



GRID LICENSING

DU-07757-001 _v4.1 (GRID) | November 2016

User Guide



TABLE OF CONTENTS

Chapter 1. Introduction to NVIDIA GRID Licensing.....	1
1.1. How GRID licensing works.....	1
1.2. NVIDIA GRID license editions.....	2
1.3. How this guide is organized.....	2
Chapter 2. Licensing GRID Virtual GPU.....	3
2.1. vGPU license requirements.....	3
2.2. Licensing GRID vGPU on Windows.....	5
2.3. Licensing GRID vGPU on Linux.....	6
Chapter 3. Licensing GRID Virtual Workstation with passthrough.....	9
3.1. GRID Virtual Workstation features.....	9
3.2. Licensing GRID Virtual Workstation on Windows.....	9
3.3. Disabling GRID Virtual Workstation on Windows.....	11
3.4. Licensing GRID Virtual Workstation on Linux.....	12
Chapter 4. Advanced topics.....	14
4.1. Licenses obtained after boot.....	14
4.2. Operating with intermittent connectivity to the license server.....	14
4.3. Applying Windows license settings via registry.....	14
Chapter 5. Troubleshooting.....	17
5.1. Known issues.....	17
5.2. Troubleshooting steps.....	17

LIST OF FIGURES

Figure 1	GRID licensing architecture	1
Figure 2	Managing vGPU licensing in NVIDIA Control Panel	6
Figure 3	Managing Virtual Workstation Licensing in NVIDIA Control Panel	10
Figure 4	Applying a GRID Virtual Workstation license	11
Figure 5	Disabling GRID Virtual Workstation	12
Figure 6	Configuring vGPU licensing via registry settings	16

LIST OF TABLES

Table 1 GRID license editions 2

Table 2 Virtual GPUs licensed on Tesla M60 3

Table 3 Virtual GPUs licensed on Tesla M10 4

Table 4 Virtual GPUs licensed on Tesla M6 4

Table 5 Licensing registry settings 15

Chapter 1.

INTRODUCTION TO NVIDIA GRID LICENSING

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 GRID 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.



Information on setting up and running the GRID License Server is provided in [GRID License Server User Guide](#).

1.2. NVIDIA GRID 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:

Table 1 GRID license editions

GRID License Edition	GRID Features
GRID Virtual Application	▶ Virtual GPUs for virtual application computing
GRID Virtual PC	▶ Virtual GPUs for business desktop computing
GRID Virtual Workstation	▶ Virtual GPUs for midrange and high-end workstation computing ▶ Workstation graphics on GPU passthrough

1.3. How this guide is organized

The remainder of this guide is organized as follows:

- ▶ [Licensing GRID Virtual GPU](#) describes licensing of GRID Virtual GPU.
- ▶ [Licensing GRID Virtual Workstation with passthrough](#) describes licensing of GRID Virtual Workstation features with GPU passthrough.
- ▶ [Advanced topics](#) discusses advanced licensing settings.
- ▶ [Troubleshooting](#) provides guidance on troubleshooting.

Chapter 2.

LICENSING 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, Tesla M10, and Tesla 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.



vGPUs that require licensing run at full capability even without a license. However, on Windows, 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 2, Table 3, and Table 4 list the vGPU types that are supported on Tesla M60, Tesla M10, and Tesla M6, and the license edition that each vGPU type requires.

Table 2 Virtual GPUs licensed on Tesla M60

GRID Virtual GPU	Frame Buffer (Mbytes)	Virtual Display Heads	Maximum Resolution per Display Head	Minimum GRID License Edition Required
M60-8Q	8192	4	4096×2160	GRID Virtual Workstation
M60-4Q	4096	4	4096×2160	GRID Virtual Workstation
M60-2Q	2048	4	4096×2160	GRID Virtual Workstation
M60-1Q	1024	2	4096×2160	GRID Virtual Workstation
M60-0Q	512	2 ¹	2560×1600	GRID Virtual Workstation
M60-1B	1024	4	2560×1600	GRID Virtual PC

¹ This virtual GPU supports only 1 virtual display head on a Windows 10 guest OS.

GRID Virtual GPU	Frame Buffer (Mbytes)	Virtual Display Heads	Maximum Resolution per Display Head	Minimum GRID License Edition Required
M60-0B	512	2 ¹	2560×1600	GRID Virtual PC
M60-8A	8192	1	1280×1024 ²	GRID Virtual Application
M60-4A	4096	1	1280×1024 ²	GRID Virtual Application
M60-2A	2048	1	1280×1024 ²	GRID Virtual Application
M60-1A	1024	1	1280×1024 ²	GRID Virtual Application

Table 3 Virtual GPUs licensed on Tesla M10

GRID Virtual GPU	Frame Buffer (Mbytes)	Virtual Display Heads	Maximum Resolution per Display Head	Minimum GRID License Edition Required
M10-8Q	8192	4	4096×2160	GRID Virtual Workstation
M10-4Q	4096	4	4096×2160	GRID Virtual Workstation
M10-2Q	2048	4	4096×2160	GRID Virtual Workstation
M10-1Q	1024	2	4096×2160	GRID Virtual Workstation
M10-0Q	512	2 ¹	2560×1600	GRID Virtual Workstation
M10-1B	1024	4	2560×1600	GRID Virtual PC
M10-0B	512	2 ¹	2560×1600	GRID Virtual PC
M10-8A	8192	1	1280×1024 ²	GRID Virtual Application
M10-4A	4096	1	1280×1024 ²	GRID Virtual Application
M10-2A	2048	1	1280×1024 ²	GRID Virtual Application
M10-1A	1024	1	1280×1024 ²	GRID Virtual Application

Table 4 Virtual GPUs licensed on Tesla M6

GRID Virtual GPU	Frame Buffer (Mbytes)	Virtual Display Heads	Maximum Resolution per Display Head	Minimum GRID License Edition Required
M6-8Q	8192	4	4096×2160	GRID Virtual Workstation
M6-4Q	4096	4	4096×2160	GRID Virtual Workstation
M6-2Q	2048	4	4096×2160	GRID Virtual Workstation

² Virtualized applications are rendered in an off-screen buffer. Therefore, the maximum resolution is independent of the maximum resolution of the display head.

GRID Virtual GPU	Frame Buffer (Mbytes)	Virtual Display Heads	Maximum Resolution per Display Head	Minimum GRID License Edition Required
M6-1Q	1024	2	4096×2160	GRID Virtual Workstation
M6-0Q	512	2 ¹	2560×1600	GRID Virtual Workstation
M6-1B	1024	4	2560×1600	GRID Virtual PC
M6-0B	512	2 ¹	2560×1600	GRID Virtual PC
M6-8A	8192	1	1280×1024 ²	GRID Virtual Application
M6-4A	4096	1	1280×1024 ²	GRID Virtual Application
M6-2A	2048	1	1280×1024 ²	GRID Virtual Application
M6-1A	1024	1	1280×1024 ²	GRID Virtual Application

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.



For more information on the virtual GPU types listed in [Table 2](#) and [Table 4](#), consult *GRID Virtual GPU User Guide*.

2.2. Licensing GRID vGPU on Windows

To license GRID vGPU on Windows, follow these steps:

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, as shown in [Figure 2](#).



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 [Applying Windows license settings via registry](#) on registry settings).

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




Figure 2 Managing vGPU licensing in NVIDIA Control Panel

3. 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`.
4. Leave the **Port Number** field unset.
It will default to 7070, which is the default port number used by NVIDIA GRID License Server.
5. Click **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 [Troubleshooting](#) for guidance on troubleshooting.

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

2.3. Licensing GRID vGPU on Linux

To license GRID vGPU on Linux:

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 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`.

- Optional: Set `ServerPort` to the port number of your local NVIDIA GRID 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 GRID License Server.

- Set `FeatureType` to 1, to license vGPU.
- 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 15:40:06 localhost nvidia-gridd: Started (4293)
Aug 5 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.

`gridd.conf` file for GRID vGPU

The following example shows a `gridd.conf` file for GRID vGPU in which `ServerAddress` is set to `gridlicense.example.com`, `ServerPort` is set to 7070, and `FeatureType` is set to 1.

```
# /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
```

Chapter 3.

LICENSING 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 2560×1600)
- ▶ Workstation-specific graphics features and accelerations
- ▶ Certified drivers for professional applications

3.2. Licensing GRID Virtual Workstation on Windows

To enable GRID Virtual Workstation on Windows, follow these steps:

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, as shown in [Figure 3](#).



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.

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




Figure 3 Managing Virtual Workstation Licensing in NVIDIA Control Panel

3. Select **GRID Virtual Workstation**, and enter the address of your local GRID License Server in the License Server field (see [Figure 4](#)).

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




Figure 4 Applying a GRID Virtual Workstation license

4. Leave the **Port Number** field unset.
It will default to 7070, which is the default port number used by NVIDIA GRID License Server.
5. Click **Apply** to assign the settings.
The **License Edition** section will indicate that your system is running in GRID Virtual Workstation 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, refer to [Troubleshooting](#) for guidance on troubleshooting.

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

3.3. Disabling GRID Virtual Workstation on Windows

To disable the GRID Virtual Workstation licensed feature on Windows, follow these steps:

1. Open **NVIDIA Control Panel**.
2. In the **Manage License** task, select **Tesla (unlicensed)**, as shown in [Figure 5](#).




Figure 5 Disabling GRID Virtual Workstation

3. Click **Apply**.

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

3.4. Licensing GRID Virtual Workstation on Linux

To license GRID Virtual Workstation on Linux, follow these steps:

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 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`.
3. Optional: Set `ServerPort` to the port number of your local NVIDIA GRID 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 GRID License Server.
4. Set `FeatureType` to 2, to license GRID Virtual Workstation.
5. Save your changes to the `/etc/nvidia/gridd.conf` file.
6. Restart the `nvidia-gridd` service:


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

The service should automatically obtain a license.

7. 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 GRID vGPU.

`gridd.conf` file for GRID Virtual Workstation

The following example shows a `gridd.conf` file for GRID Virtual Workstation in which `ServerAddress` is set to `gridlicense.example.com`, `ServerPort` is set to 7070, and `FeatureType` is set to 2.

```
# /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
```

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 [vGPU license requirements](#).

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 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 5](#).

Table 5 Licensing registry settings

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.

Figure 6 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 6 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

© 2013-2016 NVIDIA Corporation. All rights reserved.