




GRID LICENSING

DU-07757-001 | April 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
2.0	4/4/2016	PD	Release for GRID 3.0

Chapter 1. Introduction	1
1.1 How licensing works.....	1
1.2 License editions	2
1.3 How this guide is organized.....	2
Chapter 2. GRID Virtual GPU	3
2.1 vGPU License requirements	3
2.2 Licensing on Windows	5
2.3 Licensing on Linux	6
Chapter 3. GRID Virtual Workstation with passthrough	7
3.1 GRID Virtual Workstation features.....	7
3.2 Licensing on Windows	7
3.2.1 Disabling GRID Virtual Workstation	9
3.3 Licensing on Linux	10
Chapter 4. Advanced topics	11
4.1 Licenses obtained after boot.....	11
4.2 Operating with intermittent connectivity to the license server	11
4.3 Applying Windows license settings via registry	12
Chapter 5. Troubleshooting	14
5.1 Known issues	14
5.2 Troubleshooting steps	14

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 Sample gridd.conf for GRID vGPU	6
Figure 5 Managing Virtual Workstation Licensing in NVIDIA Control Panel	8
Figure 6 Applying GRID Virtual Workstation license.....	8
Figure 7 Disabling GRID Virtual Workstation	9
Figure 8 Sample gridd.conf for GRID Virtual Workstation	10
Figure 9 Configuring vGPU licensing via registry settings	13

LIST OF TABLES

Table 1 Virtual GPUs licensed on Tesla M6, M60	4
Table 2 Licensing registry settings.....	12

Chapter 1. INTRODUCTION

The following NVIDIA GRID™ products are available as licensed products on NVIDIA Tesla™ GPUs:

- ▶ Virtual Workstation
- ▶ Virtual PC
- ▶ Virtual Applications

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

1.1 HOW LICENSING WORKS

Figure 1 provides an overview of GRID licensing:

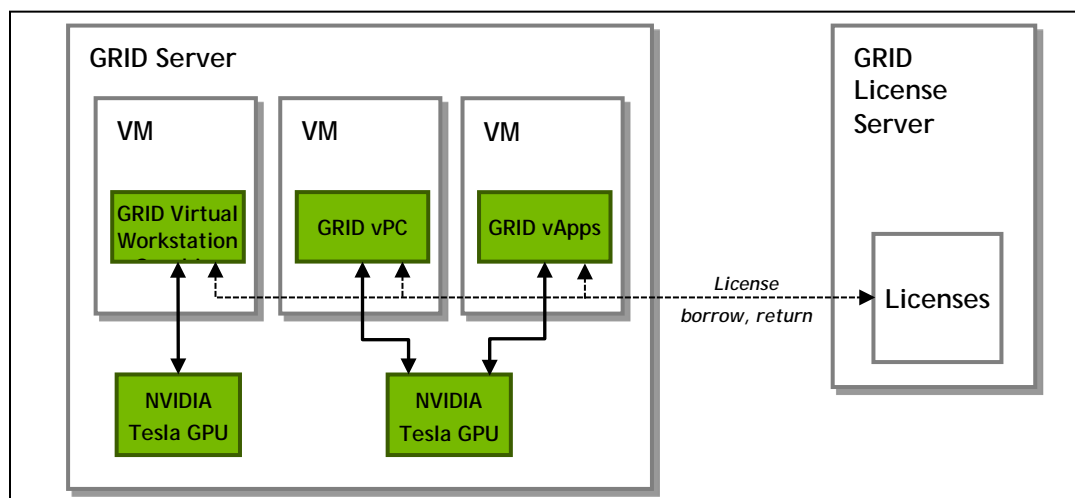


Figure 1 GRID licensing architecture

When enabled on Tesla GPUs, licensed editions of GRID Virtual Workstation or Virtual PC 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 Application	<ul style="list-style-type: none"> Virtual GPUs for virtual application computing
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 and high-end workstation computing Workstation graphics on GPU passthrough

Figure 2 GRID license editions

1.3 HOW THIS GUIDE IS ORGANIZED

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. 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.



Note: vGPUs that require licensing run at full capability even without a license. However, until a license is acquired, users are warned each time a vGPU fails to get a license. These warnings cease after 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	4096x2160	GRID Virtual Workstation
	M60-4Q	4096			
	M60-2Q	2048	4	4096x2160	
	M60-1Q	1024	2	2560x1600	
	M60-0Q	512	2	2560x1600	
	M60-1B	1024	4	2560x1600	GRID Virtual PC
	M60-0B	512	2		
	M60-8A	8192	1	1280x1024 ¹	GRID Virtual Application
	M60-4A	4096			
	M60-2A	2048			
M60-1A	1024				
Tesla M6	M6-8Q	8192	4	4096x2160	GRID Virtual Workstation
	M6-4Q	4096	4		
	M6-2Q	2048	4	4096x2160	
	M6-1Q	1024	2	2560x1600	
	M6-0Q	512	2	2560x1600	
	M6-1B	1024	4	2560x1600	GRID Virtual PC
	M6-0B	512	2		
	M6-8A	8192	1	1280x1024 ¹	GRID Virtual Application
	M6-4A	4096			
	M6-2A	2048			
M6-1A	1024				

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 license.



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

¹ Virtualized applications are rendered in an off-screen buffer. Therefore, the maximum resolution is independent of the maximum resolution of the display head.

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 4.3 on registry settings).

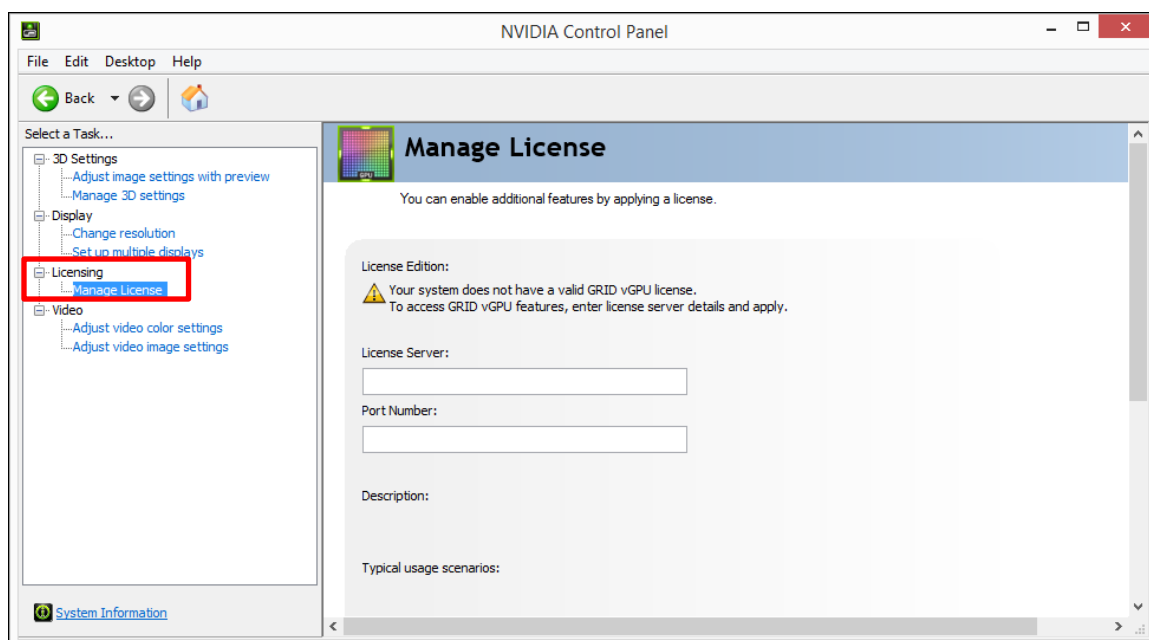


Figure 3 Managing vGPU licensing in NVIDIA Control Panel

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.mycompany.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. If the system fails to obtain a license, refer to Chapter 5 for guidance on troubleshooting.

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

2.3 LICENSING ON LINUX

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

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



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

Set `ServerURL` to the address and port number of your local NVIDIA GRID License Server. The address can be a fully-qualified domain name such as `gridlicense.mycompany.com`, or an IP address such as `10.31.20.45`. The port number is appended to the address with a colon, for example `:7070`.

Set `FeatureType` to 1, to license vGPU:

```
# /etc/nvidia/gridd.conf - Configuration file for NVIDIA Grid Daemon
# Description: Set License Server URL
# Data type: string
# Format: "<address>:<port>"
ServerUrl=hqdvls01.nvidia.com: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
```

Figure 4 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`:

```
[nvidia@localhost ~]$ sudo grep gridd /var/log/messages
...
Sep 13 15:40:06 localhost nvidia-gridd: Started (10430)
Sep 13 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 WITH PASSTHROUGH

This chapter describes how to enable GRID Virtual Workstation with passthrough 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 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 5.



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.

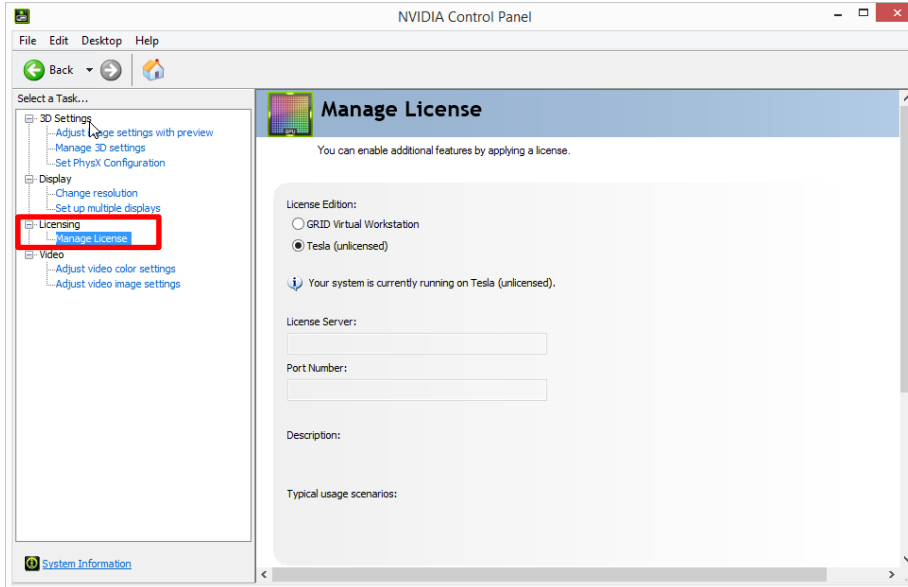


Figure 5 Managing Virtual Workstation Licensing in NVIDIA Control Panel

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 6). The address can be a fully-qualified domain name such as `gridlicense.mycompany.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.

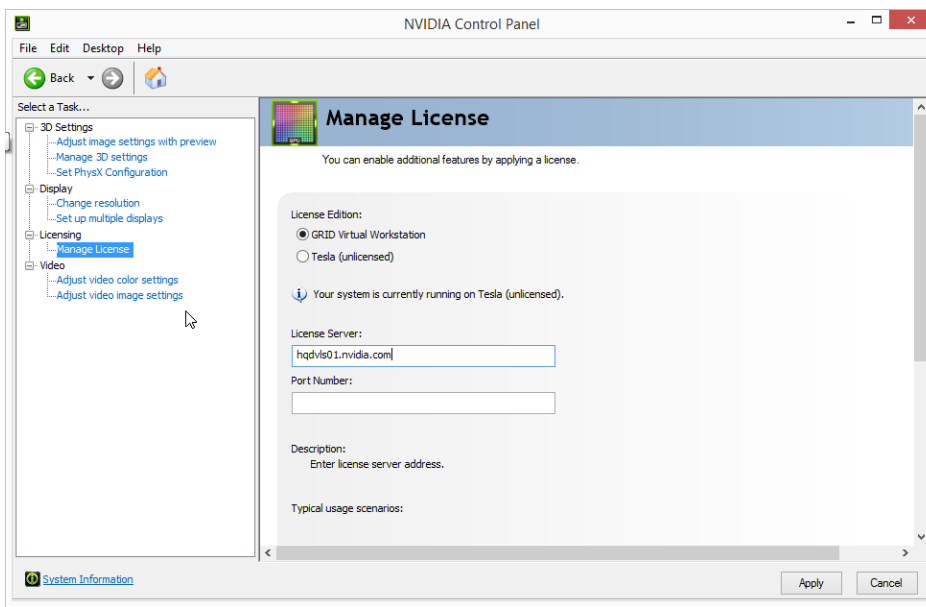


Figure 6 Applying GRID Virtual Workstation license

Select **Apply** to assign the settings. The system will request a license from the configured license server. 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.

3.2.1 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 7). The setting does not take effect until the next time the system is shutdown or rebooted; GRID Virtual Workstation features remain available until then.

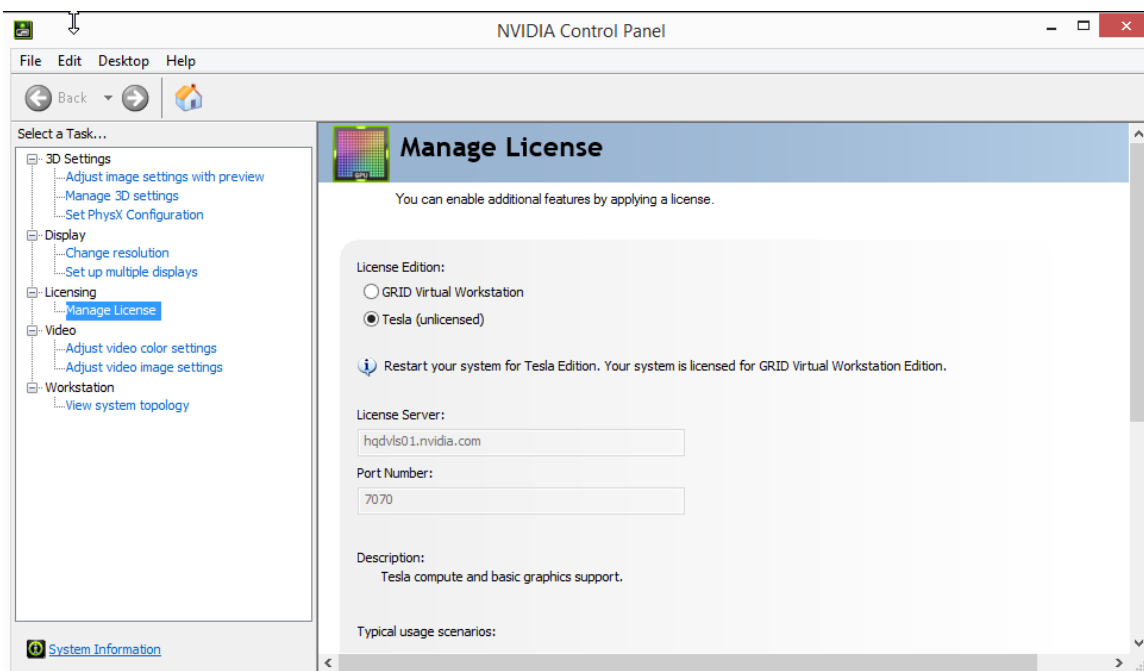


Figure 7 Disabling GRID Virtual Workstation

3.3 LICENSING ON LINUX

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

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



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

Set `ServerURL` to the address and port number of your local NVIDIA GRID License Server. The address can be a fully-qualified domain name such as `gridlicense.mycompany.com`, or an IP address such as `10.31.20.45`. The port number is appended to the address with a colon, for example `:7070`.

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

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

# Description: Set License Server URL
# Data type: string
# Format: "<address>:<port>"
ServerUrl=hqdvls01.nvidia.com: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
```

Figure 8 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`:

```
[nvidia@localhost ~]$ sudo grep gridd /var/log/messages
...
Sep 13 15:40:06 localhost nvidia-gridd: Started (10430)
Sep 13 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 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, as indicated by the license status on the Licensing tab of the NVIDIA Control Panel 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 full capability as described in section 2.1.

When a license is obtained, users are no longer warned that the vGPU has failed to obtain a license.

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 warn 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.

Table 2 Licensing registry settings

Figure 9 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`.

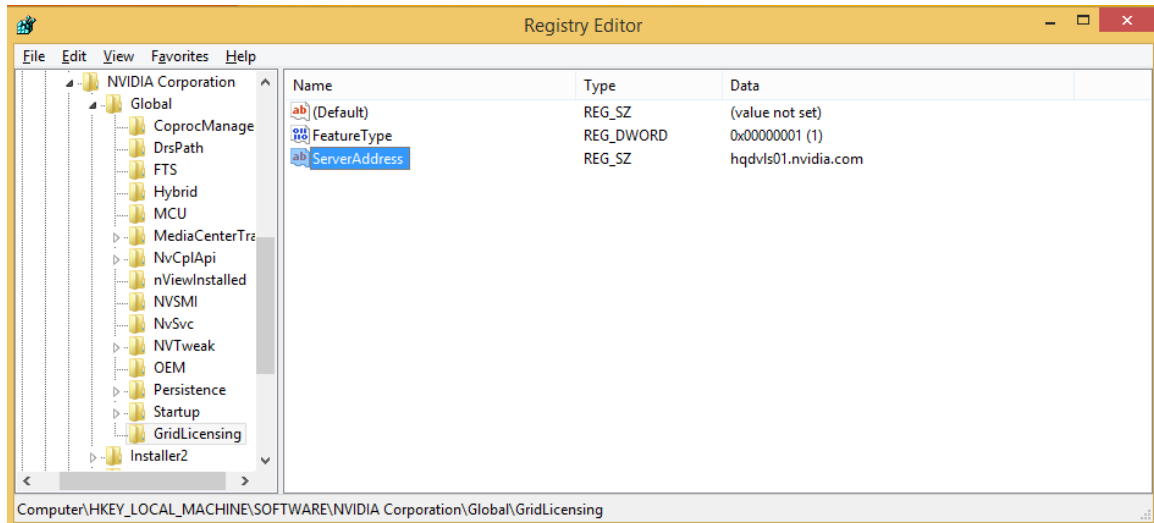


Figure 9 Configuring vGPU licensing via registry settings

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.