



VIRTUAL GPU SOFTWARE R384 FOR NUTANIX AHV

RN-08659-001 _v5.0 through 5.4 Revision 04 | May 2020

Release Notes



TABLE OF CONTENTS

Chapter 1. Release Notes	1
1.1. Updates in Release 5.1.....	1
1.2. Updates in Release 5.2.....	1
1.3. Updates in Release 5.3.....	2
1.4. Updates in Release 5.4.....	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
2.4. NVIDIA CUDA Toolkit Version Support.....	5
Chapter 3. Known Product Limitations	7
3.1. vGPU profiles with 512 Mbytes or less of frame buffer support only 1 virtual display head on Windows 10.....	7
3.2. NVENC requires at least 1 Gbyte of frame buffer.....	8
3.3. VM running older NVIDIA vGPU drivers fails to initialize vGPU when booted.....	8
3.4. Virtual GPU fails to start if ECC is enabled.....	9
3.5. Single vGPU benchmark scores are lower than pass-through GPU.....	11
3.6. nvidia-smi fails to operate when all GPUs are assigned to GPU passthrough mode.....	11
3.7. VMs configured with large memory fail to initialize vGPU when booted.....	12
Chapter 4. Resolved Issues	13
Chapter 5. NVIDIA Software Security Updates	15
Chapter 6. Known Issues	16
6.1. Cloned VMs configured with a vGPU type different than the type in the master image fail to start.....	16
6.2. 5.1, 5.2 Only: Issues in remote desktop sessions if a license is acquired after a session is started.....	17
6.3. License settings configured through a GPO are ignored.....	17
6.4. Licensing pop-up windows contain the text microsoft.explorer.notification.....	18
6.5. 5.2 Only: The license expires prematurely in Linux guest VMs.....	18
6.6. Multiple display heads are not detected by Ubuntu 14.04 guest VMs.....	19
6.7. Since 5.1: On GPUs based on the Pascal architecture, Ubuntu 16.04 VMs run slowly after acquiring a license.....	20
6.8. Resolution is not updated after a VM acquires a license and is restarted.....	20
6.9. A segmentation fault in Dbus code causes nvidia-gridd to exit on Red Hat Enterprise Linux and CentOS.....	21
6.10. Since 5.1: No Manage License option available in NVIDIA X Server Settings by default...	22
6.11. Since 5.1: The nvidia-gridd service fails because the required configuration is not provided.....	23

6.12. Since 5.1: The Apply button is disabled after change to unlicensed mode.....	24
6.13. Licenses remain checked out when VMs are forcibly powered off.....	24
6.14. Memory exhaustion can occur with vGPU profiles that have 512 Mbytes or less of frame buffer.....	25
6.15. GNOME Display Manager (GDM) fails to start on Red Hat Enterprise Linux 7.2 and CentOS 7.0.....	26

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 Nutanix AHV.

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 Virtual GPU Manager for the Nutanix AHV releases listed in Hypervisor Software Releases	Not supported	384.99	384.111	384.137	384.155
NVIDIA Windows driver	Not supported	385.90	386.09	386.37	386.57
NVIDIA Linux driver version	Not supported	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 5.1

New Features in Release 5.1

- ▶ Miscellaneous bug fixes

1.2. 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

1.3. 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

1.4. 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, Nutanix AHV hypervisor software versions, and guest operating systems. It also supports the version of NVIDIA CUDA Toolkit that is compatible with R384 drivers.

2.1. Supported NVIDIA GPUs and Validated Server Platforms

This release of NVIDIA vGPU software provides support for the following NVIDIA GPUs on Nutanix AHV, running on validated Nutanix NX series server hardware platforms:

- ▶ Tesla M10
- ▶ Tesla M60
- ▶ Tesla P40

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



Tesla M60 and M6 GPUs support compute mode and graphics mode. NVIDIA vGPU requires GPUs that support both modes to operate in graphics mode.

Recent Tesla M60 GPUs and M6 GPUs are supplied in graphics mode. However, your GPU might be in compute mode if it is an older Tesla M60 GPU or M6 GPU, or if its mode has previously been changed.

To configure the mode of Tesla M60 and M6 GPUs, use the `gpumodeswitch` tool provided with NVIDIA vGPU software releases.

2.2. Hypervisor Software Releases

This release is supported on the Nutanix AHV releases listed in the table.



Updates to a base release of Nutanix AHV are compatible with the base release and can also be used with this version of NVIDIA vGPU software unless expressly stated otherwise.

Software	Releases Supported
Nutanix AOS Hypervisor (AHV) 5.5	5.5 RTM and compatible 5.5.x updates

2.3. Guest OS Support

NVIDIA vGPU software supports several Windows releases and Linux distributions as a guest OS. The supported guest operating systems depend on the hypervisor software version.



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

NVIDIA vGPU software supports **only** the Windows releases listed in the table as a guest OS on Nutanix AHV. The releases of Nutanix AHV for which a Windows release is supported depend on whether NVIDIA vGPU or pass-through GPU is used.



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

Guest OS	NVIDIA vGPU - Nutanix AHV Releases	Pass-Through GPU - Nutanix AHV Releases
Windows Server 2016 1607, 1709	5.5	5.5
Windows Server 2012 R2	5.5	5.5
Windows Server 2008 R2	5.5	5.5

Guest OS	NVIDIA vGPU - Nutanix AHV Releases	Pass-Through GPU - Nutanix AHV Releases
Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703) (64-bit)	5.5	5.5
Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703) (32-bit)	5.5	5.5
Windows 8.1 Update (64-bit)	5.5	5.5
Windows 8.1 Update (32-bit)	5.5	5.5
Windows 8.1 (64-bit)	5.5	5.5
Windows 8.1 (32-bit)	5.5	5.5
Windows 8 (32/64-bit)	5.5	5.5
Windows 7 (32/64-bit)	5.5	5.5

2.3.2. Linux Guest OS Support

NVIDIA vGPU software supports **only** the Linux distributions listed in the table as a guest OS on Nutanix AHV. The releases of Nutanix AHV for which a Linux release is supported depend on whether NVIDIA vGPU or pass-through GPU is used.



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

Guest OS	NVIDIA vGPU - Nutanix AHV Releases	Pass-Through GPU - Nutanix AHV Releases
Since 5.4: Red Hat Enterprise Linux 7.5	5.5	5.5
Red Hat Enterprise Linux 7.0-7.4	5.5	5.5
Since 5.4: CentOS 7.5	5.5	5.5
CentOS 7.0-7.4	5.5	5.5
Ubuntu 16.04 LTS	5.5	5.5
Ubuntu 14.04 LTS	5.5	5.5

2.4. NVIDIA CUDA Toolkit Version Support

The releases in this release family of NVIDIA vGPU software support NVIDIA CUDA Toolkit 9.0.

For more information about NVIDIA CUDA Toolkit, see [CUDA Toolkit 9.0 Documentation](#).

Chapter 3.

KNOWN PRODUCT LIMITATIONS

Known product limitations for this release of NVIDIA vGPU software are described in the following sections.

3.1. vGPU profiles with 512 Mbytes or less of frame buffer support only 1 virtual display head on Windows 10

Description

To reduce the possibility of memory exhaustion, vGPU profiles with 512 Mbytes or less of frame buffer support only 1 virtual display head on a Windows 10 guest OS.

The following vGPU profiles have 512 Mbytes or less of frame buffer:

- ▶ Tesla M6-0B, M6-0Q
- ▶ Tesla M10-0B, M10-0Q
- ▶ Tesla M60-0B, M60-0Q

Workaround

Use a profile that supports more than 1 virtual display head and has at least 1 Gbyte of frame buffer.

3.2. NVENC requires at least 1 Gbyte of frame buffer

Description

Using the frame buffer for the NVIDIA hardware-based H.264/HEVC video encoder (NVENC) may cause memory exhaustion with vGPU profiles that have 512 Mbytes or less of frame buffer. To reduce the possibility of memory exhaustion, NVENC is disabled on profiles that have 512 Mbytes or less of frame buffer. Application GPU acceleration remains fully supported and available for all profiles, including profiles with 512 Mbytes or less of frame buffer. NVENC support from both Citrix and VMware is a recent feature and, if you are using an older version, you should experience no change in functionality.

The following vGPU profiles have 512 Mbytes or less of frame buffer:

- ▶ Tesla M6-0B, M6-0Q
- ▶ Tesla M10-0B, M10-0Q
- ▶ Tesla M60-0B, M60-0Q

Workaround

If you require NVENC to be enabled, use a profile that has at least 1 Gbyte of frame buffer.

3.3. VM running older NVIDIA vGPU drivers fails to initialize vGPU when booted

Description

A VM running a version of the NVIDIA guest VM drivers from a previous main release branch, for example release 4.4, will fail to initialize vGPU when booted on a Nutanix AHV platform running the current release of Virtual GPU Manager.

In this scenario, the VM boots in standard VGA mode with reduced resolution and color depth. The NVIDIA virtual GPU is present in **Windows Device Manager** but displays a warning sign, and the following device status:

```
Windows has stopped this device because it has reported problems. (Code 43)
```

Depending on the versions of drivers in use, the Nutanix AHV VM's `/var/log/messages` log file reports one of the following errors:

- ▶ An error message:

```
vmiop_log: error: Unable to fetch Guest NVIDIA driver information
```

- ▶ A version mismatch between guest and host drivers:

```
vmiop_log: error: Guest VGX version(1.1) and Host VGX version(1.2) do not match
```

- ▶ A signature mismatch:

```
vmiop_log: error: VGPU message signature mismatch.
```

Resolution

Install the current NVIDIA guest VM driver in the VM.

3.4. Virtual GPU fails to start if ECC is enabled

Description

Tesla M60, Tesla M6, and GPUs based on the Pascal GPU architecture, for example Tesla P100 or Tesla P4, support error correcting code (ECC) memory for improved data integrity. Tesla M60 and M6 GPUs in graphics mode are supplied with ECC memory disabled by default, but it may subsequently be enabled using `nvidia-smi`. GPUs based on the Pascal GPU architecture are supplied with ECC memory enabled.

However, NVIDIA vGPU does not support ECC memory. If ECC memory is enabled, NVIDIA vGPU fails to start.

The following error is logged in the Nutanix AHV VM's `/var/log/messages` log file:

```
vmiop_log: error: Initialization: VGX not supported with ECC Enabled.
```

Resolution

Ensure that ECC is disabled on all GPUs.

Before you begin, ensure that NVIDIA Virtual GPU Manager is installed on your hypervisor.

1. Use `nvidia-smi` to list the status of all GPUs, and check for ECC noted as enabled on GPUs.

```
# nvidia-smi -q
=====NVSMI LOG=====
Timestamp                : Tue Dec 19 18:36:45 2017
Driver Version           : 384.99
Attached GPUs            : 1
GPU 0000:02:00.0
[...]
```

```

Ecc Mode
  Current      : Enabled
  Pending      : Enabled

```

[...]

2. Change the ECC status to off on each GPU for which ECC is enabled.
 - ▶ If you want to change the ECC status to off for all GPUs on your host machine, run this command:

```
# nvidia-smi -e 0
```

- ▶ If you want to change the ECC status to off for a specific GPU, run this command:

```
# nvidia-smi -i id -e 0
```

id is the index of the GPU as reported by `nvidia-smi`.

This example disables ECC for the GPU with index `0000:02:00.0`.

```
# nvidia-smi -i 0000:02:00.0 -e 0
```

3. Reboot the host.
4. Confirm that ECC is now disabled for the GPU.

```

# nvidia-smi -q

=====NVSMI LOG=====

Timestamp           : Tue Dec 19 18:37:53 2017
Driver Version      : 384.99

Attached GPUs       : 1
GPU 0000:02:00.0
[...]

Ecc Mode
  Current      : Disabled
  Pending      : Disabled

[...]

```

If you later need to enable ECC on your GPUs, run one of the following commands:

- ▶ If you want to change the ECC status to on for all GPUs on your host machine, run this command:

```
# nvidia-smi -e 1
```

- ▶ If you want to change the ECC status to on for a specific GPU, run this command:

```
# nvidia-smi -i id -e 1
```

id is the index of the GPU as reported by `nvidia-smi`.

This example enables ECC for the GPU with index `0000:02:00.0`.

```
# nvidia-smi -i 0000:02:00.0 -e 1
```

After changing the ECC status to on, reboot the host.

3.5. Single vGPU benchmark scores are lower than pass-through GPU

Description

A single vGPU configured on a physical GPU produces lower benchmark scores than the physical GPU run in pass-through mode.

Aside from performance differences that may be attributed to a vGPU's smaller frame buffer size, vGPU incorporates a performance balancing feature known as Frame Rate Limiter (FRL). On vGPUs that use the best-effort scheduler, FRL is enabled. On vGPUs that use the fixed share or equal share scheduler, FRL is disabled.

FRL is used to ensure balanced performance across multiple vGPUs that are resident on the same physical GPU. The FRL setting is designed to give good interactive remote graphics experience but may reduce scores in benchmarks that depend on measuring frame rendering rates, as compared to the same benchmarks running on a pass-through GPU.

3.6. `nvidia-smi` fails to operate when all GPUs are assigned to GPU passthrough mode

Description

If all GPUs in the platform are assigned to VMs in passthrough mode, `nvidia-smi` will return an error:

```
[root@vngx-test ~]# nvidia-smi  
Failed to initialize NVML: Unknown Error
```

This is because GPUs operating in passthrough mode are not visible to `nvidia-smi` and the NVIDIA kernel driver operating in the Nutanix AHV dom0.

Resolution

N/A

3.7. VMs configured with large memory fail to initialize vGPU when booted

Description

When starting multiple VMs configured with large amounts of RAM (typically more than 32GB per VM), a VM may fail to initialize vGPU. In this scenario, the VM boots in standard VGA mode with reduced resolution and color depth. The NVIDIA vGPU software GPU is present in **Windows Device Manager** but displays a warning sign, and the following device status:

```
Windows has stopped this device because it has reported problems. (Code 43)
```

The Nutanix AHV VM's `/var/log/messages` log file contains these error messages:

```
vmiop_log: error: NVOS status 0x29
vmiop_log: error: Assertion Failed at 0x7620fd4b:179
vmiop_log: error: 8 frames returned by backtrace
...
vmiop_log: error: VGPU message 12 failed, result code: 0x29
...
vmiop_log: error: NVOS status 0x8
vmiop_log: error: Assertion Failed at 0x7620c8df:280
vmiop_log: error: 8 frames returned by backtrace
...
vmiop_log: error: VGPU message 26 failed, result code: 0x8
```


Chapter 4.

RESOLVED ISSUES

Only resolved issues that have been previously noted as known issues or had a noticeable user impact are listed. The summary and description for each resolved issue indicate the effect of the issue on NVIDIA vGPU software **before the issue was resolved**.

Issues Resolved in Release 5.1

No resolved issues are reported in this release for Nutanix AHV.

Issues Resolved in Release 5.2

Bug ID	Summary and Description
200359618	<p>On GPUs based on the Pascal architecture, Ubuntu 16.04 VMs run slowly after acquiring a license</p> <p>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.</p>
200359624	<p>The Apply button is disabled after change to unlicensed mode</p> <p>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.</p>

Issues Resolved in Release 5.3

Bug ID	Summary and Description
200334648	<p>Multiple display heads are not detected by Ubuntu 14.04 guest VMs</p>

Bug ID	Summary and Description
	After an Ubuntu 14.04 guest VM has acquired a license, multiple display heads connected to the VM are not detected.
2075467	<p>The displays flicker each time a license is requested or renewed in Linux guest VMs</p> <p>Whenever a license is requested or renewed in Linux guest VMs, the displays are reconfigured and rescanned. Rescanning the displays causes the remoting solution to momentarily drop the connection and, as a result, the displays flicker.</p>
200376678	<p>The license expires prematurely in Linux guest VMs</p> <p>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.</p>
200391532	<p>Issues in remote desktop sessions if a license is acquired after a session is started</p> <p>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.</p>

Issues Resolved in Release 5.4

No resolved issues are reported in this release for Nutanix AHV.

Chapter 5.

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 Nutanix AHV.

NVIDIA Software Security Updates in Release 5.4

No NVIDIA software security updates are reported in this release for Nutanix AHV.

Chapter 6.

KNOWN ISSUES

6.1. Cloned VMs configured with a vGPU type different than the type in the master image fail to start

Description

Cloned VMs configured with a vGPU type different than the type in the master image fail to start.

When a Windows 10 VM is booted, the VM becomes stuck in a loop and alternately displays `Getting devices ready: 50%` and `Preparation in progress`.

Workaround

Create one master image for each vGPU type that you want to use. Do not attempt to configure a cloned VM with a vGPU type different than the type in the master image.

Status

Not an NVIDIA bug

Ref.

2285306

6.2. 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

6.3. 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

6.4. 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

6.5. 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

6.6. 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

6.7. 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

6.8. 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

6.9. 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

6.10. 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

6.11. 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: 11990 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

6.12. 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

6.13. 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

6.14. Memory exhaustion can occur with vGPU profiles that have 512 Mbytes or less of frame buffer

Description

Memory exhaustion can occur with vGPU profiles that have 512 Mbytes or less of frame buffer.

This issue typically occurs in the following situations:

- ▶ Full screen 1080p video content is playing in a browser. In this situation, the session hangs and session reconnection fails.
- ▶ Multiple display heads are used with Citrix XenDesktop or VMware Horizon on a Windows 10 guest VM.
- ▶ Higher resolution monitors are used.
- ▶ Applications that are frame-buffer intensive are used.
- ▶ NVENC is in use.

To reduce the possibility of memory exhaustion, NVENC is disabled on profiles that have 512 Mbytes or less of frame buffer.

The following vGPU profiles have 512 Mbytes or less of frame buffer:

- ▶ Tesla M6-0B, M6-0Q
- ▶ Tesla M10-0B, M10-0Q

- ▶ Tesla M60-0B, M60-0Q

The root cause is a known issue associated with changes to the way that recent Microsoft operating systems handle and allow access to overprovisioning messages and errors. If your systems are provisioned with enough frame buffer to support your use cases, you should not encounter these issues.

Workaround

- ▶ Use an appropriately sized vGPU to ensure that the frame buffer supplied to a VM through the vGPU is adequate for your workloads.
- ▶ Monitor your frame buffer usage.
- ▶ If you are using Windows 10, consider these workarounds and solutions:
 - ▶ Use a profile that has 1 Gbyte of frame buffer.
 - ▶ Optimize your Windows 10 resource usage.

To obtain information about best practices for improved user experience using Windows 10 in virtual environments, complete the [NVIDIA GRID vGPU Profile Sizing Guide for Windows 10](#) download request form.

Status

Open

Ref.

- ▶ 200130864
- ▶ 1803861

6.15. GNOME Display Manager (GDM) fails to start on Red Hat Enterprise Linux 7.2 and CentOS 7.0

Description

GDM fails to start on Red Hat Enterprise Linux 7.2 and CentOS 7.0 with the following error:

```
Oh no! Something has gone wrong!
```

Workaround

Permanently enable permissive mode for Security Enhanced Linux (SELinux).

1. As root, edit the `/etc/selinux/config` file to set SELINUX to permissive.


```
SELINUX=permissive
```
2. Reboot the system.

```
~]# reboot
```

For more information, see [Permissive Mode](#) in *Red Hat Enterprise Linux 7 SELinux User's and Administrator's Guide*.

Status

Not an NVIDIA bug

Ref. #

200167868

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.