



VIRTUAL GPU SOFTWARE R390 FOR MICROSOFT WINDOWS SERVER

RN-08686-001 _v6.0 through 6.4 Revision 02 | May 2020

Release Notes



TABLE OF CONTENTS

Chapter 1. Release Notes	1
1.1. Updates in Release 6.0.....	1
1.2. Updates in Release 6.1.....	2
1.3. Updates in Release 6.2.....	2
1.4. Updates in Release 6.3.....	2
1.5. Updates in Release 6.4.....	2
Chapter 2. Validated Platforms	3
2.1. Supported NVIDIA GPUs and Validated Server Platforms.....	3
2.2. Hypervisor Software Releases.....	4
2.3. Guest OS Support.....	4
2.3.1. Windows Guest OS Support.....	4
2.3.2. Linux Guest OS Support.....	5
Chapter 3. Security Updates	6
3.1. Since 6.4: Restricting Access to GPU Performance Counters.....	6
3.1.1. Windows: Restricting Access to GPU Performance Counters for One User by Using NVIDIA Control Panel.....	6
3.1.2. Windows: Restricting Access to GPU Performance Counters Across an Enterprise by Using a Registry Key.....	7
3.1.3. Linux Guest VMs: Restricting Access to GPU Performance Counters.....	7
Chapter 4. Known Issues	9
4.1. Frame capture while the interactive logon message is displayed returns blank screen.....	9
4.2. Screen resolution reverts to a lower value after a VM is rebooted.....	10
4.3. 6.1 Only: Benign Calling load_byte_array(tra) messages are logged.....	10
4.4. 6.0-6.2 Only: Licensing pop-up windows contain the text microsoft.explorer.notification.	11
4.5. Resolution is not updated after a VM acquires a license and is restarted.....	11
4.6. A segmentation fault in Dbus code causes nvidia-gridd to exit on Red Hat Enterprise Linux and CentOS.....	12
4.7. No Manage License option available in NVIDIA X Server Settings by default.....	13
4.8. Licenses remain checked out when VMs are forcibly powered off.....	13

Chapter 1.

RELEASE NOTES

These *Release Notes* summarize current status, information on validated platforms, and known issues with NVIDIA vGPU software and associated hardware on Microsoft Windows Server.

The releases in this release family of NVIDIA vGPU software include the software listed in the following table:

Software	6.0	6.1	6.2	6.3	6.4
NVIDIA Windows driver	391.03	391.58	391.81	392.05	392.37
NVIDIA Linux driver	390.42	390.57	390.75	390.96	390.115

All releases of NVIDIA vGPU software are compatible with **all** releases of the license server.

1.1. Updates in Release 6.0

New Features in Release 6.0

- ▶ Change of maximum resolution for unlicensed GPUs based on the Pascal architecture to 1280×1024
- ▶ Plain-text logging on Windows of significant licensing events
- ▶ New setting `EnableLogging` for disabling or enabling logging of significant licensing events
- ▶ Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 6.0

- ▶ Support for GPUs based on the NVIDIA[®] Volta architecture
- ▶ Support for Windows 10 Fall Creators Update (1709) as a guest OS

1.2. Updates in Release 6.1

New Features in Release 6.1

- ▶ Change in behavior to enable logging of licensing events on Windows by default
- ▶ Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 6.1

- ▶ Support for the Tesla V100 SXM2 32GB GPU
- ▶ Support for the Tesla V100 PCIe 32GB GPU

1.3. Updates in Release 6.2

New Features in Release 6.2

- ▶ Miscellaneous bug fixes

1.4. Updates in Release 6.3

New Features in Release 6.3

- ▶ Inclusion of the name and version of the licensed product in logged license acquisition events
- ▶ Ability to disable pop-up notifications for license state changes
- ▶ Miscellaneous bug fixes

1.5. Updates in Release 6.4

New Features in Release 6.4

- ▶ Miscellaneous bug fixes
- ▶ Security updates - see [Security Updates](#)

Chapter 2.

VALIDATED PLATFORMS

This release family of NVIDIA vGPU software provides support for several NVIDIA GPUs on validated server hardware platforms, Microsoft Windows Server hypervisor software versions, and guest operating systems.

2.1. Supported NVIDIA GPUs and Validated Server Platforms

This release of NVIDIA vGPU software provides support for the following NVIDIA GPUs on Microsoft Windows Server, running on validated server hardware platforms:

- ▶ Tesla M6
- ▶ Tesla M10
- ▶ Tesla M60
- ▶ Tesla P4
- ▶ Tesla P6
- ▶ Tesla P40
- ▶ Tesla P100 PCIe 16 GB
- ▶ Tesla P100 SXM2 16 GB
- ▶ Tesla P100 PCIe 12GB
- ▶ Tesla V100 SXM2
- ▶ Tesla V100 SXM2 32GB
- ▶ Tesla V100 PCIe
- ▶ Tesla V100 PCIe 32GB
- ▶ Tesla V100 FHHL



These GPUs are supported as a secondary device in a bare-metal deployment. Tesla M6 is also supported as the primary display device in a bare-metal deployment.

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

2.2. Hypervisor Software Releases

This release supports **only** the hypervisor software versions listed in the table.



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

Software	Version Supported
Microsoft Windows Server	Windows Server 2016 1709 with Hyper-V role Windows Server 2016 1607 with Hyper-V role

2.3. Guest OS Support

NVIDIA vGPU software supports several Windows releases and Linux distributions as a guest OS using GPU pass-through.

Microsoft Windows Server with Hyper-V role supports GPU pass-through over Microsoft Virtual PCI bus. This bus is supported through paravirtualized drivers.



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

In pass-through mode, GPUs based on the Pascal architecture or Volta architecture support only 64-bit guest operating systems. No 32-bit guest operating systems are supported in pass-through mode for these GPUs.

2.3.1. Windows Guest OS Support

NVIDIA vGPU software supports **only** the Windows releases listed as a guest OS on Microsoft Windows Server.



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

- ▶ Windows Server 2016 1607, 1709
- ▶ Windows Server 2012 R2 with patch `Windows8.1-KB3133690-x64.msu`
- ▶ Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703), Fall Creators Update (1709) (64-bit)
- ▶ Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703), Fall Creators Update (1709) (32-bit) on Tesla M6, Tesla M10, and Tesla M60 GPUs only

2.3.2. Linux Guest OS Support

NVIDIA vGPU software supports **only** the 64-bit Linux distributions listed as a guest OS on Microsoft Windows Server.



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

- ▶ Red Hat Enterprise Linux 7.0-7.4
- ▶ CentOS 7.0-7.4
- ▶ Ubuntu 16.04 LTS
- ▶ SUSE Linux Enterprise Server 12 SP2

Chapter 3.

SECURITY UPDATES

3.1. Since 6.4: Restricting Access to GPU Performance Counters

The NVIDIA graphics driver contains a vulnerability (CVE-2018-6260) that may allow access to application data processed on the GPU through a side channel exposed by the GPU performance counters. To address this vulnerability, update the driver and restrict access to GPU performance counters to allow access only by administrator users and users who need to use CUDA profiling tools.

The GPU performance counters that are affected by this vulnerability are the hardware performance monitors used by the CUDA profiling tools such as CUPTI, Nsight Graphics, and Nsight Compute. These performance counters are exposed on the hypervisor host and in guest VMs only as follows:

- ▶ On the hypervisor host, they are always exposed. However, the Virtual GPU Manager does not access these performance counters and, therefore, is not affected.
- ▶ In Windows and Linux guest VMs, they are exposed **only** in VMs configured for GPU pass through. They are not exposed in VMs configured for NVIDIA vGPU.

3.1.1. Windows: Restricting Access to GPU Performance Counters for One User by Using NVIDIA Control Panel

Perform this task from the guest VM to which the GPU is passed through.

Ensure that you are running **NVIDIA Control Panel** version 8.1.950.

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 GPU Performance Counters** task in the **Developer** section of the navigation pane.
3. Complete the task by following the instructions in the **Manage GPU Performance Counters > Developer** topic in the **NVIDIA Control Panel** help.

3.1.2. Windows: Restricting Access to GPU Performance Counters Across an Enterprise by Using a Registry Key

You can use a registry key to restrict access to GPU Performance Counters for all users who log in to a Windows guest VM. By incorporating the registry key information into a script, you can automate the setting of this registry for all Windows guest VMs across your enterprise.

Perform this task from the guest VM to which the GPU is passed through.



Caution Only enterprise administrators should perform this task. Changes to the Windows registry must be made with care and system instability can result if registry keys are incorrectly set.

1. Set the `RmProfilingAdminOnly` Windows registry key to 1.

```
[HKLM\SYSTEM\CurrentControlSet\Services\nvlddmkm\Global\NVTweak]
Value: "RmProfilingAdminOnly"
Type: DWORD
Data: 00000001
```

The data value 1 restricts access, and the data value 0 allows access, to application data processed on the GPU through a side channel exposed by the GPU performance counters.

2. Restart the VM.

3.1.3. Linux Guest VMs: Restricting Access to GPU Performance Counters

On systems where unprivileged users don't need to use GPU performance counters, restrict access to these counters to system administrators, namely users with the `CAP_SYS_ADMIN` capability set. By default, the GPU performance counters are not restricted to users with the `CAP_SYS_ADMIN` capability.

Perform this task from the guest VM to which the GPU is passed through.

This task requires `sudo` privileges.

1. Log in to the guest VM.
2. Set the kernel module parameter `NVreg_RestrictProfilingToAdminUsers` to 1 by adding this parameter to the `/etc/modprobe.d/nvidia.conf` file.
 - ▶ If you are setting only this parameter, add an entry for it to the `/etc/modprobe.d/nvidia.conf` file as follows:

```
options nvidia
NVreg_RegistryDwords="NVreg_RestrictProfilingToAdminUsers=1"
```

- ▶ If you are setting multiple parameters, set them in a single entry as in the following example:

```
options nvidia NVreg_RegistryDwords="RmPVMRL=0x0"  
"NVreg_RestrictProfilingToAdminUsers=1"
```

If the `/etc/modprobe.d/nvidia.conf` file does not already exist, create it.

3. Restart the VM.

Chapter 4.

KNOWN ISSUES

4.1. Frame capture while the interactive logon message is displayed returns blank screen

Description

Because of a known limitation with NvFBC, a frame capture while the interactive logon message is displayed returns a blank screen.

An NvFBC session can capture screen updates that occur after the session is created. Before the logon message appears, there is no screen update after the message is shown and, therefore, a black screen is returned instead. If the NvFBC session is created after this update has occurred, NvFBC cannot get a frame to capture.

Workaround

Press **Enter** or wait for the screen to update for NvFBC to capture the frame.

Status

Not a bug

Ref.

2115733

4.2. Screen resolution reverts to a lower value after a VM is rebooted

Description

When a VM is booted, the NVIDIA vGPU software graphics driver is initially unlicensed. Screen resolution is limited to a maximum of 1280×1024 until the VM requires a license for NVIDIA vGPU software. Because the higher resolutions are not available, the OS falls back to next available resolution in its mode list (for example, 1366×768) even if the resolution for the VM had previously been set to a higher value (for example, 1920×1080). After the license has been acquired, the OS does not attempt to set the resolution to a higher value.

This behavior is the expected behavior for licensed NVIDIA vGPU software products.

Workaround

Manually set the screen resolution to the required higher value after the VM has acquired the NVIDIA vGPU software license.

Status

Resolved in NVIDIA vGPU software release 6.2.

Ref.

2104867

4.3. 6.1 Only: Benign Calling `load_byte_array(tra)` messages are logged

Description

In Linux guest VMs, the following messages from the `nvidia-gridd` daemon are logged in `/var/log/syslog`:

```
May 21 18:36:39 test-HVM-domU nvidia-gridd: Started (657)
May 21 18:36:39 test-HVM-domU nvidia-gridd: Ignore Service Provider Licensing.
May 21 18:36:39 test-HVM-domU nvidia-gridd: Calling load_byte_array(tra)
May 21 18:36:41 test-HVM-domU nvidia-gridd: Acquiring license for GRID vGPU
Edition.
May 21 18:36:41 test-HVM-domU nvidia-gridd: Calling load_byte_array(tra)
May 21 18:36:43 test-HVM-domU nvidia-gridd: License acquired successfully.
Server URL : http://192.0.2.117:7070/request
```

Workaround

Ignore these messages as they are benign.

Status

Resolved in NVIDIA vGPU software release 6.2.

Ref. #

200407382

4.4. 6.0-6.2 Only: Licensing pop-up windows contain the text `microsoft.explorer.notification`

Description

On Windows 10 Creators Update (1703), licensing pop-up windows contain the text `microsoft.explorer.notification`.

Version

Windows 10 Creators Update (1703)

Status

Resolved in NVIDIA vGPU software release 6.3

Ref. #

200346607

4.5. Resolution is not updated after a VM acquires a license and is restarted

Description

In a Red Enterprise Linux 7.3 guest VM, an increase in resolution from 1024×768 to 2560×1600 is not applied after a license is acquired and the `gridd` service is restarted. This issue occurs if the `multimonitor` parameter is added to the `xorg.conf` file.

Version

Red Enterprise Linux 7.3

Status

Open

Ref. #

200275925

4.6. A segmentation fault in DBus code causes `nvidia-gridd` to exit on Red Hat Enterprise Linux and CentOS

Description

On Red Hat Enterprise Linux 6.8 and 6.9, and CentOS 6.8 and 6.9, a segmentation fault in DBus code causes the `nvidia-gridd` service to exit.

The `nvidia-gridd` service uses DBus for communication with **NVIDIA X Server Settings** to display licensing information through the **Manage License** page. Disabling the GUI for licensing resolves this issue.

To prevent this issue, the GUI for licensing is disabled by default. You might encounter this issue if you have enabled the GUI for licensing and are using Red Hat Enterprise Linux 6.8 or 6.9, or CentOS 6.8 and 6.9.

Version

Red Hat Enterprise Linux 6.8 and 6.9

CentOS 6.8 and 6.9

Status

Open

Ref. #

- ▶ 200358191
- ▶ 200319854
- ▶ 1895945

4.7. No Manage License option available in NVIDIA X Server Settings by default

Description

By default, the **Manage License** option is not available in **NVIDIA X Server Settings**. This option is missing because the GUI for licensing on Linux is disabled by default to work around the issue that is described in [A segmentation fault in Dbus code causes nvidia-gridd to exit on Red Hat Enterprise Linux and CentOS](#).

Workaround

This workaround requires `sudo` privileges.



Do not use this workaround 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.

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.

Status

Open

4.8. Licenses remain checked out when VMs are forcibly powered off

Description

NVIDIA vGPU software licenses remain checked out on the license server when non-persistent VMs are forcibly powered off.

The NVIDIA service running in a VM returns checked out licenses when the VM is shut down. In environments where non-persistent licensed VMs are not cleanly shut down, licenses on the license server can become exhausted. For example, this issue can occur in automated test environments where VMs are frequently changing and are not guaranteed to be cleanly shut down. The licenses from such VMs remain checked out against their MAC address for seven days before they time out and become available to other VMs.

Resolution

If VMs are routinely being powered off without clean shutdown in your environment, you can avoid this issue by shortening the license borrow period. To shorten the license borrow period, set the `LicenseInterval` configuration setting in your VM image. For details, refer to *Virtual GPU Client Licensing User Guide*.

Status

Closed

Ref.

1694975

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.