



# VIRTUAL GPU SOFTWARE R384 FOR RED HAT ENTERPRISE LINUX WITH KVM

RN-08687-001 \_v5.0 through 5.4 Revision 04 | May 2020

## Release Notes



# TABLE OF CONTENTS

<b>Chapter 1. Release Notes</b> .....	<b>1</b>
1.1. Updates in Release 5.0.....	1
1.2. Updates in Release 5.1.....	1
1.3. Updates in Release 5.2.....	2
1.4. Updates in Release 5.3.....	2
1.5. Updates in Release 5.4.....	2
<b>Chapter 2. Validated Platforms</b> .....	<b>3</b>
2.1. Supported NVIDIA GPUs and Validated Server Platforms.....	3
2.2. Hypervisor Software Releases.....	3
2.3. Guest OS Support.....	4
2.3.1. Windows Guest OS Support.....	4
2.3.2. Linux Guest OS Support.....	5
<b>Chapter 3. NVIDIA Software Security Updates</b> .....	<b>6</b>
<b>Chapter 4. Known Issues</b> .....	<b>7</b>
4.1. 5.1, 5.2 Only: Issues in remote desktop sessions if a license is acquired after a session is started.....	7
4.2. License settings configured through a GPO are ignored.....	8
4.3. Licensing pop-up windows contain the text microsoft.explorer.notification.....	8
4.4. 5.2 Only: The license expires prematurely in Linux guest VMs.....	9
4.5. Since 5.1: The link to NVIDIA Control Panel is missing after installation.....	9
4.6. Multiple display heads are not detected by Ubuntu 14.04 guest VMs.....	10
4.7. Since 5.1: On GPUs based on the Pascal architecture, Ubuntu 16.04 VMs run slowly after acquiring a license.....	11
4.8. Resolution is not updated after a VM acquires a license and is restarted.....	12
4.9. A segmentation fault in Dbus code causes nvidia-gridd to exit on Red Hat Enterprise Linux and CentOS.....	12
4.10. Since 5.1: No Manage License option available in NVIDIA X Server Settings by default...	13
4.11. Since 5.1: The nvidia-gridd service fails because the required configuration is not provided.....	14
4.12. Since 5.1: The Apply button is disabled after change to unlicensed mode.....	15
4.13. Licenses remain checked out when VMs are forcibly powered off.....	16

# 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 Red Hat Enterprise Linux with KVM.

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

Software	5.0	5.1	5.2	5.3	5.4
NVIDIA Windows driver	385.41	385.90	386.09	386.37	386.57
NVIDIA Linux driver version	384.73	384.99	384.111	384.137	384.155

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

## 1.1. Updates in Release 5.0

### New Features in Release 5.0

- ▶ Software enforcement of licensing requirements
- ▶ Miscellaneous bug fixes

### Hardware and Software Support Introduced in Release 5.0

- ▶ Support for GPUs based on the NVIDIA<sup>®</sup> Pascal<sup>™</sup> architecture

## 1.2. Updates in Release 5.1

### New Features in Release 5.1

- ▶ Miscellaneous bug fixes

## 1.3. Updates in Release 5.2

### New Features in Release 5.2

- ▶ New default values for the license borrow time and license linger time:
  - ▶ The default license borrow time is reduced from 7 days to 1 day.
  - ▶ The default license linger time is reduced from 10 minutes to 0 minutes.
- ▶ New setting `LingerInterval` for overriding the default license linger time
- ▶ Miscellaneous bug fixes

## 1.4. Updates in Release 5.3

### New Features in Release 5.3

- ▶ Plain-text logging on Windows of significant licensing events
- ▶ New setting `EnableLogging` for disabling or enabling logging of significant licensing events
- ▶ Miscellaneous bug fixes

## 1.5. Updates in Release 5.4

### New Features in Release 5.4

- ▶ Miscellaneous bug fixes

### Hardware and Software Support Introduced in Release 5.4

- ▶ Support for Red Hat Enterprise Linux 7.5 and CentOS 7.5 as a guest OS

# Chapter 2.

## VALIDATED PLATFORMS

This release family of NVIDIA vGPU software provides support for several NVIDIA GPUs on validated server hardware platforms, Red Hat Enterprise Linux with KVM 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 Red Hat Enterprise Linux with KVM, running on validated server hardware platforms:

- ▶ Tesla M6
- ▶ Tesla M10
- ▶ Tesla M60
- ▶ Tesla P4
- ▶ Tesla P6
- ▶ Tesla P40
- ▶ Tesla P100
- ▶ **Since 5.1:** Tesla P100 12GB

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 release listed in the table.



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

Software	Releases Supported	Notes
Red Hat Enterprise Linux with KVM	7.0, 7.1	Only the following NVIDIA GPUs are supported: <ul style="list-style-type: none"> <li>▶ Tesla M6</li> <li>▶ Tesla M10</li> <li>▶ Tesla M60</li> </ul>
Red Hat Enterprise Linux with KVM	7.2 through 7.4	All NVIDIA GPUs that support NVIDIA vGPU software are supported.

## 2.3. Guest OS Support

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



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 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



Red Hat Enterprise Linux with KVM supports Windows guest operating systems only under specific Red Hat subscription programs. For details, see [Certified guest operating systems for Red Hat Enterprise Linux with KVM](#).

NVIDIA vGPU software supports **only** the Windows releases listed as a guest OS on Red Hat Enterprise Linux with KVM.



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

- ▶ Windows Server 2012 R2
- ▶ Windows Server 2008 R2 on Tesla M6, Tesla M10, and Tesla M60 GPUs only
- ▶ Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703) (64-bit)
- ▶ Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703) (32-bit) on Tesla M6, Tesla M10, and Tesla M60 GPUs only
- ▶ Windows 8.1 Update (64-bit)
- ▶ Windows 8.1 Update (32-bit) on Tesla M6, Tesla M10, and Tesla M60 GPUs only
- ▶ Windows 7 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 Red Hat Enterprise Linux with KVM.



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

CentOS is not a certified guest OS for Red Hat Enterprise Linux with KVM.

- ▶ **Since 5.4:** Red Hat Enterprise Linux 7.5
- ▶ Red Hat Enterprise Linux 7.0-7.4
- ▶ **Since 5.4:** CentOS 7.5
- ▶ CentOS 7.0-7.4
- ▶ Red Hat Enterprise Linux 6.6
- ▶ CentOS 6.6

# Chapter 3.

## NVIDIA SOFTWARE SECURITY UPDATES

For more information about NVIDIA's vulnerability management, visit the [NVIDIA Product Security](#) page.

### NVIDIA Software Security Updates in Release 5.2

CVE ID	NVIDIA Issue Number	Description
CVE-2017-5753	CVE-2017-5753	Computer systems with microprocessors utilizing speculative execution and branch prediction may allow unauthorized disclosure of information to an attacker with local user access via a side-channel analysis.

### NVIDIA Software Security Updates in Release 5.3

No NVIDIA software security updates are reported in this release for Red Hat Enterprise Linux with KVM.

### NVIDIA Software Security Updates in Release 5.4

No NVIDIA software security updates are reported in this release for Red Hat Enterprise Linux with KVM.



# Chapter 4.

## KNOWN ISSUES

### 4.1. 5.1, 5.2 Only: Issues in remote desktop sessions if a license is acquired after a session is started

#### Description

A VM might acquire a license for NVIDIA vGPU software after a remote desktop session has connected to the VM. In this situation, some licensed features and capabilities are not available to a properly licensed vGPU or pass-through GPU in the session. For example, the updated maximum resolution supported is not available.

#### Workaround

Before attempting this workaround, confirm that the VM has obtained the correct license for NVIDIA vGPU software.

1. After installing the guest VM driver package and configuring required license settings on the VM (or on the master image used for VM deployment), set the `IgnoreSP` property to 1.

- ▶ On Windows, add the following registry setting:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\NVIDIA Corporation\Global\GridLicensing]
"IgnoreSP"=dword:00000001
```

- ▶ On Linux, add the following setting to the file `/etc/nvidia/gridd.conf`:

```
IgnoreSP=1
```

2. Restart the VM.

#### Status

Resolved in NVIDIA vGPU software release 5.3.

**Ref. #**

200391532

## 4.2. License settings configured through a GPO are ignored

**Description**

License settings configured through a Windows Group Policy Object (GPO) are ignored. Windows Registry settings applied through a GPO are set after the NVIDIA vGPU software graphics driver service is started. Therefore, NVIDIA vGPU software cannot be configured through a GPO.

**Workaround**

Use the Registry Editor to set the Windows Registry keys for license settings individually.

**Status**

Open

**Ref. #**

2010398

## 4.3. 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**

Open

**Ref. #**

200346607

## 4.4. 5.2 Only: The license expires prematurely in Linux guest VMs

**Description**

In Linux guest VMs, the license expires before the default borrow period has elapsed. In normal operation, the license is renewed periodically at an interval that depends on the license borrow period. As a result, a failure to renew the license may cause the license to expire before the default borrow period has elapsed.

**Workaround**

To reduce the possibility of license-renewal failures caused by transient network issues, increase the license borrow period to a value of about 7 days.

**Status**

Resolved in NVIDIA vGPU software release 5.3.

**Ref. #**

200376678

## 4.5. Since 5.1: The link to NVIDIA Control Panel is missing after installation

**Description**

NVIDIA Control Panel is missing from the desktop context menu on Windows VMs running on Red Hat Enterprise Linux Kernel-based Virtual Machine (KVM).

**Version**

All supported Windows guest OS versions on Red Hat Enterprise Linux with KVM (version 7.x)

**Workaround**

1. Log in to the host.
2. Get the serial number of the system.

```
# dmidecode -t system | grep "Serial Number"
```

3. In `virsh`, open for editing the XML file of your VM.

```
# virsh edit vm-name
```

*vm-name*

The name of your VM.

4. Inside the `os` element, add an `smbios` element for which `mode` is set to `sysinfo`.

```
<os>
  <smbios mode='sysinfo' />
</os>
```

5. Add an `smbios` system information entry element named `serial` that contains the serial number of the system.

```
<sysinfo type='smbios'>
  <system>
    <entry name='serial'>system-serial-number</entry>
  </system>
</sysinfo>
```

*system-serial-number*

The serial number of the system that you got by using the `dmidecode` command.

6. Start your VM.

```
# virsh start vm-name
```

*vm-name*

The name of your VM.

NVIDIA Control Panel should now be available in the desktop context menu on Windows.

### Status

Open

### Ref. #

200360518

## 4.6. Multiple display heads are not detected by Ubuntu 14.04 guest VMs

### Description

After an Ubuntu 14.04 guest VM has acquired a license, multiple display heads connected to the VM are not detected.

**Version**

Ubuntu 14.04

**Workaround**

To see all the connected display heads after the VM has acquired a license, open the **Displays** settings window and click **Detect displays**.

**Status**

Resolved in NVIDIA vGPU software release 5.3.

**Ref. #**

200334648

## 4.7. Since 5.1: On GPUs based on the Pascal architecture, Ubuntu 16.04 VMs run slowly after acquiring a license

**Description**

On GPUs based on the Pascal architecture, Ubuntu VMs to which an NVIDIA vGPU or pass-through GPU is assigned run slowly after acquiring a license. Ubuntu VMs that have not been assigned an NVIDIA vGPU or pass-through GPU run noticeably faster.

**Workaround**

After the VM has acquired a license, restart the `lightdm` service.

**Status**

Resolved in NVIDIA vGPU software release 5.2.

**Ref. #**

200359618

## 4.8. 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

## 4.9. 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.

**Since 5.1:** The GUI for licensing is disabled by default.

### Version

Red Hat Enterprise Linux 6.8 and 6.9

CentOS 6.8 and 6.9

NVIDIA vGPU software 5.0

## 5.0 Only: Workaround

This workaround requires `sudo` privileges.

1. As root, edit the `/etc/nvidia/gridd.conf` file to set the `EnableUI` option to `FALSE`.
2. Start the `nvidia-gridd` service.

```
# sudo service nvidia-gridd start
```

3. Confirm that the `nvidia-gridd` service has obtained a license by examining the log messages written to `/var/log/messages`.

```
# 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.
```

### Status

Open

### Ref. #

- ▶ 200358191
- ▶ 200319854
- ▶ 1895945

## 4.10. Since 5.1: 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](#).

### Version

NVIDIA vGPU software 5.1

### 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

## 4.11. Since 5.1: The `nvidia-gridd` service fails because the required configuration is not provided

### Description

The `nvidia-gridd` service exits with an error because the required configuration is not provided.

The known issue described in [A segmentation fault in Dbus code causes nvidia-gridd to exit on Red Hat Enterprise Linux and CentOS](#) causes the **NVIDIA X Server Settings** page for managing licensing settings through a GUI to be disabled by default. As a result, if the required license configuration is not provided through the configuration file, the service exits with an error.

Details of the error can be obtained by checking the status of the `nvidia-gridd` service.

```
# service nvidia-gridd status
nvidia-gridd.service - NVIDIA Grid Daemon
Loaded: loaded (/usr/lib/systemd/system/nvidia-gridd.service; enabled; vendor preset: disabled)
   Active: failed (Result: exit-code) since Wed 2017-11-01 19:25:07 IST; 27s ago
     Process: 1190 ExecStopPost=/bin/rm -rf /var/run/nvidia-gridd (code=exited, status=0/SUCCESS)
     Process: 11905 ExecStart=/usr/bin/nvidia-gridd (code=exited, status=0/SUCCESS)
  Main PID: 11906 (code=exited, status=1/FAILURE)
Nov 01 19:24:35 localhost.localdomain systemd[1]: Starting NVIDIA Grid Daemon...
Nov 01 19:24:35 localhost.localdomain nvidia-gridd[11906]: Started (11906)
Nov 01 19:24:35 localhost.localdomain systemd[1]: Started NVIDIA Grid Daemon.
```



```
Nov 01 19:24:36 localhost.localdomain nvidia-gridd[11906]: Failed to open
config file : /etc/nvidia/gridd.conf error :No such file or directory
Nov 01 19:25:07 localhost.localdomain nvidia-gridd[11906]: Service provider
detection complete.
Nov 01 19:25:07 localhost.localdomain nvidia-gridd[11906]: Shutdown (11906)
Nov 01 19:25:07 localhost.localdomain systemd[1]: nvidia-gridd.service: main
process exited, code=exited, status=1/FAILURE
Nov 01 19:25:07 localhost.localdomain systemd[1]: Unit nvidia-gridd.service
entered failed state.
Nov 01 19:25:07 localhost.localdomain systemd[1]: nvidia-gridd.service failed.
```

### Workaround

Use a configuration file to license NVIDIA vGPU software on Linux as explained in *Virtual GPU Client Licensing User Guide*.

### Status

Open

### Ref. #

200359469

## 4.12. Since 5.1: The Apply button is disabled after change to unlicensed mode

### Description

After the mode is changed from licensed **Quadro Virtual Datacenter Workstation Edition** mode to **Unlicensed Tesla** mode, the **Apply** button on the **Manage GRID License** page is disabled. As a result, **NVIDIA X Server Settings** cannot be used to switch to **Tesla (Unlicensed)** mode on a licensed system.

### Workaround

1. Start **NVIDIA X Server Settings** by using the method for launching applications provided by your Linux distribution.
2. In the **NVIDIA X Server Settings** window that opens, click **Manage GRID License**.
3. Clear the **Primary Server** field.
4. Select the **Tesla (unlicensed)** option.
5. Click **Apply**.

### Status

Resolved in NVIDIA vGPU software release 5.2.

**Ref. #**

200359624

## 4.13. 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.