



# VIRTUAL GPU SOFTWARE R384 FOR CITRIX XENSERVER

RN-06927-001 \_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.....	2
1.2. Updates in Release 5.1.....	2
1.3. Updates in Release 5.2.....	2
1.4. Updates in Release 5.3.....	3
1.5. Updates in Release 5.4.....	3
<b>Chapter 2. Validated Platforms</b> .....	<b>4</b>
2.1. Supported NVIDIA GPUs and Validated Server Platforms.....	4
2.2. Hypervisor Software Releases.....	5
2.3. Guest OS Support.....	6
2.3.1. Windows Guest OS Support.....	6
2.3.2. Linux Guest OS Support.....	8
2.4. NVIDIA CUDA Toolkit Version Support.....	8
<b>Chapter 3. Known Product Limitations</b> .....	<b>9</b>
3.1. vGPU profiles with 512 Mbytes or less of frame buffer support only 1 virtual display head on Windows 10.....	9
3.2. NVENC requires at least 1 Gbyte of frame buffer.....	10
3.3. VM running older NVIDIA vGPU drivers fails to initialize vGPU when booted.....	10
3.4. Virtual GPU fails to start if ECC is enabled.....	11
3.5. Single vGPU benchmark scores are lower than pass-through GPU.....	13
3.6. nvidia-smi fails to operate when all GPUs are assigned to GPU passthrough mode.....	14
3.7. Windows Aero is disabled on XenDesktop session using 3 or 4 monitors in 2560×1600 resolution.....	15
3.8. VMs configured with large memory fail to initialize vGPU when booted.....	15
3.9. vGPU host driver RPM upgrade fails.....	16
<b>Chapter 4. Resolved Issues</b> .....	<b>17</b>
<b>Chapter 5. NVIDIA Software Security Updates</b> .....	<b>20</b>
<b>Chapter 6. Known Issues</b> .....	<b>21</b>
6.1. Cloned VMs configured with a vGPU type different than the type in the master image fail to start.....	21
6.2. NVIDIA Control Panel Crashes in a VM connected to two 4K displays.....	22
6.3. vGPU guest VM driver not properly loaded on servers with more than 512 GB or 1 TB or more of system memory.....	22
6.4. 5.1, 5.2 Only: Issues in remote desktop sessions if a license is acquired after a session is started.....	23
6.5. License settings configured through a GPO are ignored.....	24
6.6. Licensing pop-up windows contain the text microsoft.explorer.notification.....	24
6.7. 5.2 Only: The license expires prematurely in Linux guest VMs.....	25
6.8. Multiple display heads are not detected by Ubuntu 14.04 guest VMs.....	25
6.9. CILP is not working in Windows VMs.....	26

6.10. Luxmark causes a segmentation fault on an unlicensed Linux client.....	26
6.11. Ubuntu 16.04 VMs run slowly after acquiring a license.....	27
6.12. Since 5.1: On GPUs based on the Pascal architecture, Ubuntu 16.04 VMs run slowly after acquiring a license.....	28
6.13. Resolution is not updated after a VM acquires a license and is restarted.....	28
6.14. On Tesla P40, P6, and P4 GPUs, the default ECC setting prevents NVIDIA vGPU from starting.....	29
6.15. NVIDIA vGPU utilization in guest VMs is incorrectly reported for vGPUs using a fixed share scheduler.....	29
6.16. NVIDIA vGPU encoder and process utilization counters don't work with Windows Performance Counters.....	30
6.17. A segmentation fault in DBus code causes nvidia-gridd to exit on Red Hat Enterprise Linux and CentOS.....	31
6.18. Since 5.1: No Manage License option available in NVIDIA X Server Settings by default....	32
6.19. Since 5.1: The nvidia-gridd service fails because the required configuration is not provided.....	33
6.20. Since 5.1: The Apply button is disabled after change to unlicensed mode.....	34
6.21. Licenses remain checked out when VMs are forcibly powered off.....	34
6.22. Memory exhaustion can occur with vGPU profiles that have 512 Mbytes or less of frame buffer.....	35
6.23. VM bug checks after the guest VM driver for Windows 10 RS2 is installed.....	36
6.24. On XenServer 7.0, VMs unexpectedly reboot and XenServer crashes or freezes.....	37
6.25. With no NVIDIA driver installed, XenServer misidentifies Tesla M10 cards.....	38
6.26. GNOME Display Manager (GDM) fails to start on Red Hat Enterprise Linux 7.2 and CentOS 7.0.....	38
6.27. Video goes blank when run in loop in Windows Media Player.....	39
6.28. Local VGA console is momentarily unblanked when XenDesktop changes resolution of the VM desktop.....	40
6.29. VM bugchecks on shutdown/restart when XenDesktop is installed and NVIDIA driver is uninstalled or upgraded.....	40
6.30. Application frame rate may drop when running XenDesktop at 2560×1600 resolution.....	41
6.31. Windows VM BSOD.....	41
6.32. Windows VM BSOD when upgrading NVIDIA drivers over a XenDesktop session.....	42
6.33. XenCenter does not allow vGPUs to be selected as a GPU type for Linux VMs.....	43
6.34. If X server is killed on a RHEL7 VM running vGPU, XenCenter console may not automatically switch to text console.....	43
6.35. XenDesktop shows only a black screen when connected to a vGPU VM.....	44



# 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 Citrix XenServer.

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 Citrix XenServer releases listed in <a href="#">Hypervisor Software Releases</a>	384.73	384.99	384.111	384.137	384.155
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



### Caution

If you install the wrong package for the version of Citrix XenServer you are using, NVIDIA Virtual GPU Manager will fail to load.

The releases of the vGPU Manager and guest VM drivers that you install must be compatible. Different versions of the vGPU Manager and guest VM driver from within the same main release branch can be used together. For example, you can use the vGPU Manager from release 5.1 with guest VM drivers from release 5.0. However, versions of the vGPU Manager and guest VM driver from different main release branches cannot be used together. For example, you cannot use the vGPU Manager from release 5.1 with guest VM drivers from release 4.4. See [VM running older NVIDIA vGPU drivers fails to initialize vGPU when booted](#).

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

- ▶ New NVIDIA vGPU schedulers for GPUs based on the NVIDIA Pascal architecture
- ▶ Support for NVML and `nvidia-smi` on 32-bit Windows VMs
- ▶ Application-level monitoring of NVIDIA vGPU engine utilization
- ▶ Encoder session monitoring
- ▶ Support for NVENC on Linux NVIDIA vGPUs
- ▶ 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
- ▶ Compute mode support for NVIDIA vGPU and pass-through GPU on GPUs based on the NVIDIA Pascal architecture

### Feature Support Withdrawn in Release 5.0

- ▶ GRID K1 and GRID K2 GPUs are no longer supported.

## 1.2. Updates in Release 5.1

### New Features in Release 5.1

- ▶ Miscellaneous bug fixes

### Hardware and Software Support Introduced in Release 5.1

- ▶ Support for Citrix XenDesktop 7.15

## 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 Citrix XenDesktop 7.18
- ▶ Support for Red Hat Enterprise Linux 7.5 and CentOS 7.5 as a guest OS

### Feature Support Withdrawn in Release 5.4

- ▶ Citrix XenServer 7.3, 7.2, and 6.5 are no longer supported.
- ▶ Citrix XenDesktop 7.6, 7.8, and 7.11 are no longer supported.

# Chapter 2.

## VALIDATED PLATFORMS

This release family of NVIDIA vGPU software provides support for several NVIDIA GPUs on validated server hardware platforms, Citrix XenServer 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 Citrix XenServer, 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](#).



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

### Supported Citrix XenServer Releases

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

 Support for NVIDIA vGPU software requires the Premium Edition (previously Enterprise Edition) of Citrix XenServer. For details, see [Licensing](#) in the Citrix documentation.

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

Software	Release Supported	Notes
<b>5.1-5.3 only:</b> Citrix XenServer 7.3	RTM build is supported.	All NVIDIA GPUs that support NVIDIA vGPU software are supported.
<b>5.0-5.3 only:</b> Citrix XenServer 7.2	RTM build is supported.	All NVIDIA GPUs that support NVIDIA vGPU software are supported.
Citrix XenServer 7.1	RTM build is supported.	All NVIDIA GPUs that support NVIDIA vGPU software are supported.
Citrix XenServer 7.0	RTM build 125380 is supported.	All NVIDIA GPUs that support NVIDIA vGPU software are supported.
<b>5.0-5.3 only:</b> Citrix XenServer 6.5	Version 6.5 with XS65ESP1. The NVIDIA Virtual GPU Manager included in this release will not install without XenServer 6.5 SP1.	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>

### Supported Virtual Desktop Software Releases

This release supports **only** the virtual desktop software releases listed in the table.

Software	Releases Supported
Citrix XenDesktop	<p><b>Since 5.4:</b> Version 7.18 in HDX 3D Pro mode.</p> <p><b>Since 5.1:</b> Version 7.15 in HDX 3D Pro mode.</p>

Software	Releases Supported
	<b>5.0-5.3 only:</b> Version 7.6, 7.8, and 7.11 in HDX 3D Pro mode.

## 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 Citrix XenServer. The releases of Citrix XenServer 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 - Citrix XenServer Releases	Pass-Through GPU - Citrix XenServer Releases
Windows Server 2016 1607, 1709	<b>5.1-5.3 only:</b> 7.3 <b>5.0-5.3 only:</b> 7.2 7.1, 7.0	<b>5.1-5.3 only:</b> 7.3 <b>5.0-5.3 only:</b> 7.2 7.1, 7.0
Windows Server 2012 R2	<b>5.1-5.3 only:</b> 7.3 <b>5.0-5.3 only:</b> 7.2 7.1, 7.0 <b>5.0-5.3 only:</b> 6.5 SP1	<b>5.1-5.3 only:</b> 7.3 <b>5.0-5.3 only:</b> 7.2 7.1, 7.0 <b>5.0-5.3 only:</b> 6.5 SP1
Windows Server 2008 R2	<b>5.1-5.3 only:</b> 7.3 <b>5.0-5.3 only:</b> 7.2 7.1, 7.0 <b>5.0-5.3 only:</b> 6.5 SP1	<b>5.1-5.3 only:</b> 7.3 <b>5.0-5.3 only:</b> 7.2 7.1, 7.0 <b>5.0-5.3 only:</b> 6.5 SP1

Guest OS	NVIDIA vGPU - Citrix XenServer Releases	Pass-Through GPU - Citrix XenServer Releases
Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703) (64-bit)	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 <b>5.0-5.3 only: 6.5 SP1</b>	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 <b>5.0-5.3 only: 6.5 SP1</b>
Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703) (32-bit)	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 <b>5.0-5.3 only: 6.5 SP1</b>	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 <b>5.0-5.3 only: 6.5 SP1</b>
Windows 8.1 Update (64-bit)	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 <b>5.0-5.3 only: 6.5 SP1</b>	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 <b>5.0-5.3 only: 6.5 SP1</b>
Windows 8.1 Update (32-bit)	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 6.5 SP1	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 6.5 SP1
Windows 8.1 (64-bit)	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 <b>5.0-5.3 only: 6.5 SP1</b>	-
Windows 8.1 (32-bit)	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 <b>5.0-5.3 only: 6.5 SP1</b>	-
Windows 8 (32/64-bit)	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 <b>5.0-5.3 only: 6.5 SP1</b>	-
Windows 7 (32/64-bit)	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 <b>5.0-5.3 only: 6.5 SP1</b>	<b>5.1-5.3 only: 7.3</b> <b>5.0-5.3 only: 7.2</b> 7.1, 7.0 <b>5.0-5.3 only: 6.5 SP1</b> Supported only on Tesla M6, Tesla M10, and Tesla M60 GPUs

## 2.3.2. Linux Guest OS Support

NVIDIA vGPU software supports **only** the Linux distributions listed in the table as a guest OS on Citrix XenServer. The releases of Citrix XenServer 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 - Citrix XenServer Releases	Pass-Through GPU - Citrix XenServer Releases
Since 5.4: Red Hat Enterprise Linux 7.5	7.1, 7.0	7.1, 7.0
Red Hat Enterprise Linux 7.0-7.4	5.1-5.3 only: 7.3 5.0-5.3 only: 7.2 7.1, 7.0	5.1-5.3 only: 7.3 5.0-5.3 only: 7.2 7.1, 7.0
Since 5.4: CentOS 7.5	7.1, 7.0	7.1, 7.0
CentOS 7.0-7.4	5.1-5.3 only: 7.3 5.0-5.3 only: 7.2 7.1, 7.0	5.1-5.3 only: 7.3 5.0-5.3 only: 7.2 7.1, 7.0
Ubuntu 16.04 LTS	5.1-5.3 only: 7.3 5.0-5.3 only: 7.2 7.1, 7.0	5.1-5.3 only: 7.3 5.0-5.3 only: 7.2 7.1, 7.0
Ubuntu 14.04 LTS	5.1-5.3 only: 7.3 5.0-5.3 only: 7.2 7.1, 7.0	5.1-5.3 only: 7.3 5.0-5.3 only: 7.2 7.1, 7.0

## 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 Citrix XenServer 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 Citrix XenServer 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.

Citrix XenCenter displays the following error message:

```
An emulator required to run this VM failed to start
```

The following error is logged in the Citrix XenServer 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.

```
# shutdown -r now
```

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.

### Resolution

FRL is controlled by an internal vGPU setting. On vGPUs that use the best-effort scheduler, NVIDIA does not validate vGPU with FRL disabled, but for validation of benchmark performance, FRL can be temporarily disabled by specifying `frame_rate_limiter=0` in the VM's `platform:vgpu_extra_args` parameter:

```
[root@xenserver ~]# xe vm-param-set uuid=e71afda4-53f4-3a1b-6c92-a364a7f619c2 platform:vgpu_extra_args="frame_rate_limiter=0"
[root@xenserver ~]#
```

The setting takes effect the next time the VM is started or rebooted.

With this setting in place, the VM's vGPU will run without any frame rate limit. The FRL can be reverted back to its default setting in one of the following ways:

- ▶ Removing the `vgpu_extra_args` key from the `platform` parameter
- ▶ Removing `frame_rate_limiter=0` from the `vgpu_extra_args` key
- ▶ Setting `frame_rate_limiter=1`. For example:

```
[root@xenserver ~]# xe vm-param-set uuid=e71afda4-53f4-3a1b-6c92-a364a7f619c2 platform:vgpu_extra_args="frame_rate_limiter=1"
[root@xenserver ~]#
```

## 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@xenserver-vgx-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 Citrix XenServer dom0.

To confirm that all GPUs are operating in passthrough, use XenCenter's GPU tab to review current GPU assignment:

The screenshot shows the XenCenter interface for a VM named 'xenserver-vgx-test'. The 'GPU' tab is selected, displaying the GPU assignment configuration. The placement policy is set to 'Maximum density: put as many VMs as possible on the same GPU'. The GPU model is 'GK104GL [GRID K2]'. The allowed vGPU types are listed as follows:

- Pass-through whole GPU
- GRID K260Q vGPU (2 per GPU)
- GRID K240Q vGPU (4 per GPU)
- GRID K220Q vGPU (8 per GPU)
- GRID K200 vGPU (8 per GPU)

The current assignment shows six VMs assigned to the GPU in passthrough mode:

- Test VM passthrough
- Test VM passthrough
- vgx-base-image-win7-64 passthrough
- Test VM passthrough
- Test VM passthrough
- Test VM passthrough

Buttons for 'Select All' and 'Clear All' are visible at the bottom of the assignment list.

**Resolution**

N/A

## 3.7. Windows Aero is disabled on XenDesktop session using 3 or 4 monitors in 2560×1600 resolution

**Description**

Windows Aero may be disabled when XenDesktop is connected to a VM with a vGPU or passthrough GPU, with 3 or 4 monitors at 2560×1600 resolution.

This limitation is a limitation of Windows 7. For details, see the Microsoft knowledge base article [Desktop background disappears with very large extended desktop on Windows 7](#).

## 3.8. 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 Citrix XenServer 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
```

## Resolution

vGPU reserves a portion of the VM's framebuffer for use in GPU mapping of VM system memory. The reservation is sufficient to support up to 32GB of system memory, and may be increased to accommodate up to 64GB by specifying `enable_large_sys_mem=1` in the VM's `platform:vgpu_extra_args` parameter:

```
[root@xenserver ~]# xe vm-param-set uuid=e71afda4-53f4-3a1b-6c92-a364a7f619c2 platform:vgpu_extra_args="enable_large_sys_mem=1"
```

The setting takes effect the next time the VM is started or rebooted. With this setting in place, less GPU FB is available to applications running in the VM. To accommodate system memory larger than 64GB, the reservation can be further increased by specifying `extra_fb_reservation` in the VM's `platform:vgpu_extra_args` parameter, and setting its value to the desired reservation size in megabytes. The default value of 64M is sufficient to support 64GB of RAM. We recommend adding 2M of reservation for each additional 1GB of system memory. For example, to support 96GB of RAM, set `extra_fb_reservation` to 128:

```
platform:vgpu_extra_args="enable_large_sys_mem=1, extra_fb_reservation=128"
```

The reservation can be reverted back to its default setting in one of the following ways:

- ▶ Removing the `vgpu_extra_args` key from the `platform` parameter
- ▶ Removing `enable_large_sys_mem` from the `vgpu_extra_args` key
- ▶ Setting `enable_large_sys_mem=0`

## 3.9. vGPU host driver RPM upgrade fails

### Description

Upgrading vGPU host driver RPM fails with an error message about failed dependencies on the console.

```
[root@xenserver ~]# rpm -U NVIDIA-vGPU-xenserver-6.5-352.46.x86_64.rpm
error: Failed dependencies:
    NVIDIA-vgx-xenserver conflicts with NVIDIA-vGPU-xenserver-6.5-352.46.x86_64
[root@xenserver ~]#
```

### Resolution

Uninstall the older vGPU RPM before installing the latest driver.

Use the following command to uninstall the older vGPU RPM:

```
[root@xenserver ~]# rpm -e NVIDIA-vgx-xenserver
```

# 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 5.0

No resolved issues are reported in this release for Citrix XenServer.

### Issues Resolved in Release 5.1

Bug ID	Summary and Description
200307528	<p><b>Ubuntu 16.04 VMs run slowly after acquiring a license</b></p> <p>Ubuntu 16.04 VMs to which an NVIDIA vGPU is assigned run slowly after acquiring a license. Ubuntu 16.04 VMs that have not been assigned an NVIDIA vGPU run noticeably faster.</p>
1967625	<p><b>NVIDIA vGPU utilization in guest VMs is incorrectly reported for vGPUs using a fixed share scheduler</b></p> <p>3D engine utilization in guest VMs should be reported as a percentage of the maximum physical GPU utilization the vGPU can consume. However, the utilization for vGPUs using a fixed share scheduler does not correctly account for the vGPU's fixed share of the physical GPU cycles. Furthermore, the utilization reported by <code>nvidia-smi pmon</code> is inconsistent with the utilization reported by <code>nvidia-smi dmon</code>.</p>

## Issues Resolved in Release 5.2

Bug ID	Summary and Description
200359618	<p><b>On GPUs based on the Pascal architecture, Ubuntu 16.04 VMs run slowly after acquiring a license</b></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><b>The Apply button is disabled after change to unlicensed mode</b></p> <p>After the mode is changed from licensed <b>Quadro Virtual Datacenter Workstation Edition</b> mode to <b>Unlicensed Tesla</b> mode, the <b>Apply</b> button on the <b>Manage GRID License</b> page is disabled. As a result, <b>NVIDIA X Server Settings</b> cannot be used to switch to <b>Tesla (Unlicensed)</b> mode on a licensed system.</p>

## Issues Resolved in Release 5.3

Bug ID	Summary and Description
200334648	<p><b>Multiple display heads are not detected by Ubuntu 14.04 guest VMs</b></p> <p>After an Ubuntu 14.04 guest VM has acquired a license, multiple display heads connected to the VM are not detected.</p>
2075467	<p><b>The displays flicker each time a license is requested or renewed in Linux guest VMs</b></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><b>The license expires prematurely in Linux guest VMs</b></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><b>Issues in remote desktop sessions if a license is acquired after a session is started</b></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</p>

Bug ID	Summary and Description
	in the session. For example, the updated maximum resolution supported is not available.

#### Issues Resolved in Release 5.4

No resolved issues are reported in this release for Citrix XenServer.

# 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 Citrix XenServer.

### NVIDIA Software Security Updates in Release 5.4

No NVIDIA software security updates are reported in this release for Citrix XenServer.



# 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. NVIDIA Control Panel Crashes in a VM connected to two 4K displays

### Description

When **NVIDIA Control Panel** is started in a VM connected through a pass-through GPU to two 4K displays, a **Citrix HDX 3D Pro Warning** pop-up window opens containing the following error message:

```
Full screen applications are not supported on Citrix HDX 3D Pro. Press Ok to exit
```

**NVIDIA Control Panel** then crashes.

After the VM is rebooted, **NVIDIA Control Panel** is available again.

### Version

Citrix XenDesktop 7.16

### Status

Not an NVIDIA bug

### Ref. #

200393766

## 6.3. vGPU guest VM driver not properly loaded on servers with more than 512 GB or 1 TB or more of system memory

### Description

If PV IOMMU is enabled, support for vGPU is limited to servers with a maximum of 512 GB of system memory. On servers with more than 512 GB of system memory and PV IOMMU enabled, the guest VM driver is not properly loaded. **Device Manager** marks the vGPU with a yellow exclamation point.

If PV IOMMU is disabled, support for vGPU is limited to servers with less than 1 TB of system memory. This limitation applies only to systems with supported GPUs based on the Maxwell architecture: Tesla M6, Tesla M10, and Tesla M60. On servers with 1 TB or more of system memory, VMs configured with vGPU fail to power on. However, support for GPU pass through is not affected by this limitation.

## Resolution

1. If PV IOMMU is enabled, disable it.

```
[root@xenserver ~]# /opt/xensource/libexec/xen-cmdline --set-xen iommu=dom0-passthrough
```

2. If the server has 1 TB or more of system memory, limit the amount of system memory on the server to 1 TB minus 16 GB.

If the server has less than 1 TB or more of system memory, omit this step.

3. Reboot the server.

## Status

Not an NVIDIA bug

## Ref. #

1799582

## 6.4. 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

## 6.5. 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

## 6.6. 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

## 6.7. 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

## 6.8. 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

## 6.9. CILP is not working in Windows VMs

**Description**

CILP is not working in Windows VMs. If a CUDA application is running in one VM and a graphics application is running in another VM, the following errors occur:

- ▶ The CUDA application times out.
- ▶ A TDR is triggered in the VM that is running graphics application, which may cause flickering or an application crash.

CILP isn't expected to work on Windows until Windows 10 RS3.

**Version**

Windows 10 RS2 Creators Update

**Status**

Open

**Ref. #**

200333574

## 6.10. Luxmark causes a segmentation fault on an unlicensed Linux client

**Description**

If the Luxmark application is run on a Linux guest VM configured with NVIDIA vGPU that is booted without acquiring a license, a segmentation fault occurs and the application core dumps. The fault occurs when the application cannot allocate a

CUDA object on NVIDIA vGPUs where CUDA is disabled. On NVIDIA vGPUs that can support CUDA, CUDA is disabled in unlicensed mode.

**Status**

Not an NVIDIA bug.

**Ref. #**

200330956

## 6.11. Ubuntu 16.04 VMs run slowly after acquiring a license

**Description**

Ubuntu 16.04 VMs to which an NVIDIA vGPU is assigned run slowly after acquiring a license. Ubuntu 16.04 VMs that have not been assigned an NVIDIA vGPU run noticeably faster.

**Version**

Ubuntu 16.04

**Workaround**

After the VM has acquired a license, restart the Compiz compositing window manager

**Status**

Resolved in version 5.1

**Ref. #**

200307528

## 6.12. 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

## 6.13. 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



## 6.14. On Tesla P40, P6, and P4 GPUs, the default ECC setting prevents NVIDIA vGPU from starting

### Description

On Tesla P40, Tesla P6, and Tesla P4 GPUs, the default error-correcting code (ECC) memory setting prevents NVIDIA vGPU from starting. By default, ECC memory is enabled on these GPUs, but NVIDIA vGPU does not support ECC memory.

### Workaround

Before running NVIDIA vGPU, disable ECC memory as explained in *Virtual GPU Software User Guide*.

### Status

Closed.

### Ref. #

200269717

## 6.15. NVIDIA vGPU utilization in guest VMs is incorrectly reported for vGPUs using a fixed share scheduler

### Description

NVIDIA vGPU utilization in guest VMs is incorrectly reported for vGPUs using a fixed share scheduler.

3D engine utilization in guest VMs should be reported as a percentage of the maximum physical GPU utilization the vGPU can consume. For example, a vGPU that is allocated a 25% share of the physical GPU should be reported as fully utilized (100%) when it is consuming 25% of the physical GPU's cycles. However, the utilization for vGPUs using a fixed share scheduler does not correctly account for the vGPU's fixed share of the physical GPU cycles.

Furthermore, the utilization reported by `nvidia-smi pmon` is inconsistent with the utilization reported by `nvidia-smi dmon`.

**Status**

Resolved in version 5.1

**Ref. #**

1967625

## 6.16. NVIDIA vGPU encoder and process utilization counters don't work with Windows Performance Counters

**Description**

GPU encoder and process utilization counter groups are listed in Windows Performance Counters, but no instances of the counters are available. The counters are disabled by default and must be enabled.

**Workaround**

Enable the counters by running the following sequence of commands from a command shell:

```
wmic /namespace:nv path System call enableProcessUtilizationPerfCounter
```

```
wmic /namespace:nv path System call enableEncoderSessionsPerfCounter
```

If you need to disable the counters, run the following sequence of commands from a command shell:

```
wmic /namespace:nv path System call disableProcessUtilizationPerfCounter
```

```
wmic /namespace:nv path System call disableEncoderSessionsPerfCounter
```

**Status**

Open

**Ref. #**

1971698

## 6.17. 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

## 6.18. 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

## 6.19. 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

## 6.20. 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

## 6.21. 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

## 6.22. 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.

When memory exhaustion occurs, the NVIDIA host driver reports Xid error 31 and Xid error 43 in XenServer's `/var/log/messages` file.

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](#).

For more information, see also [Windows 10 Optimization for XenDesktop](#) on the Citrix blog.

### Status

Open

### Ref. #

- ▶ 200130864
- ▶ 1803861

## 6.23. VM bug checks after the guest VM driver for Windows 10 RS2 is installed

### Description

When the VM is rebooted after the guest VM driver for Windows 10 RS2 is installed, the VM bug checks. When Windows boots, it selects one of the standard supported video modes. If Windows is booted directly with a display that is driven by an NVIDIA driver, for example a vGPU on Citrix XenServer, a blue screen crash occurs.



This issue occurs when the screen resolution is switched from VGA mode to a resolution that is higher than 1920×1200.

**Fix**

Download and install [Microsoft Windows Update KB4020102](#) from the Microsoft Update Catalog.

**Workaround**

If you have applied the fix, ignore this workaround.

Otherwise, you can work around this issue until you are able to apply the fix by not using resolutions higher than 1920×1200.

1. Choose a GPU profile in Citrix XenCenter that does not allow resolutions higher than 1920×1200.
2. Before rebooting the VM, set the display resolution to 1920×1200 or lower.

**Status**

Not an NVIDIA bug

**Ref. #**

200310861

## 6.24. On XenServer 7.0, VMs unexpectedly reboot and XenServer crashes or freezes

**Description**

On XenServer 7.0, VMs to which a vGPU is attached unexpectedly reboot and XenServer crashes or freezes.

The event log in XenServer's `/var/log/crash/xen.log` file lists the following errors:

- ▶ A fatal bus error on a component at the slot where the GPU card is installed
- ▶ A fatal error on a component at bus 0, device 2, function 0

This issue occurs when page-modification logging (PML) is enabled on Intel Broadwell CPUs running XenServer 7.0. Citrix is aware of this issue and is working on a permanent fix.

**Workaround**

Disable page-modification logging (PML) as explained in [XenServer 7 host crash while starting multiple virtual machines](#) in the Citrix Support Knowledge Center.

**Status**

Not an NVIDIA bug

**Ref. #**

1853248

## 6.25. With no NVIDIA driver installed, XenServer misidentifies Tesla M10 cards

**Description**

An erroneous entry in the `pci.ids` database causes Citrix XenServer to identify Tesla M10 cards as GRID M40 when no NVIDIA driver is installed.

**Version**

Citrix XenServer 6.5 and 7.0

**Workaround**

None

**Status**

Not an NVIDIA bug

**Ref. #**

NVIDIA-420/1792341

## 6.26. 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

## 6.27. Video goes blank when run in loop in Windows Media Player

### Description

When connected to a vGPU-enabled VM using Citrix XenDesktop, a video played back in looping mode on Windows Media Player goes blank or freezes after a few iterations.

### Workaround

None

### Status

Not an NVIDIA bug

### Ref. #

1306623

## 6.28. Local VGA console is momentarily unblanked when XenDesktop changes resolution of the VM desktop

### Description

When XenDesktop establishes a remote connection to a VM using vGPU, the VM's local VGA console display in XenCenter is blanked (assuming the VM local console has not been disabled by setting `platform:vgpu_extra_args="disable_vnc=1"`). If the XenDesktop session changes resolution of the VM's desktop, the local VGA console momentarily unblanks, allowing a XenCenter user to briefly view the desktop.

### Workaround

Disable the VM's local VGA console

```
xe vm-param-set uuid=vm-uuid platform:vgpu_extra_args="disable_vnc=1"
```

### Status

Open

### Ref. #

NVIDIA-145/1375164

## 6.29. VM bugchecks on shutdown/restart when XenDesktop is installed and NVIDIA driver is uninstalled or upgraded.

### Description

If the XenDesktop agent is installed in a VM before any NVIDIA GPU driver is installed, the VM will bugcheck (bluescreen) when the NVIDIA driver is subsequently upgraded or uninstalled. The bugcheck code is 0x7E, `SYSTEM_THREAD_EXCEPTION_NOT_HANDLED`.

### Workaround

Use one of the following workarounds:

- ▶ Do a force shutdown of the VM and restart it.

- ▶ Install the NVIDIA driver in guest VMs before installing XenDesktop.

**Status**

Open

**Ref. #**

NVIDIA-295/200018125

## 6.30. Application frame rate may drop when running XenDesktop at 2560×1600 resolution.

**Description**

An application's rendering frame rate may drop when running XenDesktop at 2560×1600 resolution, relative to the frame rate obtained at lower resolutions.

**Fix**

Using the Windows `regedit` utility within the VM, open the `HKLM\SOFTWARE\Citrix\Graphics` registry key and create a new `DWORD` value, `EncodeSpeed`, with a value of 2. Reboot the VM. This setting may improve the delivered frame rate at the expense of a reduction in image quality.

**Status**

Open

**Ref. #**

NVIDIA-190/1416336

## 6.31. Windows VM BSOD

**Description**

Windows VM bugchecks on XenServer when running a large number of vGPU based VMs.

XenServer's `/var/log/messages` file contains these error messages:

```
NVRM: Xid (PCI:0000:08:00): 31, Ch 0000001e, engmask 00000111, intr 10000000
NVRM: Xid (PCI:0000:08:00): 31, Ch 00000016, engmask 00000111, intr 10000000
...
vmiop_log: error: Assertion Failed at 0xb5b898d8:4184
vmiop_log: error: 8 frames returned by backtrace
```

```

vmiop_log: error: /usr/lib/libnvidia-vgx.so(_nv000793vgx+0x69d) [0xb5b8064d]
vmiop_log: error: /usr/lib/libnvidia-vgx.so(_nv000479vgx+0x118) [0xb5b898d8]
vmiop_log: error: /usr/lib/libnvidia-vgx.so(_nv000782vgx+0x59) [0xb5b85f49]
vmiop_log: error: /usr/lib/libnvidia-vgx.so(_nv000347vgx+0x3db) [0xb5b932db]
vmiop_log: error: /usr/lib/libnvidia-vgx.so [0xb5b78e4a]
vmiop_log: error: /usr/lib/xen/bin/vgpu [0x80554be]
vmiop_log: error: /lib/libpthread.so.0 [0xb7612912]
vmiop_log: error: /lib/libc.so.6(clone+0x5e) [0xb76fc5ee]
vmiop_log: error: failed to initialize guest PTE entries
vmiop_log: error: failed to fill up guest PTE entries 3
vmiop_log: error: VGPU message 27 failed, result code: 0xff000003
vmiop_log: error:          0xc1d00001, 0xff010000, 0x1a77ba000, 0x0, 0x1,
vmiop_log: error:          0x1, 0x1000, 0x10202, 0xc1d00001, 0xff010000,
vmiop_log: error:          0xcaf00004, 0x0
vmiop_log: error: Timeout occurred, reset initiated.

```

**Version**

XenServer 6.2

**Fix**

Ensure that you are running the latest OEM firmware for your NVIDIA vGPU software boards.

**Status**

Closed

**Ref. #**

NVIDIA-327/1632120

## 6.32. Windows VM BSOD when upgrading NVIDIA drivers over a XenDesktop session

**Description**

Windows VM bugchecks when NVIDIA guest drivers are upgraded over a XenDesktop session.

If the VM is restarted after the bugcheck, the upgraded driver loads correctly and full functionality is available.

**Fix**

Upgrade XenDesktop to 7.6 Feature Pack 3

**Status**

Closed

**Ref. #**

NVIDIA-370/200130780

## 6.33. XenCenter does not allow vGPUs to be selected as a GPU type for Linux VMs

**Description**

When creating a new Linux VM or editing the properties of an existing Linux VM, XenCenter does not allow vGPUs to be selected as a GPU type.

vGPU on Linux VMs is supported as a technical preview on XenServer 6.5, and does include XenCenter integration.

**Version**

Affects the XenCenter integration with XenServer 6.5 only.

Resolved in the XenCenter integration with XenServer 7.0.

**Workaround**

Refer to XenServer vGPU Management in *Virtual GPU Software User Guide* for how to configure vGPU by using the `xm` CLI.

**Status**

Closed

**Ref. #**

NVIDIA-360

## 6.34. If X server is killed on a RHEL7 VM running vGPU, XenCenter console may not automatically switch to text console

**Description**

If X server is killed on a RHEL7 VM running vGPU, XenCenter console may display a corrupted image and fail to switchover to text console.

The failure to switchover to text console is due to a bug in RHEL7, which causes X server to not start correctly under certain configurations.

**Workaround**

Use **CTRL+ALT+F1**, **F2**, or **F3** to switch between Linux terminals.

**Status**

Closed

**Ref. #**

NVIDIA-350/200123378

## 6.35. XenDesktop shows only a black screen when connected to a vGPU VM

**Description**

XenDesktop sometimes displays only a black screen when it is connected to an NVIDIA vGPU VM. The probable cause is that the display that is connected to the NVIDIA vGPU is entering a lower power state.

**Fix**

Disable all display-related power management settings.

For detailed instructions, visit [Microsoft power plans frequently asked questions](#) and from the list, select your OS version.

**Status**

Not an NVIDIA bug

**Ref. #**

1719877



## 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.