Licensing in NVIDIA X Server Settings

3. Select the license type that you want.
 - ▶ **Quadro vDWS**
 - ▶ **NVIDIA vComputeServer**
4. In the **Primary Server** field, enter the address of your primary NVIDIA vGPU software License Server.
 The address can be a fully-qualified domain name such as `gridlicense1.example.com`, or an IP address such as `10.31.20.45`.
 If you have only one license server configured, enter its address in this field.
5. Leave the **Port Number** field under the **Primary Server** field unset.
 The port defaults to `7070`, which is the default port number used by NVIDIA vGPU software License Server.
6. In the **Secondary Server** field, enter the address of your secondary NVIDIA vGPU software License Server.
 If you have only one license server configured, leave this field unset.
 The address can be a fully-qualified domain name such as `gridlicense2.example.com`, or an IP address such as `10.31.20.46`.
7. Leave the **Port Number** field under the **Secondary Server** field unset.
 The port defaults to `7070`, which is the default port number used by NVIDIA vGPU software License Server.
8. Click **Apply** to assign the settings.

The **License Edition** section indicates that your system is running in one of the following modes, depending on the license type that you selected:

- ▶ Quadro vDWS
- ▶ NVIDIA vComputeServer

The system also requests the appropriate license for the current GPU from the configured license server. If unsuccessful, the system warns you that it could not obtain a license. If necessary, see [Troubleshooting](#) for guidance on troubleshooting.

Once configured in **NVIDIA X Server Settings**, licensing settings persist across reboots and need only be modified if the license server address changes, or the VM is switched to running NVIDIA vGPU.

Chapter 4.

ADVANCED TOPICS

This chapter discusses advanced topics and settings for NVIDIA vGPU software licensing.

4.1. Setting the Client Host Identifier for License Checkouts

If a VM is shut down abruptly and then restarted after a change to its network configuration, two licenses might be checked out for the same VM. You can avoid this situation by setting the client host identifier for license checkouts.

By default, the license server selects the first valid MAC address it finds to identify the VM for license checkouts. If a VM that has a license checked out is shut down abruptly, it might not release the license back to the license server. If the network configuration of the VM is changed after the shutdown and the VM is restarted, the license server might use a different MAC address to identify the VM, which will cause two licenses to be checked out for the same VM. Setting the client host identifier avoids this situation because the license server always uses identifier that you set to identify the VM.

1. Set `LicenseHostId` to the MAC address of a network interface controller (NIC) on the VM.

- ▶ On Windows, add the following Windows registry key setting:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\NVIDIA Corporation\Global\GridLicensing]"LicenseHostId"=vm-nic-mac-address
```

- ▶ On Linux, add the following entry to `/etc/nvidia/gridd.conf`:

```
LicenseHostId==vm-nic-mac-address
```

vm-nic-mac-address

The MAC address of the network interface controller (NIC) on the VM that the license server will use to identify the VM for license checkouts. Specify **only** the numbers and letters in the address **without** any spaces or punctuation, for example, **00005E0053FF**.

2. Apply the changed setting to the VM.

- ▶ If the vGPU or physical GPU assigned to the VM has already been licensed, restart the VM.
- ▶ Otherwise, license the vGPU or physical GPU assigned to the VM as explained in the following topics:
 - ▶ [Licensing an NVIDIA vGPU](#)
 - ▶ [Licensing a Physical GPU](#)

4.2. Licenses Obtained After Boot

Under normal operation, an NVIDIA vGPU software license is obtained by a platform during boot, before user login and launch of applications. If a license is not available, the system will periodically retry its license request to the license server. During this time, NVIDIA vGPU runs at reduced capability as described in [Licensing an NVIDIA vGPU](#). Similarly, Quadro vDWS features described in [Licensing a Physical GPU](#) are not available.

On Windows, unavailability of a license is indicated by a pop-up window or by the license status on the **Licensing** tab of the **NVIDIA Control Panel**. On Linux, unavailability of a license is indicated by log messages.

When a license is obtained, the licensed features are dynamically enabled and become available for immediate use. However, any application software launched before the license became available may need to be restarted in order to recognize and utilize the licensed features.

4.3. Operating with Intermittent Connectivity to the License Server

To check out a license, Quadro vDWS, GRID Virtual PC, and GRID Virtual Applications clients require connectivity to a license server when booting.

After they have booted, clients may operate without connectivity to the license server for a period of up to 1 day, after which time the client will be warned of license expiration.

4.4. Applying Windows License Settings Through the Windows Registry

NVIDIA vGPU software licensing settings can be controlled via the Windows Registry, removing the need for manual interaction with **NVIDIA Control Panel**. Settings are stored in this registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\NVIDIA Corporation\Global\GridLicensing
```

Registry values are summarized in [Table 3](#).

Table 3 Licensing Registry Settings

Name	Type	Description
ServerAddress	String (REG_SZ)	License server address
BackupServerAddress	String (REG_SZ)	Secondary license server address
ServerPort	String (REG_SZ)	License server port number (Optional: defaults to 7070)
BackupServerPort	String (REG_SZ)	Secondary license server port number (Optional: defaults to 7070)
FeatureType	DWord (REG_DWORD)	The feature type of a GPU in pass-through mode or a bare-metal deployment: 0: GRID Virtual Apps 2: Quadro vDWS All other values reserved. Do not change the value of this registry key in a VM configured with NVIDIA vGPU.
NvCplDisableManageLicensePage	DWord (REG_DWORD)	If this registry key is absent, licensing controls are shown in NVIDIA Control Panel for GPUs supporting licensing. When present: 0: Show licensing controls in NVIDIA Control Panel 1: Hide licensing controls in NVIDIA Control Panel All other values reserved.
LicenseInterval	DWord (REG_DWORD)	An integer in the range 10-10080 that specifies the period of time in minutes for which a license can be borrowed after it is checked out After this period has elapsed, the client must obtain a new license from the server. The default is 1440 minutes, which corresponds to a period of 1 day.
LingerInterval	DWord (REG_DWORD)	An integer in the range 0-10080 that specifies the period of time in minutes for which a license will remain checked out against a client VM after the VM is shut down. During this period, a VM that is rebooting can reclaim the

Name	Type	Description
		<p>same license after the reboot is complete.</p> <p>After this period has elapsed, the license is freed and available for use by other clients.</p> <p>The default is 0 minutes, which instantly frees licenses from a VM that is shut down cleanly.</p> <p>If you oversubscribe your licenses and want to protect against accidental loss of a license during a reboot, increase the <code>LingerInterval</code> to a value that reflects the amount of time required for a reboot.</p>
DisableExpirationPopups	DWord (REG_DWORD)	<p>Disables or enables pop-up notifications for NVIDIA vGPU software license expiration.</p> <p>0: Enable license expiration pop-up notifications</p> <p>1: Disable license expiration pop-up notifications</p> <p>(Optional: defaults to 0)</p>
DisableSpecificPopups	DWord (REG_DWORD)	<p>If this registry key is absent, pop-up notifications are enabled for changes in NVIDIA vGPU software license state to the following states:</p> <ul style="list-style-type: none"> ▶ Unlicensed ▶ Requesting license ▶ License acquired ▶ License request failed ▶ License renewal failed ▶ License expired (if <code>DisableExpirationPopups</code> is 0) <p>When present:</p> <p>0: Disable none of the pop-up notifications for license state changes</p> <p>1: Disable pop-up notifications for changes to the following license states:</p> <ul style="list-style-type: none"> ▶ Requesting license ▶ License acquired ▶ License renewal failed <p>(Optional: defaults to 0)</p>

Name	Type	Description
EnableLogging	DWord (REG_DWORD)	If this registry key is absent, significant licensing events are logged. When present: 0: Disable logging of significant licensing events 1: Enable logging of significant licensing events (Optional: defaults to 1)
LicenseHostId	String (REG_SZ)	If this registry key is absent, the license server selects the first valid MAC address it finds to identify the VM for license checkouts. When present, this registry key must be the MAC address of the network interface controller (NIC) on the VM that the license server will use to identify the VM for license checkouts. Specify only the numbers and letters in the address without any spaces or punctuation, for example, 00005E0053FF. Setting this registry prevents two license checkouts by a single VM if the VM is shut down abruptly and then restarted after a change to its network configuration.

Figure 5 shows an example of configuring virtual GPU licensing settings in the registry. Note it is sufficient to simply set the license server address in `ServerAddress`. Leave `FeatureType` at its preset value.

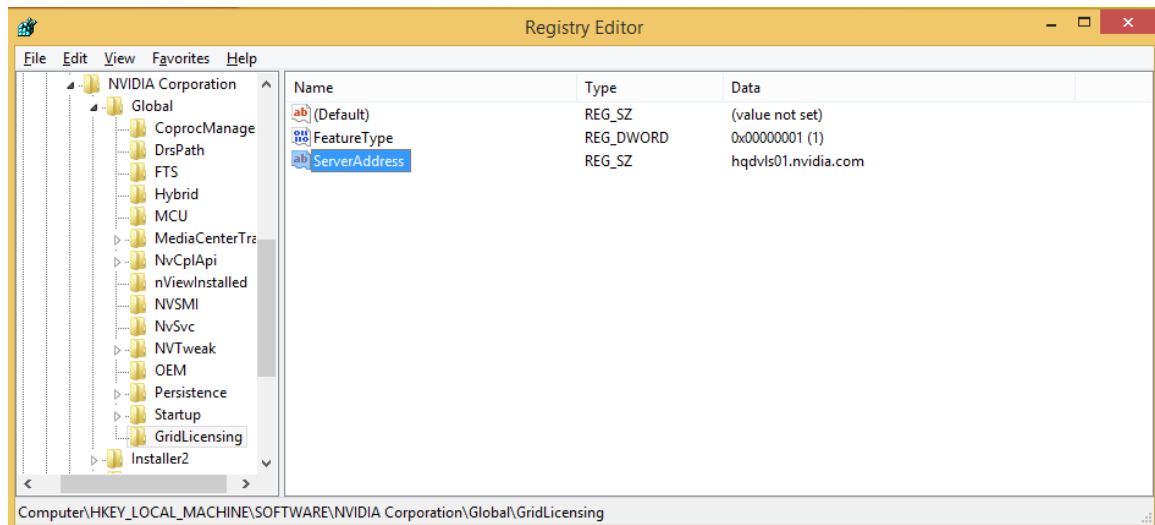


Figure 5 Configuring vGPU Licensing Through Windows Registry Settings

Chapter 5.

TROUBLESHOOTING

This chapter describes basic troubleshooting steps.

5.1. Known Issues

Before troubleshooting or filing a bug report, review the release notes that accompany each driver release, for information about known issues with the current release, and potential workarounds.

5.2. Licensing Event Logs

Significant licensing events, for example, acquisition of a license, return of a license, expiration of a license, or failure to acquire a license, are logged to help troubleshoot issues. Error events are logged with an error code to help you to determine the underlying cause of the error. License acquisition events are logged with the name and version of the licensed product, for example, `GRID-Virtual-WS, 2.0`.

Windows Licensing Event Logs

On Windows, licensing events are logged in the plain-text file `%SystemDrive%\Users\Public\Documents\NvidiaLogging\Log.NVDisplay.Container.exe.log`.

The log file is rotated when its size reaches 16 MB. A new log file is created and the old log file is renamed to `Log.NVDisplay.Container.exe.log1`. Each time the log file is rotated, the number in the file name of each existing old log file is increased by 1. The oldest log file is deleted when the number of log files exceeds 16.

Logging is enabled by default. To disable logging, create the Windows Registry key `EnableLogging` (DWORD) with the value zero at the path `HKLM\SOFTWARE\NVIDIA Corporation\Global\GridLicensing`.

Linux Licensing Event Logs

On Linux, logging is enabled by default. The licensing events are logged in the standard activity log `/var/log/messages`. You can use the standard features of the operating system to rotate and purge these logs.

Log Messages for Acquisition and Return of a License and Transient Network Failure

This example shows log messages for the following sequence of events:

1. Acquisition of a license when a VM is started
2. Return of a license when the VM is shut down
3. Failure to communicate with the license server when the VM is started when the license server is unreachable
4. Acquisition of a license after communications with the license server are restored

```
Mon Oct 22 20:53:29 2018:<1>:License acquired successfully. (Info: http://
gridlicense1.example.com:7070/request; GRID-Virtual-WS,2.0)
Mon Oct 22 20:54:22 2018:<1>:End Logging
Mon Oct 22 20:54:24 2018:<0>:License returned successfully
Mon Oct 22 20:54:24 2018:<0>:End Logging
Mon Oct 22 20:55:36 2018:<1>:Failed server communication. Please ensure your
system has network connectivity to the license server
Mon Oct 22 20:55:49 2018:<1>:Failed server communication. Please ensure your
system has network connectivity to the license server
Mon Oct 22 20:56:01 2018:<1>:Failed server communication. Please ensure your
system has network connectivity to the license server
Mon Oct 22 20:56:19 2018:<1>:Failed server communication. Please ensure your
system has network connectivity to the license server
Mon Oct 22 20:56:31 2018:<1>:Failed server communication. Please ensure your
system has network connectivity to the license server
Mon Oct 22 20:56:44 2018:<1>:Failed server communication. Please ensure your
system has network connectivity to the license server
Mon Oct 22 20:56:47 2018:<1>:License acquired successfully. (Info: http://
gridlicense1.example.com:7070/request; GRID-Virtual-WS,2.0)
Mon Oct 22 21:36:11 2018:<1>:End Logging
Mon Oct 22 21:36:12 2018:<0>:License returned successfully
Mon Oct 22 21:36:12 2018:<0>:End Logging
```

Log Messages for Failure to Acquire a License

This example shows log messages for the failure to acquire a license because no licenses of the requested type are available.

```
Mon Oct 22 22:26:36 2018:<1>:Error : Server URL : http://
gridlicense1.example.com:7070/request -
[1,7E2,2,0[70000030,0,702C3]]
Insufficient count for the requested feature.
Mon Oct 22 22:28:15 2018:<1>:Error : Server URL : http://
gridlicense1.example.com:7070/request -
[1,7E2,2,0[70000030,0,702C3]]
Insufficient count for the requested feature.
Mon Oct 22 22:29:55 2018:<1>:Error : Server URL : http://
gridlicense1.example.com:7070/request -
[1,7E2,2,0[70000030,0,702C3]]
Insufficient count for the requested feature.
Mon Oct 22 22:31:34 2018:<1>:Error : Server URL : http://
gridlicense1.example.com:7070/request -
[1,7E2,2,0[70000030,0,702C3]]
```

Insufficient count for the requested feature.

5.3. Troubleshooting Steps

If an NVIDIA vGPU software system fails to obtain a license, troubleshoot the failure as follows:

- ▶ Examine the log of significant licensing events for information about the cause of the failure. For details, see [Licensing Event Logs](#).
- ▶ Check that the license server address and port number are correctly configured.
- ▶ Run a network ping test from the NVIDIA vGPU software system to the license server address to verify that the system has network connectivity to the license server.
- ▶ Verify that the date and time are configured correctly on the NVIDIA vGPU software system. If the time is set inaccurately or is adjusted backwards by a large amount, the system may fail to obtain a license.
- ▶ Verify that the license server in use has available licenses of the type required by the NVIDIA vGPU software feature the NVIDIA vGPU software system is configured to use.

Appendix A.

LICENSING NVIDIA vGPU SOFTWARE ON LINUX BY USING A CONFIGURATION FILE

By default, the **Manage License** option is not available in **NVIDIA X Server Settings** and you must enable this option if you want to manage licensing by using **NVIDIA X Server Settings**. However, this option must be disabled with Red Hat Enterprise Linux 6.8 and 6.9 or CentOS 6.8 and 6.9.

If you do not want to or cannot enable the **Manage License** option, you can license NVIDIA vGPU software on Linux by using the configuration file `/etc/nvidia/gridd.conf` instead of using **NVIDIA X Server Settings**.

A.1. Licensing an NVIDIA vGPU on Linux by Using a Configuration File

1. As root, open the file `/etc/nvidia/gridd.conf` in a plain-text editor, such as `vi`.

```
[nvidia@localhost ~]$ sudo vi /etc/nvidia/gridd.conf
```



You can create the `/etc/nvidia/gridd.conf` file by copying the supplied template file `/etc/nvidia/gridd.conf.template`.

2. Set `ServerAddress` to the address of your primary NVIDIA vGPU software License Server.

The address can be a fully-qualified domain name such as `gridlicense1.example.com`, or an IP address such as `10.31.20.45`.

If you have only one license server configured, set `ServerAddress` to the address of the server.

3. Optional: Set `ServerPort` to the port number of your primary NVIDIA vGPU software License Server.

If you do not set `ServerPort`, it will default to `7070`, which is the default port number that is used by the NVIDIA vGPU software License Server.

4. Set `BackupServerAddress` to the address of your secondary NVIDIA vGPU software License Server.

The address can be a fully-qualified domain name such as `gridlicense2.example.com`, or an IP address such as `10.31.20.46`.

If you have only one license server configured, do not set `BackupServerAddress`.

5. Optional: Set `BackupServerPort` to the port number of your secondary NVIDIA vGPU software License Server.

If you do not set `BackupServerPort`, it will default to `7070`, which is the default port number that is used by the NVIDIA vGPU software License Server.

If you have only one license server configured, do not set `BackupServerPort`.

6. Set `FeatureType` to `1`, to license vGPU.



You do not need to specify the type of the license. NVIDIA vGPU software automatically selects the correct type of license based on the vGPU type.

7. Save your changes to the `/etc/nvidia/gridd.conf` file.
8. Restart the `nvidia-gridd` service:

```
[nvidia@localhost ~]$ sudo service nvidia-gridd restart
```

The `nvidia-gridd` service requests the appropriate license for the current vGPU from the configured license server.

9. Optional: Confirm that the service has obtained a license by examining the log messages written to `/var/log/messages`:

```
[nvidia@localhost ~]$ sudo grep gridd /var/log/messages
```

```
...
Aug 5 15:40:06 localhost nvidia-gridd: Started (4293)
Aug 5 15:40:24 localhost nvidia-gridd: License acquired successfully.
```

The vGPU within the VM should now exhibit full frame rate, resolution, and display output capabilities.

`gridd.conf` File for NVIDIA vGPU

The following example shows a `gridd.conf` file for NVIDIA vGPU in which options are set as follows:

- ▶ `ServerAddress` is set to `gridlicense1.example.com`.
- ▶ `BackupServerAddress` is set to `gridlicense2.example.com`.
- ▶ `ServerPort` and `BackupServerPort` are both set to `7070`.
- ▶ `FeatureType` is set to `1`.

```
# /etc/nvidia/gridd.conf.template - Configuration file for NVIDIA Grid Daemon

# This is a template for the configuration file for NVIDIA Grid Daemon.
# For details on the file format, please refer to the nvidia-gridd(1)
# man page.

# Description: Set License Server Address
# Data type: string
# Format: "<address>"
ServerAddress= gridlicense1.example.com
```

```

# Description: Set License Server port number
# Data type: integer
# Format: <port>, default is 7070
ServerPort=7070

# Description: Set Backup License Server Address
# Data type: string
# Format: "<address>"
BackupServerAddress= gridlicense2.example.com

# Description: Set Backup License Server port number
# Data type: integer
# Format: <port>, default is 7070
BackupServerPort=7070

# Description: Set Feature to be enabled
# Data type: integer
# Possible values:
#   0 => for unlicensed state
#   1 => for GRID vGPU
#   2 => for Quadro Virtual Datacenter Workstation
#   4 => for NVIDIA vComputeServer
FeatureType=1

# Description: Parameter to enable or disable Grid Licensing tab in nvidia-
settings
# Data type: boolean
# Possible values: TRUE or FALSE, default is FALSE
#EnableUI=TRUE

# Description: Set license borrow period in minutes
# Data type: integer
# Possible values: 10 to 10080 mins(7 days), default is 1440 mins(1 day)
#LicenseInterval=1440

# Description: Set license linger period in minutes
# Data type: integer
# Possible values: 0 to 10080 mins(7 days), default is 0 mins
#LingerInterval=10

```

If the system fails to obtain a license, see [Troubleshooting](#) for guidance on troubleshooting.

A.2. Licensing a Physical GPU on Linux by Using a Configuration File

Licensing of a physical GPU in pass-through mode or a bare-metal deployment for Quadro vDWS or vComputeServer is enforced through software. To enable you to apply the correct type of license, options are provided for specifying these types of licenses.



Licensing of a physical GPU in pass-through mode or a bare-metal deployment for GRID Virtual Apps is enforced through the EULA only. Therefore, no option is provided to specify the GRID Virtual Apps license type.

1. As root, open the file `/etc/nvidia/gridd.conf` in a plain-text editor, such as `vi`.

```
[nvidia@localhost ~]$ sudo vi /etc/nvidia/gridd.conf
```



You can create the `/etc/nvidia/gridd.conf` file by copying the supplied template file `/etc/nvidia/gridd.conf.template`.

- Set `ServerAddress` to the address of your primary NVIDIA vGPU software License Server.
The address can be a fully-qualified domain name such as `gridlicense1.example.com`, or an IP address such as `10.31.20.45`.
If you have only one license server configured, set `ServerAddress` to the address of the server.
- Optional: Set `ServerPort` to the port number of your primary NVIDIA vGPU software License Server.
If you do not set `ServerPort`, it will default to `7070`, which is the default port number that is used by the NVIDIA vGPU software License Server.
- Set `BackupServerAddress` to the address of your secondary NVIDIA vGPU software License Server.
The address can be a fully-qualified domain name such as `gridlicense2.example.com`, or an IP address such as `10.31.20.46`.
If you have only one license server configured, do not set `BackupServerAddress`.
- Optional: Set `BackupServerPort` to the port number of your secondary NVIDIA vGPU software License Server.
If you do not set `BackupServerPort`, it will default to `7070`, which is the default port number that is used by the NVIDIA vGPU software License Server.
If you have only one license server configured, do not set `BackupServerPort`.
- Set `FeatureType` to the integer that specifies the type of license that you want to apply to the GPU.

License Type	FeatureType
Quadro vDWS	2
vComputeServer	4

- Save your changes to the `/etc/nvidia/gridd.conf` file.
- Restart the `nvidia-gridd` service:

```
[nvidia@localhost ~]$ sudo service nvidia-gridd restart
```

The service should automatically obtain a license.

- Optional: Confirm that the service has obtained a license by examining the log messages written to `/var/log/messages`:

```
[nvidia@localhost ~]$ sudo grep gridd /var/log/messages
```

```
...
Aug 5 17:12:15 localhost nvidia-gridd: Started (10430)
Aug 5 17:12:34 localhost nvidia-gridd: License acquired successfully.
```


Once configured in `gridd.conf`, licensing settings persist across reboots and need only be modified if the license server address changes, or the VM is switched to running NVIDIA vGPU.

`gridd.conf` File for Quadro vDWS

The following example shows a `gridd.conf` file for Quadro vDWS in which options are set as follows:

- ▶ `ServerAddress` is set to `gridlicense1.example.com`.
- ▶ `BackupServerAddress` is set to `gridlicense2.example.com`.
- ▶ `ServerPort` and `BackupServerPort` are both set to 7070.
- ▶ `FeatureType` is set to 2.

```
# /etc/nvidia/gridd.conf.template - Configuration file for NVIDIA Grid Daemon

# This is a template for the configuration file for NVIDIA Grid Daemon.
# For details on the file format, please refer to the nvidia-gridd(1)
# man page.

# Description: Set License Server Address
# Data type: string
# Format: "<address>"
ServerAddress= gridlicense1.example.com

# Description: Set License Server port number
# Data type: integer
# Format: <port>, default is 7070
ServerPort=7070

# Description: Set Backup License Server Address
# Data type: string
# Format: "<address>"
BackupServerAddress= gridlicense2.example.com

# Description: Set Backup License Server port number
# Data type: integer
# Format: <port>, default is 7070
BackupServerPort=7070

# Description: Set Feature to be enabled
# Data type: integer
# Possible values:
#   0 => for unlicensed state
#   1 => for GRID vGPU
#   2 => for Quadro Virtual Datacenter Workstation
#   4 => for NVIDIA vComputeServer
FeatureType=2

# Description: Parameter to enable or disable Grid Licensing tab in nvidia-
settings
# Data type: boolean
# Possible values: TRUE or FALSE, default is FALSE
#EnableUI=TRUE

# Description: Set license borrow period in minutes
# Data type: integer
# Possible values: 10 to 10080 mins(7 days), default is 1440 mins(1 day)
#LicenseInterval=1440

# Description: Set license linger period in minutes
# Data type: integer
```

```
# Possible values: 0 to 10080 mins(7 days), default is 0 mins
#LingerInterval=10
```

gridd.conf File for vComputeServer

The following example shows a `gridd.conf` file for `vComputeServer` in which options are set as follows:

- ▶ `ServerAddress` is set to `gridlicense1.example.com`.
- ▶ `BackupServerAddress` is set to `gridlicense2.example.com`.
- ▶ `ServerPort` and `BackupServerPort` are both set to `7070`.
- ▶ `FeatureType` is set to `4`.

```
# /etc/nvidia/gridd.conf.template - Configuration file for NVIDIA Grid Daemon

# This is a template for the configuration file for NVIDIA Grid Daemon.
# For details on the file format, please refer to the nvidia-gridd(1)
# man page.

# Description: Set License Server Address
# Data type: string
# Format: "<address>"
ServerAddress= gridlicense1.example.com

# Description: Set License Server port number
# Data type: integer
# Format: <port>, default is 7070
ServerPort=7070

# Description: Set Backup License Server Address
# Data type: string
# Format: "<address>"
BackupServerAddress= gridlicense2.example.com

# Description: Set Backup License Server port number
# Data type: integer
# Format: <port>, default is 7070
BackupServerPort=7070

# Description: Set Feature to be enabled
# Data type: integer
# Possible values:
#   0 => for unlicensed state
#   1 => for GRID vGPU
#   2 => for Quadro Virtual Datacenter Workstation
#   4 => for NVIDIA vComputeServer
FeatureType=4

# Description: Parameter to enable or disable Grid Licensing tab in nvidia-
settings
# Data type: boolean
# Possible values: TRUE or FALSE, default is FALSE
EnableUI=TRUE

# Description: Set license borrow period in minutes
# Data type: integer
# Possible values: 10 to 10080 mins(7 days), default is 1440 mins(1 day)
LicenseInterval=1440

# Description: Set license linger period in minutes
# Data type: integer
# Possible values: 0 to 10080 mins(7 days), default is 0 mins
LingerInterval=10
```

Notice

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.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

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, vGPU, Pascal, 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

© 2013-2020 NVIDIA Corporation. All rights reserved.