




GRID LICENSING

DU-07757-001 | August 2016

User Guide



DOCUMENT CHANGE HISTORY

DU-07757-001

Version	Date	Authors	Description of Change
1.0	9/1/2015	AC	Release for GRID 2.0
1.1	8/12/2016	PD	Specify server host and port separately on Linux Specify the period of time for which a license can be borrowed

Chapter 1. Introduction	1
1.1 How licensing works.....	1
1.2 License editions	2
Chapter 2. GRID Virtual GPU	3
2.1 vGPU License requirements	3
2.2 Licensing on Windows	4
2.3 Licensing on Linux	7
Chapter 3. GRID Virtual Workstation	9
3.1 GRID Virtual Workstation features	9
3.2 Licensing on Windows	9
3.3 Disabling GRID Virtual Workstation	12
3.4 Licensing on Linux	14
Chapter 4. Advanced topics	16
4.1 Licenses obtained after boot.....	16
4.2 Operating with intermittent connectivity to the license server	16
4.3 Applying Windows license settings via registry	17
Chapter 5. Troubleshooting	19
5.1 Known issues	19
5.2 Troubleshooting steps	19

LIST OF FIGURES

Figure 1 GRID licensing architecture	1
Figure 2 GRID license editions	2
Figure 3 Managing vGPU licensing in NVIDIA Control Panel	5
Figure 4 Successful acquire of vGPU license	6
Figure 5 Sample gridd.conf for GRID vGPU	7
Figure 6 Managing Virtual Workstation Licensing in NVIDIA Control Panel	10
Figure 7 Applying GRID Virtual Workstation license.....	11
Figure 8 Success acquire of Virtual Workstation license	12
Figure 9 Disabling GRID Virtual Workstation	13
Figure 10 Sample gridd.conf for GRID Virtual Workstation.....	14
Figure 11 Configuring vGPU licensing via registry settings	18

LIST OF TABLES

Table 1 Virtual GPUs licensed on Tesla M6, M60	4
Table 2 Licensing registry settings.....	18

Chapter 1. INTRODUCTION

Certain NVIDIA GRID™ products, such as GRID vGPU™ and GRID Virtual Workstation, are available as licensed features on NVIDIA Tesla™ GPUs. This guide describes the licensed features and how to enable and use them on supported hardware.

1.1 HOW LICENSING WORKS

Figure 1 provides an overview of GRID licensing:

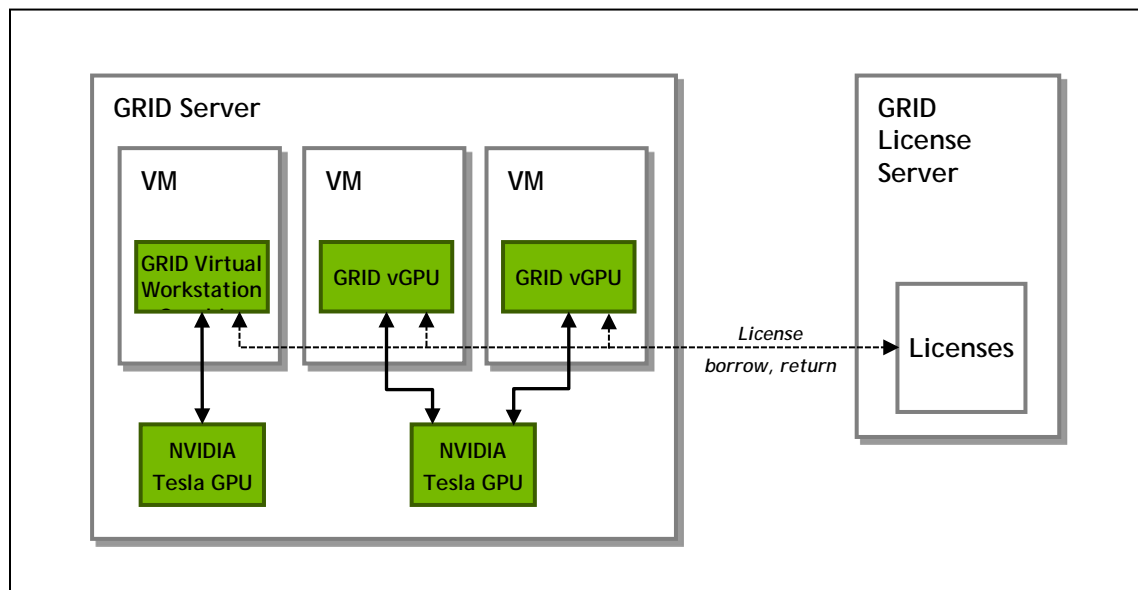


Figure 1 GRID licensing architecture

When enabled on Tesla GPUs, licensed features such as Virtual GPU are activated by obtaining a license over the network from an NVIDIA GRID License Server. The license is “checked out” or “borrowed” at the time the Virtual Machine (VM) is booted, and returned when the VM is shut down.



Note: Information on setting up and running the GRID License Server is provided in the NVIDIA GRID License Server User Guide.

1.2 LICENSE EDITIONS

GRID licenses come in three editions that enable different classes of GRID features. The GRID software automatically selects the right license edition based on the features being used:

GRID License Edition	GRID Features
GRID Virtual PC	<ul style="list-style-type: none"> Virtual GPUs for business desktop computing
GRID Virtual Workstation	<ul style="list-style-type: none"> Virtual GPUs for midrange workstation computing
GRID Virtual Workstation - Extended	<ul style="list-style-type: none"> Virtual GPUs for high-end workstation computing Workstation graphics on GPU passthrough

Figure 2 GRID license editions

The remainder of this guide is organized as follows:

- ▶ Chapter 2 describes licensing of GRID Virtual GPU.
- ▶ Chapter 3 describes licensing of GRID Virtual Workstation features with GPU passthrough.
- ▶ Chapter 4 discusses advanced licensing settings.
- ▶ Chapter 5 provides guidance on troubleshooting.

Chapter 2. GRID VIRTUAL GPU

This chapter describes licensing of NVIDIA GRID vGPU.

2.1 VGPU LICENSE REQUIREMENTS

NVIDIA GRID Virtual GPU (vGPU) is offered as a licensable feature on Tesla M6 and M60 GPUs. When a vGPU is booted on these GPUs, a license must be obtained by the Virtual Machine (VM) in order to enable the full features of the vGPU. The VM retains the license until it is shut down; it then releases the license back to the license server.



Note: vGPUs that require licensing run at a reduced capability until a license is acquired. Screen resolution is limited to no higher than 1280x1024, frame rate is capped at 3 frames per second, and GPU resource allocations are limited, which will prevent some applications from running correctly. These restrictions are removed once a license is acquired.

NVIDIA GRID K1 and K2 GPUs do not require a license to run vGPU.

Table 1 lists the vGPU types that are supported on Tesla M6 / M60, and the license edition that each vGPU type requires.

Card	GRID Virtual GPU	Frame Buffer (Mbytes)	Virtual Display Heads	Max Resolution per Display Head	Minimum GRID License Edition Required
Tesla M60	M60-8Q	8192	4	3840x2160	GRID Virtual Workstation - Extended
	M60-4Q	4096			
	M60-2Q	2048	4	2560x1600	GRID Virtual Workstation
	M60-1Q	1024	2	2560x1600	
	M60-0Q	512			
	M60-2B	2048	2	2560x1600	GRID Virtual PC
	M60-1B	1024			
M60-0B	512				
Tesla M6	M6-8Q	8192	4	3840x2160	GRID Virtual Workstation - Extended
	M6-4Q	4096	4		
	M6-2Q	2048	4	2560x1600	GRID Virtual Workstation
	M6-1Q	1024	2	2560x1600	
	M6-0Q	512			
	M6-2B	2048	2	2560x1600	GRID Virtual PC
	M6-1B	1024			
M6-0B	512				

Table 1 Virtual GPUs licensed on Tesla M6, M60

The higher-end GRID license editions are inclusive of lower editions: for example virtual GPUs that require a GRID Virtual PC license are also usable with a GRID Virtual Workstation or GRID Virtual Workstation Extended license.



Note: For more information on the virtual GPU types listed in Table 1, consult the NVIDIA GRID vGPU User Guide.

2.2 LICENSING ON WINDOWS

To license vGPU, open NVIDIA Control Panel by right-clicking on the Windows desktop and selecting **NVIDIA Control Panel** from the menu, or by opening Windows Control Panel and double-clicking the NVIDIA Control Panel icon.

In NVIDIA Control Panel, select **Manage License** task in the **Licensing** section of the navigation pane, as shown in Figure 3.



Note: If the Licensing section and Manage License task are not displayed in NVIDIA Control Panel, the vGPU you are running on does not require a license, or the system has been configured to hide licensing controls in NVIDIA Control Panel (see section [Error! Reference source not found.](#) on registry settings).

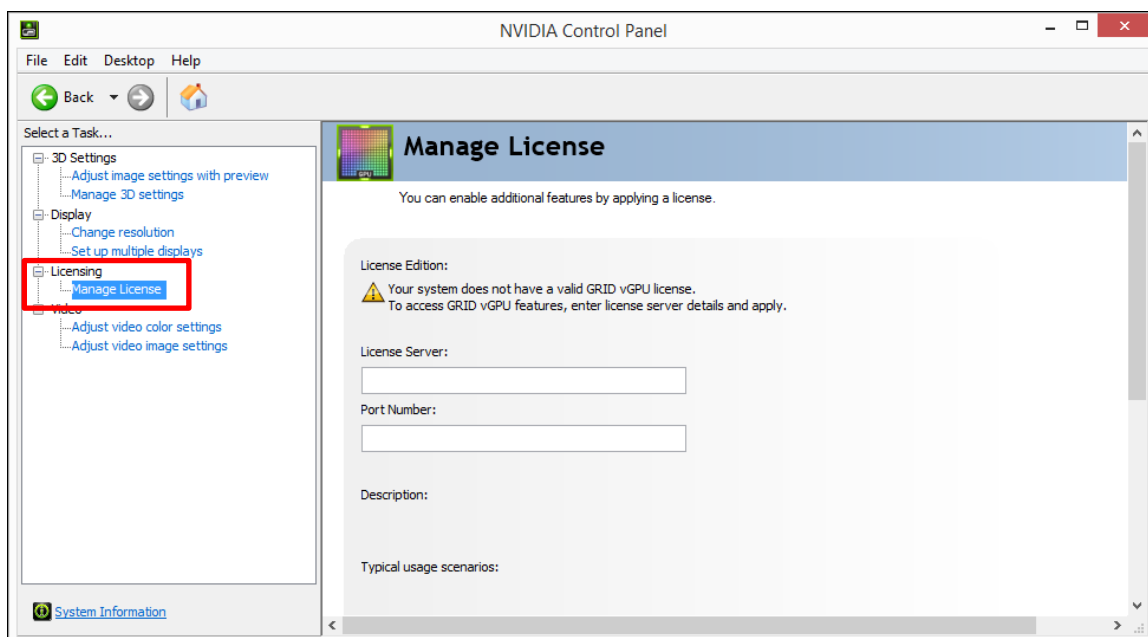


Figure 3 Managing vGPU licensing in NVIDIA Control Panel

The **Manage License** task pane shows that GRID vGPU is currently unlicensed. Enter the address of your local GRID License Server in the **License Server** field. The address can be a fully-qualified domain name such as `gridlicense.example.com`, or an IP address such as `10.31.20.45`.

The **Port Number** field can be left unset and will default to `7070`, which is the default port number used by NVIDIA GRID License Server.

Select **Apply** to assign the settings. The system will request the appropriate license for the current vGPU from the configured license server and, if successful, vGPU's full capabilities are enabled (see Figure 4). If the system fails to obtain a license, refer to 4.3 for guidance on troubleshooting.

Once configured in NVIDIA Control Panel, licensing settings persist across reboots.

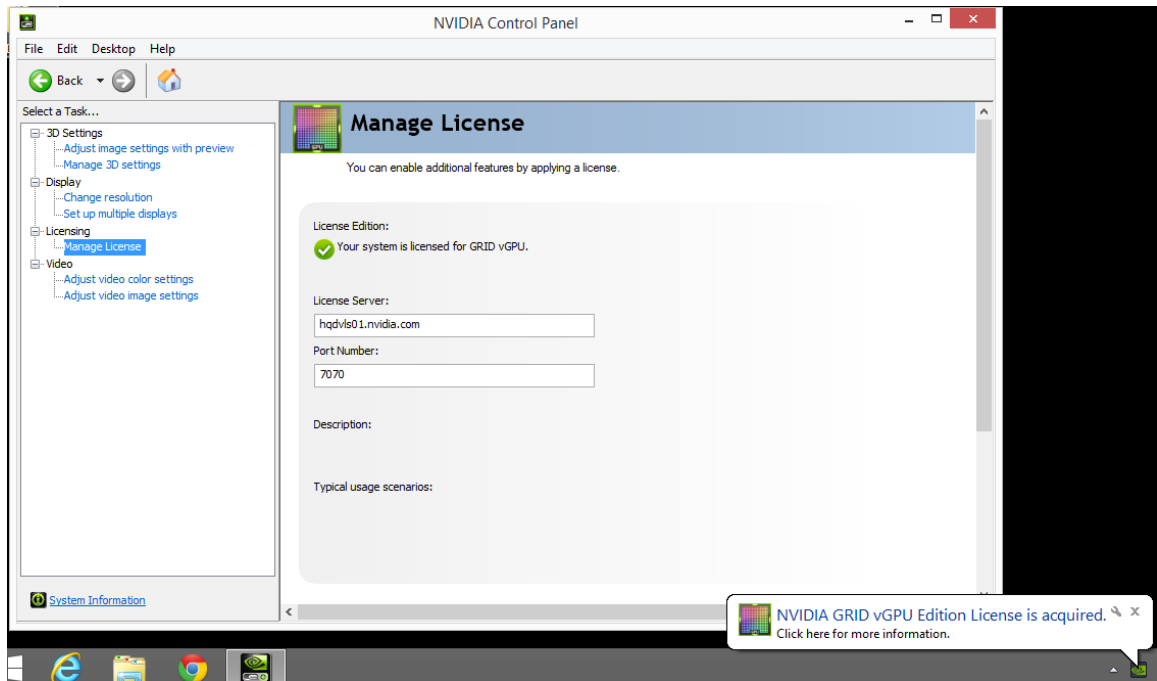


Figure 4 Successful acquire of vGPU license

2.3 LICENSING ON LINUX

To license GRID vGPU, edit `/etc/nvidia/gridd.conf`:

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

Set `ServerAddress` to the address and `ServerPort` to the port number of your local NVIDIA GRID License Server. The address can be a fully-qualified domain name such as `gridlicense.example.com`, or an IP address such as `10.31.20.45`. If you do not set `ServerPort`, it will default to `7070`, which is the default port number that is used by the NVIDIA GRID License Server.

Set `FeatureType` to `1`, to license vGPU:

```
# /etc/nvidia/gridd.conf - 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=gridlicense.example.com

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

# Description: Set Feature to be enabled
# Data type: integer
# Possible values:
#   1 => for GRID vGPU
#   2 => for GRID Virtual Workstation
FeatureType=1

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

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

Figure 5 Sample `gridd.conf` for GRID vGPU

Restart the `nvidia-gridd` service:

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

The service should automatically obtain a license. This can be confirmed with `log`

messages written to `/var/log/messages`, and the vGPU within the VM should now exhibit full framerate, resolution, and display output capabilities:

```
[nvidia@localhost ~]$ sudo grep gridd /var/log/messages
...
Aug 12 15:40:06 localhost nvidia-gridd: Started (4293)
Aug 12 15:40:24 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 GPU passthrough.

Chapter 3. GRID VIRTUAL WORKSTATION

This chapter describes how to enable the GRID Virtual Workstation feature on supported Tesla GPUs.

3.1 GRID VIRTUAL WORKSTATION FEATURES

GRID Virtual Workstation is available on Tesla GPUs running in GPU passthrough mode to Windows and Linux VMs. Virtual Workstation requires a GRID Virtual Workstation – Extended license edition, and provides these features:

- ▶ Up to four virtual display heads at 4k resolution (unlicensed Tesla GPUs support a single virtual display head with maximum resolution of 2560x1600).
- ▶ Workstation-specific graphics features and accelerations.
- ▶ Certified drivers for professional applications

3.2 LICENSING ON WINDOWS

To enable GRID Virtual Workstation, open NVIDIA Control Panel by right-clicking on the Windows desktop and selecting **NVIDIA Control Panel** from the menu, or by opening Windows Control Panel and double-clicking the NVIDIA Control Panel icon.

In NVIDIA Control Panel, select **Manage License** task in the **Licensing** section of the navigation pane, as shown in Figure 6.



Note: 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 GRID licensed features.

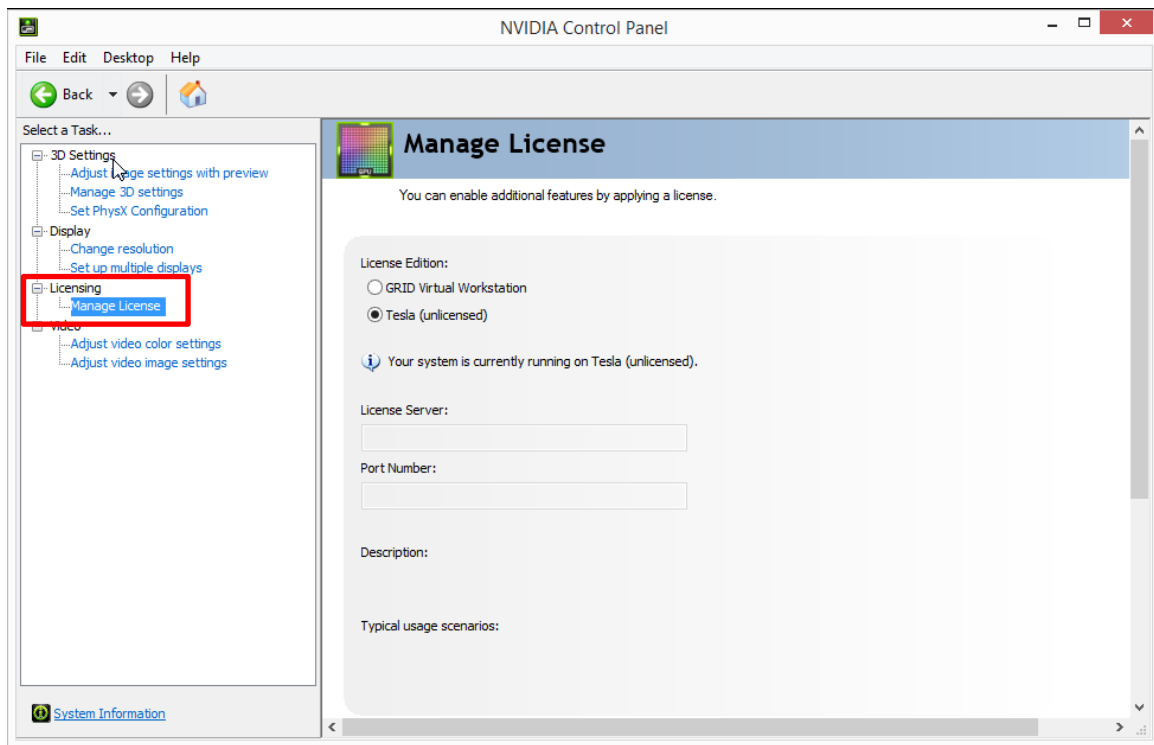


Figure 6 Managing Virtual Workstation Licensing in NVIDIA Control Panel

The **Manage License** task pane shows the current License Edition being used, and defaults to unlicensed.

Select **GRID Virtual Workstation**, and enter the address of your local GRID License Server in the **License Server** field (see Figure 7). The address can be a fully-qualified domain name such as `gridlicense.example.com`, or an IP address such as `10.31.20.45`.

The **Port Number** field can be left unset and will default to `7070`, which is the default port number used by NVIDIA GRID License Server.

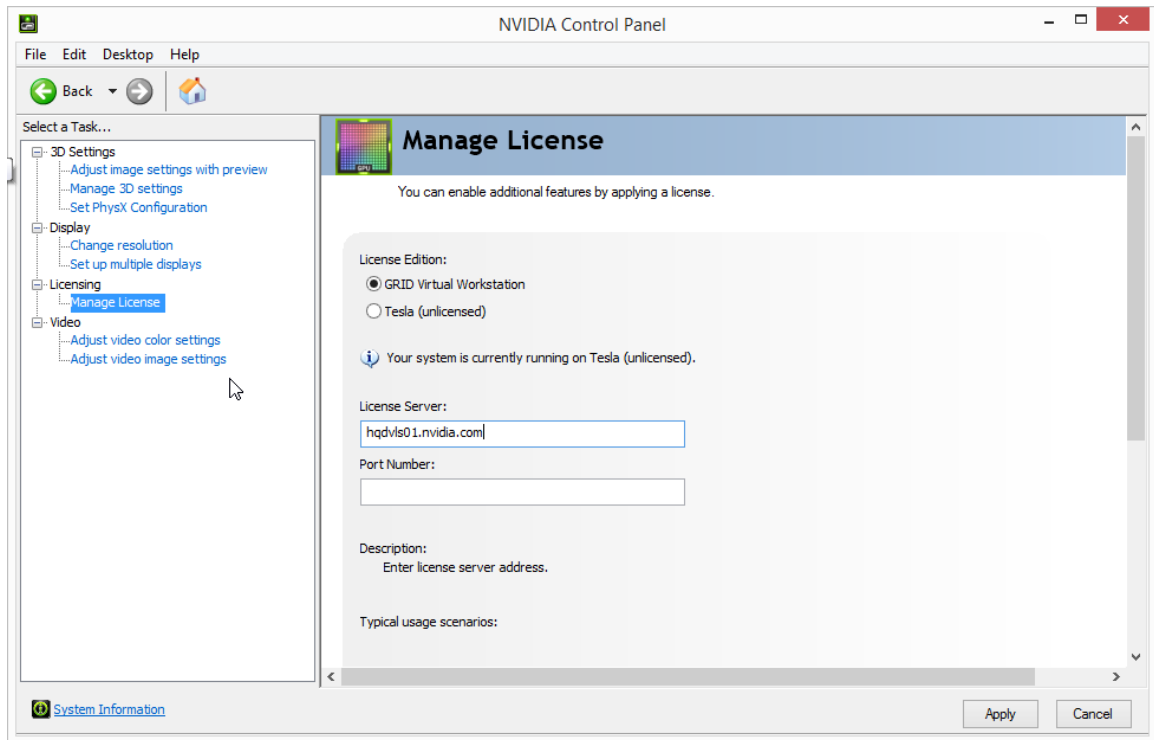


Figure 7 Applying GRID Virtual Workstation license

Select **Apply** to assign the settings. The system will request a license from the configured license server and, if successful, GRID Virtual Workstation features are enabled (see Figure 8). If the system fails to obtain a license, refer to 4.3 for guidance on troubleshooting.

Once configured in NVIDIA Control Panel, licensing settings persist across reboots.

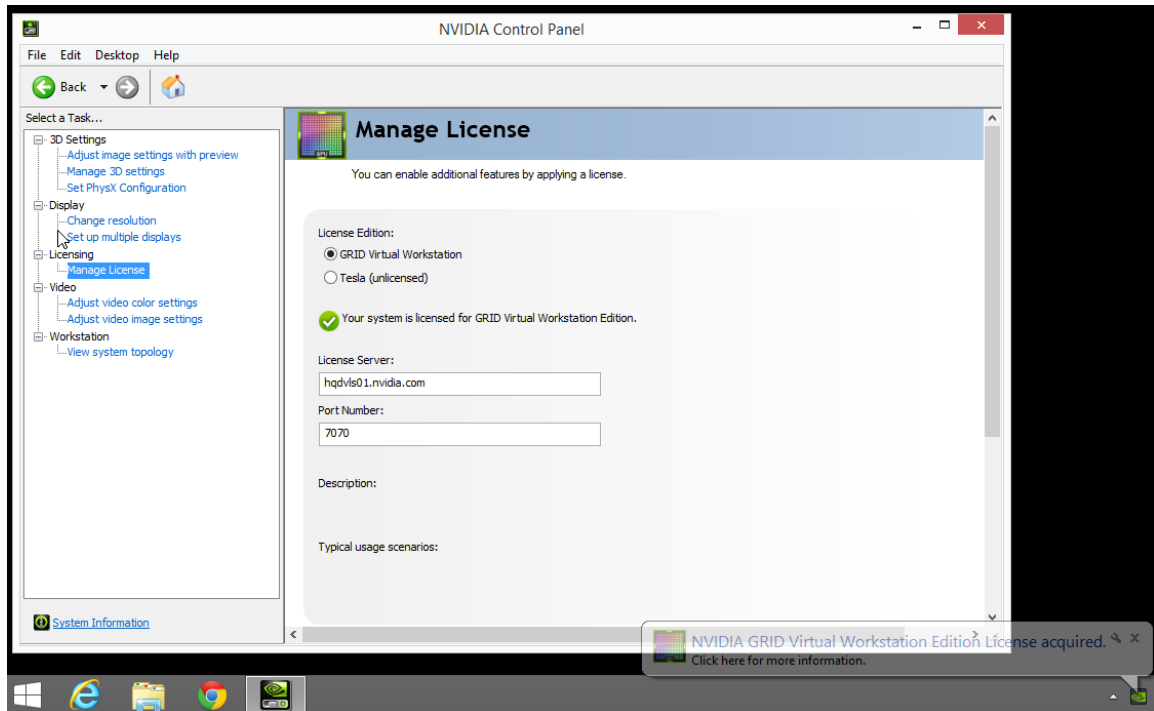


Figure 8 Success acquire of Virtual Workstation license

3.3 DISABLING GRID VIRTUAL WORKSTATION

To disable the GRID Virtual Workstation licensed feature, open NVIDIA Control Panel; in the Manage License task, select **Tesla (unlicensed)** and select **Apply** (see Figure 9). The setting does not take effect until the next time the system is shutdown or rebooted; GRID Virtual Workstation features remain available until then.




Figure 9 Disabling GRID Virtual Workstation

3.4 LICENSING ON LINUX

To license GRID Virtual Workstation, edit `/etc/nvidia/gridd.conf`:

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

Set `ServerAddress` to the address and `ServerPort` to the port number of your local NVIDIA GRID License Server. The address can be a fully-qualified domain name such as `gridlicense.example.com`, or an IP address such as `10.31.20.45`. If you do not set `ServerPort`, it will default to `7070`, which is the default port number that is used by the NVIDIA GRID License Server.

Set `FeatureType` to `2`, to license GRID Virtual Workstation:

```
# /etc/nvidia/gridd.conf - 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=gridlicense.example.com

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

# Description: Set Feature to be enabled
# Data type: integer
# Possible values:
#   1 => for GRID vGPU
#   2 => for GRID Virtual Workstation
FeatureType=2

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

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

Figure 10 Sample `gridd.conf` for GRID Virtual Workstation

Restart the `nvidia-gridd` service:

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

The service should automatically obtain a license. This can be confirmed with `log`

messages written to `/var/log/messages`, and the GPU should now exhibit Virtual Workstation display output and resolution capabilities:

```
[nvidia@localhost ~]$ sudo grep gridd /var/log/messages
...
Aug 12 15:40:06 localhost nvidia-gridd: Started (4293)
Aug 12 15:40:24 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 GRID vGPU.

Chapter 4. ADVANCED TOPICS

This chapter discusses advanced topics and settings for GRID licensing.

4.1 LICENSES OBTAINED AFTER BOOT

Under normal operation, a GRID license is obtained by a platform during boot, prior to user login and launch of applications. If a license is not available, indicated by a popup on Windows or log messages on Linux, the system will periodically retry its license request to the license server. During this time, GRID vGPU runs at reduced capability described in section 2.1; similarly, GRID Virtual Workstation features described in section 3.1 are not available.

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

4.2 OPERATING WITH INTERMITTENT CONNECTIVITY TO THE LICENSE SERVER

GRID vGPU and Virtual Workstation clients require connectivity to a license server when booting, in order to check out a license. Once booted, clients may operate without connectivity to the license server for a period of up to 7 days, after which time the client will be warned of license expiration.

4.3 APPLYING WINDOWS LICENSE SETTINGS VIA REGISTRY

GRID 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 2.

Name	Type	Description
ServerAddress	String (REG_SZ)	License server address
ServerPort	String (REG_SZ)	License server port number (Optional: defaults to 7070)
FeatureType	DWord (REG_DWORD)	0: Unlicensed (Tesla edition) 1: GRID vGPU 2: GRID Virtual Workstation All other values reserved.
NvCplDisableManageLicensePage	DWord (REG_DWORD)	If this regkey 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 10080 minutes, which corresponds to a period of 7 days.

Table 2 Licensing registry settings

Figure 11 shows an example of configuring virtual GPU licensing settings in the registry. Note it is sufficient to simply configure `FeatureType = 1` (GRID vGPU) and set the license server address in `ServerAddress`.

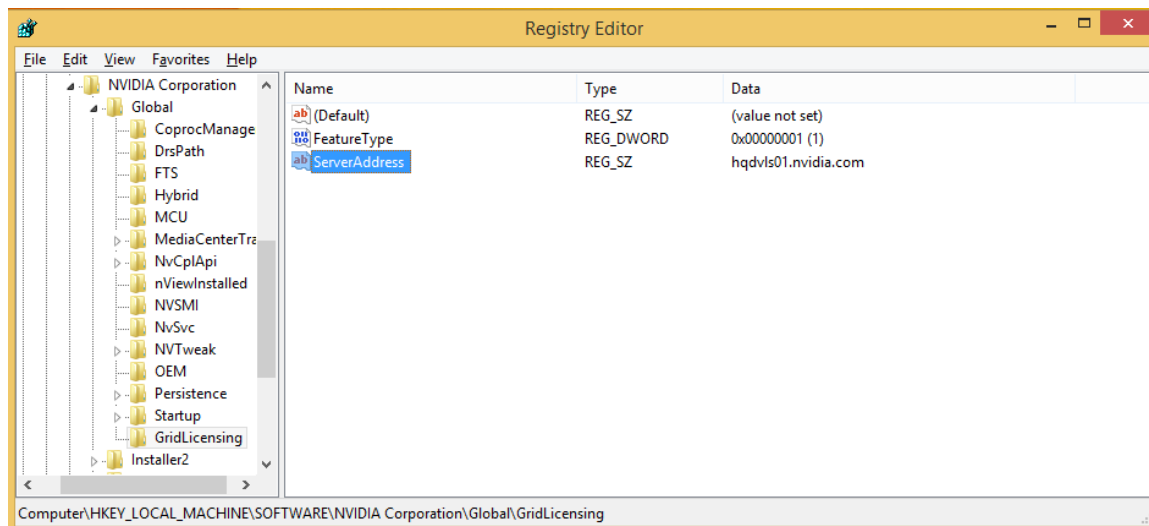


Figure 11 Configuring vGPU licensing via registry settings

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 TROUBLESHOOTING STEPS

If a GRID system fails to obtain a license, investigate the following as potential causes for the failure:

- ▶ Check that the license server address and port number are correctly configured.
- ▶ Run a network ping test from the GRID 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 GRID 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 GRID feature the GRID system is configured to use.

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 and the NVIDIA logo 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

© 2016 NVIDIA Corporation. All rights reserved.