



GRID VGAMING R440 GA7.5

DA-09680-001 _v03/20 | March 2020

Application Note



TABLE OF CONTENTS

Chapter 1. GRID vGaming Release Information.....	1
Chapter 2. Validated Platforms.....	2
2.1. Supported NVIDIA GPUs and Validated Server Platforms.....	2
2.2. Hypervisor Software Releases.....	3
2.3. Guest OS Support.....	3
2.3.1. Windows Guest OS Support.....	4
2.3.2. Linux Guest OS Support.....	4
2.4. NVIDIA CUDA Toolkit Version Support.....	4
Chapter 3. GRID vGaming Virtual GPU Types.....	5
3.1. Tesla P40 Virtual GPU Types.....	5
3.2. Tesla P4 Virtual GPU Types.....	6
3.3. Tesla T10 16GB Virtual GPU Types.....	6
3.4. Tesla T4 Virtual GPU Types.....	7
3.5. Tesla V100 PCIe 32GB Virtual GPU Types.....	7
3.6. Quadro RTX 8000 Virtual GPU Types.....	7
3.7. Quadro RTX 6000 Virtual GPU Types.....	8
Chapter 4. GRID vGaming Licensing.....	9
4.1. GRID vGaming Licensed Deployments.....	9
4.2. How GRID vGaming Licensing Is Enforced.....	9
4.3. Licensing an NVIDIA vGPU or Physical GPU.....	9
Appendix A. GRID vGaming Resources.....	11

Chapter 1.

GRID VGAMING RELEASE INFORMATION

This *Application Note* supplements the documentation for NVIDIA vGPU software with information that is specific to GRID vGaming.

GRID vGaming supports NVIDIA vGPU, GPU pass through, and bare-metal deployments.

This release of GRID vGaming includes the software listed in the following table:

Software	Version
NVIDIA Virtual GPU Manager	440.68
NVIDIA Windows driver	442.59
NVIDIA Linux driver	440.68



Caution

If you install the wrong GRID vGaming packages for the version of the hypervisor you are using, NVIDIA Virtual GPU Manager will fail to load.

The vGPU Manager and guest VM drivers must be installed together. Older VM drivers will not function correctly with this release of vGPU Manager. Similarly, older releases of vGPU Manager will not function correctly with this release of the guest VM drivers.

This requirement does not apply to the NVIDIA vGPU software license sever. All releases of GRID vGaming are compatible with **all** releases of the license server.

Chapter 2.

VALIDATED PLATFORMS

This release of GRID vGaming provides support for several NVIDIA GPUs on validated server hardware platforms, hypervisor software versions, and guest operating systems. It also supports NVIDIA CUDA Toolkit 10.2.

2.1. Supported NVIDIA GPUs and Validated Server Platforms

This release of GRID vGaming provides support for the following NVIDIA GPUs running on validated server hardware platforms:

- ▶ GPUs based on the NVIDIA Maxwell™ graphic architecture:
 - ▶ Tesla M60 in GPU pass through mode and bare-metal deployments only (**not** supported with NVIDIA vGPU)
- ▶ GPUs based on the NVIDIA Pascal™ architecture:
 - ▶ Tesla P4
 - ▶ Tesla P40
- ▶ GPUs based on the NVIDIA Volta architecture:
 - ▶ Tesla V100 PCIe 32GB
- ▶ GPUs based on the NVIDIA Turing™ architecture:
 - ▶ Tesla T4
 - ▶ Tesla T10
 - ▶ Quadro RTX 6000 in displayless mode
 - ▶ Quadro RTX 8000 in displayless mode



These GPUs are supported as a secondary device in a bare-metal deployment.

In displayless mode, local physical display connectors are disabled.

For a list of validated server platforms, refer to [NVIDIA GRID Certified Servers](#).



Tesla M60 GPUs support compute mode and graphics mode. NVIDIA vGPU requires GPUs that support both modes to operate in graphics mode.

Recent Tesla M60 GPUs are supplied in graphics mode. However, your GPU might be in compute mode if it is an older Tesla M60 GPU, or if its mode has previously been changed.

To configure the mode of Tesla M60 GPUs, use the `gpumodeswitch` tool provided with NVIDIA vGPU software releases.

2.2. Hypervisor Software Releases

This release of GRID vGaming is supported on the hypervisor software releases listed in the table.



If a specific release, even an update release, is not listed, it's **not** supported.

Software	Releases Supported	Notes
Citrix Hypervisor	8.0 7.6	The RTM build and compatible cumulative update releases are supported.
Red Hat Enterprise Linux with KVM	7.6, 7.5	All NVIDIA GPUs that GRID vGaming supports are supported with vGPU and in pass-through mode.
	7.2 through 7.4	All NVIDIA GPUs that GRID vGaming supports are supported in pass-through mode only.
	7.0, 7.1	Only the Tesla M60 GPU is supported in pass-through mode only.

2.3. Guest OS Support

GRID vGaming supports several Windows releases and Linux distributions as a guest OS.



Use only a guest OS release that is listed as supported by GRID vGaming with your virtualization software. To be listed as supported, a guest OS release must be supported not only by GRID vGaming, but also by your virtualization software. NVIDIA cannot support guest OS releases that your virtualization software does not support.

GRID vGaming supports **only** 64-bit guest operating systems. No 32-bit guest operating systems are supported.

2.3.1. Windows Guest OS Support

GRID vGaming supports **only** the 64-bit Windows releases listed as a guest OS.



If a specific release, even an update release, is not listed, it's **not** supported.

- ▶ Windows Server 2019 (**not** supported on Citrix Hypervisor)
- ▶ Windows Server 2016
- ▶ Windows Server 2012 R2

2.3.2. Linux Guest OS Support

GRID vGaming supports **only** the Linux distributions listed as a guest OS.



If a specific release, even an update release, is not listed, it's **not** supported.

- ▶ Ubuntu 18.04

2.4. NVIDIA CUDA Toolkit Version Support

This release of GRID vGaming supports NVIDIA CUDA Toolkit 10.2.

For more information about NVIDIA CUDA Toolkit, see [CUDA Toolkit 10.2 Documentation](#).



If you are using NVIDIA vGPU software with CUDA on Linux, avoid conflicting installation methods by installing CUDA from a distribution-independent runfile package. Do not install CUDA from distribution-specific RPM or Deb package.

To ensure that the NVIDIA vGPU software graphics driver is not overwritten when CUDA is installed, deselect the CUDA driver when selecting the CUDA components to install.

For more information, see [NVIDIA CUDA Installation Guide for Linux](#).

Chapter 3.

GRID VGAMING VIRTUAL GPU TYPES

A physical GPU can support several different types of virtual GPU (vGPU). Different GRID vGaming vGPU types have different amounts of frame buffer, but they all support a single virtual display head with a maximum resolution of 4K. They are grouped into different series according to the different classes of workload for which they are optimized. Each series is identified by a three-letter prefix in the vGPU type name.

Series	Optimal Workload
GTX series	GRID vGaming on GPUs based on the NVIDIA Pascal™ architecture and the NVIDIA Volta architecture
RTX series	GRID vGaming on GPUs based on the NVIDIA Turing™ architecture

The number after the board type in the vGPU type name denotes the amount of frame buffer that is allocated to a vGPU of that type. For example, a vGPU of type RTX T10-4 is allocated 4096 Mbytes of frame buffer on a Tesla T10 16GB board.

Due to their differing resource requirements, the maximum number of vGPUs that can be created simultaneously on a physical GPU varies according to the vGPU type. For example, a Tesla T10 16GB board can support up to four RTX T10-4 vGPUs on its physical GPU, but only two RTX T10-8 vGPUs, and only one RTX T10-16 vGPU.



GRID vGaming is a licensed product on all supported GPU boards. A GRID vGaming software license is required to enable all vGPU features within the guest VM.

3.1. Tesla P40 Virtual GPU Types

Physical GPUs per board: 1

Virtual GPU Type	Intended Use Case	Frame Buffer (MB)	Virtual Display Heads	Maximum Resolution per Display Head	Maximum vGPUs per GPU	Maximum vGPUs per Board	Required License Edition
GTX P40-24	GRID vGaming	24576	1	4096×2160	1	1	GRID vGaming
GTX P40-12	GRID vGaming	12288	1	4096×2160	2	2	GRID vGaming
GTX P40-8	GRID vGaming	8192	1	4096×2160	3	3	GRID vGaming
GTX P40-6	GRID vGaming	6144	1	4096×2160	4	4	GRID vGaming

3.2. Tesla P4 Virtual GPU Types

Physical GPUs per board: 1

Virtual GPU Type	Intended Use Case	Frame Buffer (MB)	Virtual Display Heads	Maximum Resolution per Display Head	Maximum vGPUs per GPU	Maximum vGPUs per Board	Required License Edition
GTX P4-8	GRID vGaming	8192	1	4096×2160	1	1	GRID vGaming
GTX P4-4	GRID vGaming	4096	1	4096×2160	2	2	GRID vGaming
GTX P4-2	GRID vGaming	2048	1	4096×2160	4	4	GRID vGaming

3.3. Tesla T10 16GB Virtual GPU Types

Physical GPUs per board: 1

Virtual GPU Type	Intended Use Case	Frame Buffer (MB)	Virtual Display Heads	Maximum Resolution per Display Head	Maximum vGPUs per GPU	Maximum vGPUs per Board	Required License Edition
RTX T10-16	GRID vGaming	16384	1	4096×2160	1	1	GRID vGaming
RTX T10-8	GRID vGaming	8192	1	4096×2160	2	2	GRID vGaming
RTX T10-4	GRID vGaming	4096	1	4096×2160	4	4	GRID vGaming

3.4. Tesla T4 Virtual GPU Types

Physical GPUs per board: 1

Virtual GPU Type	Intended Use Case	Frame Buffer (MB)	Virtual Display Heads	Maximum Resolution per Display Head	Maximum vGPUs per GPU	Maximum vGPUs per Board	Required License Edition
RTX T4-16	GRID vGaming	16384	1	4096×2160	1	1	GRID vGaming
RTX T4-8	GRID vGaming	8192	1	4096×2160	2	2	GRID vGaming
RTX T4-4	GRID vGaming	4096	1	4096×2160	4	4	GRID vGaming

3.5. Tesla V100 PCIe 32GB Virtual GPU Types

Physical GPUs per board: 1

Virtual GPU Type	Intended Use Case	Frame Buffer (MB)	Virtual Display Heads	Maximum Resolution per Display Head	Maximum vGPUs per GPU	Maximum vGPUs per Board	Required License Edition
GTX V100D-32	GRID vGaming	32768	1	4096×2160	1	1	GRID vGaming
GTX V100D-16	GRID vGaming	16384	1	4096×2160	2	2	GRID vGaming
GTX V100D-8	GRID vGaming	8192	1	4096×2160	4	4	GRID vGaming

3.6. Quadro RTX 8000 Virtual GPU Types

Physical GPUs per board: 1

Virtual GPU Type	Intended Use Case	Frame Buffer (MB)	Virtual Display Heads	Maximum Resolution per Display Head	Maximum vGPUs per GPU	Maximum vGPUs per Board	Required License Edition
RTX8000-48	GRID vGaming	49152	1	4096×2160	1	1	GRID vGaming
RTX8000-24	GRID vGaming	24576	1	4096×2160	2	2	GRID vGaming

Virtual GPU Type	Intended Use Case	Frame Buffer (MB)	Virtual Display Heads	Maximum Resolution per Display Head	Maximum vGPUs per GPU	Maximum vGPUs per Board	Required License Edition
RTX8000-16	GRID vGaming	16384	1	4096×2160	3	3	GRID vGaming
RTX8000-12	GRID vGaming	12288	1	4096×2160	4	4	GRID vGaming

3.7. Quadro RTX 6000 Virtual GPU Types

Physical GPUs per board: 1

Virtual GPU Type	Intended Use Case	Frame Buffer (MB)	Virtual Display Heads	Maximum Resolution per Display Head	Maximum vGPUs per GPU	Maximum vGPUs per Board	Required License Edition
RTX6000-24	GRID vGaming	24576	1	4096×2160	1	1	GRID vGaming
RTX6000-12	GRID vGaming	12288	1	4096×2160	2	2	GRID vGaming
RTX6000-8	GRID vGaming	8192	1	4096×2160	3	3	GRID vGaming
RTX6000-6	GRID vGaming	6144	1	4096×2160	4	4	GRID vGaming

Chapter 4.

GRID VGAMING LICENSING

GRID vGaming is a licensed product on NVIDIA GPUs that support it.

4.1. GRID vGaming Licensed Deployments

The GRID vGaming license supports different classes of deployments.

Licensed Product	Target Users	Supported Deployments
GRID vGaming	Online gamers using a cloud gaming service	<ul style="list-style-type: none">▶ GTX series NVIDIA vGPUs▶ RTX series NVIDIA vGPUs▶ GPU pass through▶ Bare metal

4.2. How GRID vGaming Licensing Is Enforced

All GRID vGaming deployments require a GRID vGaming license. Licensing for all GRID vGaming deployments is enforced through software.

4.3. Licensing an NVIDIA vGPU or Physical GPU

Licensing through **NVIDIA Control Panel** on Windows and **NVIDIA X Server Settings** on Linux is disabled for GRID vGaming guest drivers. How to license an NVIDIA vGPU or a pass-through or bare-metal physical GPU for GRID vGaming depends on your guest OS:

- ▶ On Windows, you must use the Windows registry key `HKEY_LOCAL_MACHINE\SOFTWARE\NVIDIA Corporation\Global\GridLicensing`.
- ▶ On Linux, you must use the `/etc/nvidia/gridd.conf` configuration file.

If you do not require high availability for the license server and are using the default license server port (7070), you need set only the configuration parameters in the following table.

Configuration Parameter	Windows Data Type	Setting
ServerAddress	String (REG_SZ)	<p>The address of your license server.</p> <p>The address can be a fully-qualified domain name such as <code>gridlicense1.example.com</code>, or an IP address such as <code>10.31.20.45</code>.</p>
FeatureType	DWord (REG_DWORD)	<p>NVIDIA vGPU:</p> <ul style="list-style-type: none"> ▶ Windows: Do not change the value of this registry key. ▶ Linux: Set to 1.
		<p>Physical GPU: Set to 3 to specify a GRID vGaming license.</p>

Appendix A.

GRID VGAMING RESOURCES

Getting Started

- ▶ *Virtual GPU Software Quick Start Guide*
- ▶ *Virtual GPU Software for Citrix Hypervisor Release Notes*
- ▶ *Virtual GPU Software for Red Hat Enterprise Linux with KVM Release Notes*

Game Ready Driver Release Information

- ▶ *Release 440 Driver for Windows Release Notes*
- ▶ *NVIDIA Accelerated Linux Graphics Driver README and Installation Guide*

License Server Installation and Setup

- ▶ *Virtual GPU License Server Release Notes*
- ▶ *Virtual GPU License Server User Guide*

GRID vGaming Installation, Setup, and Client-Side Licensing

- ▶ *Virtual GPU Software User Guide*
- ▶ *Virtual GPU Client Licensing User Guide*
- ▶ *gpumodeswitch User Guide*

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.