



# VIRTUAL GPU SOFTWARE R390 FOR CITRIX XENSERVER

RN-06927-001 \_v6.0 through 6.4 Revision 02 | May 2020

## Release Notes



# TABLE OF CONTENTS

<b>Chapter 1. Release Notes</b> .....	<b>1</b>
1.1. Updates in Release 6.0.....	2
1.2. Updates in Release 6.1.....	2
1.3. Updates in Release 6.2.....	3
1.4. Updates in Release 6.3.....	3
1.5. Updates in Release 6.4.....	4
<b>Chapter 2. Validated Platforms</b> .....	<b>5</b>
2.1. Supported NVIDIA GPUs and Validated Server Platforms.....	5
2.2. Hypervisor Software Releases.....	6
2.3. Guest OS Support.....	8
2.3.1. Windows Guest OS Support.....	8
2.3.2. Linux Guest OS Support.....	11
2.4. NVIDIA CUDA Toolkit Version Support.....	12
2.5. XenMotion with vGPU Support.....	12
<b>Chapter 3. Known Product Limitations</b> .....	<b>14</b>
3.1. vGPU profiles with 512 Mbytes or less of frame buffer support only 1 virtual display head on Windows 10.....	14
3.2. NVENC requires at least 1 Gbyte of frame buffer.....	15
3.3. VM running older NVIDIA vGPU drivers fails to initialize vGPU when booted.....	15
3.4. Virtual GPU fails to start if ECC is enabled.....	16
3.5. Single vGPU benchmark scores are lower than pass-through GPU.....	18
3.6. nvidia-smi fails to operate when all GPUs are assigned to GPU pass-through mode.....	19
3.7. Windows Aero is disabled on XenDesktop session using 3 or 4 monitors in 2560×1600 resolution.....	20
3.8. VMs configured with large memory fail to initialize vGPU when booted.....	20
3.9. vGPU host driver RPM upgrade fails.....	21
<b>Chapter 4. Resolved Issues</b> .....	<b>22</b>
<b>Chapter 5. Security Updates</b> .....	<b>25</b>
5.1. Since 6.4: Restricting Access to GPU Performance Counters.....	25
5.1.1. Windows: Restricting Access to GPU Performance Counters for One User by Using NVIDIA Control Panel.....	25
5.1.2. Windows: Restricting Access to GPU Performance Counters Across an Enterprise by Using a Registry Key.....	26
5.1.3. Linux Guest VMs and Hypervisor Host: Restricting Access to GPU Performance Counters.....	26
<b>Chapter 6. Known Issues</b> .....	<b>28</b>
6.1. NVIDIA Notification Icon prevents log off of Citrix Published Application user sessions.....	28
6.2. Frame capture while the interactive logon message is displayed returns blank screen.....	29
6.3. Cloned VMs configured with a vGPU type different than the type in the master image fail to start.....	29

6.4. Migrating a VM configured with vGPU can cause errors.....	30
6.5. Even when the scheduling policy is equal share, unequal GPU utilization is reported.....	30
6.6. When the scheduling policy is fixed share, GPU utilization is reported as higher than expected.....	31
6.7. Since 6.1: nvidia-smi reports that vGPU migration is supported on all hypervisors.....	32
6.8. NVIDIA Control Panel Crashes in a VM connected to two 4K displays.....	33
6.9. Screen resolution reverts to a lower value after a VM is rebooted.....	33
6.10. 6.1 Only: Benign Calling load_byte_array(tra) messages are logged.....	34
6.11. 6.0 Only: Migrating a VM Configured with vGPU Fails.....	35
6.12. vGPU guest VM driver not properly loaded on servers with more than 512 GB or 1 TB or more of system memory.....	35
6.13. 6.0 Only: Issues in remote desktop sessions if a license is acquired after a session is started.....	36
6.14. 6.0-6.2 Only: Licensing pop-up windows contain the text microsoft.explorer.notification.	37
6.15. Luxmark causes a segmentation fault on an unlicensed Linux client.....	37
6.16. Resolution is not updated after a VM acquires a license and is restarted.....	38
6.17. On Tesla P40, P6, and P4 GPUs, the default ECC setting prevents NVIDIA vGPU from starting.....	38
6.18. NVIDIA vGPU encoder and process utilization counters don't work with Windows Performance Counters.....	39
6.19. A segmentation fault in DBus code causes nvidia-gridd to exit on Red Hat Enterprise Linux and CentOS.....	39
6.20. No Manage License option available in NVIDIA X Server Settings by default.....	40
6.21. Licenses remain checked out when VMs are forcibly powered off.....	41
6.22. Memory exhaustion can occur with vGPU profiles that have 512 Mbytes or less of frame buffer.....	42
6.23. VM bug checks after the guest VM driver for Windows 10 RS2 is installed.....	43
6.24. On XenServer 7.0, VMs unexpectedly reboot and XenServer crashes or freezes.....	44
6.25. With no NVIDIA driver installed, XenServer misidentifies Tesla M10 cards.....	44
6.26. GNOME Display Manager (GDM) fails to start on Red Hat Enterprise Linux 7.2 and CentOS 7.0.....	45
6.27. Video goes blank when run in loop in Windows Media Player.....	46
6.28. Local VGA console is momentarily unblanked when XenDesktop changes resolution of the VM desktop.....	46
6.29. VM bugchecks on shutdown/restart when XenDesktop is installed and NVIDIA driver is uninstalled or upgraded.....	47
6.30. Application frame rate may drop when running XenDesktop at 2560×1600 resolution....	47
6.31. Windows VM BSOD.....	48
6.32. Windows VM BSOD when upgrading NVIDIA drivers over a XenDesktop session.....	49
6.33. XenCenter does not allow vGPUs to be selected as a GPU type for Linux VMs.....	49
6.34. If X server is killed on a RHEL7 VM running vGPU, XenCenter console may not automatically switch to text console.....	50
6.35. XenDesktop shows only a black screen when connected to a vGPU VM.....	50



# 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	6.0	6.1	6.2	6.3	6.4
NVIDIA Virtual GPU Manager for the Citrix XenServer releases listed in <a href="#">Hypervisor Software Releases</a>	390.42	390.57	390.72	390.94	390.113
NVIDIA Windows driver	391.03	391.58	391.81	392.05	392.37
NVIDIA Linux driver	390.42	390.57	390.75	390.96	390.115



### Caution

If you install the wrong NVIDIA vGPU software packages 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 6.1 with guest VM drivers from release 6.0. 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 6.1 with guest VM drivers from release 5.2.

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 6.0

### New Features in Release 6.0

- ▶ XenMotion support for VMs that are configured with vGPU
- ▶ New -2B vGPU type for each supported GPU
- ▶ vGPU support for NVML accounting functions
- ▶ vGPU support for `nvidia-smi` accounting modes
- ▶ Change of default scheduler to best effort scheduler for GPUs based on the NVIDIA<sup>®</sup> Pascal<sup>™</sup> architecture
- ▶ Change of maximum resolution for unlicensed GPUs based on the Pascal architecture to 1280×1024
- ▶ Plain-text logging on Windows of significant licensing events
- ▶ New setting `EnableLogging` for disabling or enabling logging of significant licensing events
- ▶ Miscellaneous bug fixes

### Hardware and Software Support Introduced in Release 6.0

- ▶ Support for GPUs based on the NVIDIA<sup>®</sup> Volta architecture
- ▶ Support for Citrix XenServer 7.4
- ▶ Support for Windows 10 Fall Creators Update (1709) as a guest OS

### Feature Support Withdrawn in Release 6.0

- ▶ Citrix XenServer 6.5 SP1 is no longer supported.

## 1.2. Updates in Release 6.1

### New Features in Release 6.1

- ▶ Change in behavior to enable XenMotion support for VMs that are configured with vGPU by default
- ▶ Change in behavior to enable logging of licensing events on Windows by default
- ▶ Miscellaneous bug fixes

### Hardware and Software Support Introduced in Release 6.1

- ▶ Support for Citrix XenDesktop 7.17
- ▶ Support for the Tesla V100 SXM2 32GB GPU
- ▶ Support for the Tesla V100 PCIe 32GB GPU

### Feature Support Withdrawn in Release 6.1

- ▶ Citrix XenDesktop 7.6, 7.8, and 7.11 are no longer supported.

## 1.3. Updates in Release 6.2

### New Features in Release 6.2

- ▶ New -2B4 vGPU type, which supports up to four displays at resolutions up to 2560×1600 for each supported GPU
- ▶ Relaxation of restrictions on unlicensed vGPUs to allow screen resolutions higher than 1280×1024
- ▶ Miscellaneous bug fixes

### Hardware and Software Support Introduced in Release 6.2

- ▶ Support for Citrix XenServer 7.5
- ▶ Support for Citrix XenDesktop 7.18

### Feature Support Withdrawn in Release 6.2

- ▶ Citrix XenServer 7.2 and 7.3 are no longer supported.

## 1.4. Updates in Release 6.3

### New Features in Release 6.3

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

### Hardware and Software Support Introduced in Release 6.3

- ▶ Support for Citrix XenServer 7.6
- ▶ Support for Citrix Virtual Apps and Desktops (formerly Citrix XenDesktop) 7 1808

### Feature Support Withdrawn in Release 6.3

- ▶ Citrix XenServer 7.4 is no longer supported.

## 1.5. Updates in Release 6.4

### New Features in Release 6.4

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



# Chapter 2.

## VALIDATED PLATFORMS

This release family of NVIDIA vGPU software provides support for several NVIDIA GPUs on validated server hardware platforms, Citrix XenServer hypervisor software versions, and guest operating systems. It also supports the version of NVIDIA CUDA Toolkit that is compatible with R390 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 PCIe 16 GB (XenMotion with vGPU is **not** supported.)
- ▶ Tesla P100 SXM2 16 GB (XenMotion with vGPU is **not** supported.)
- ▶ Tesla P100 PCIe 12GB (XenMotion with vGPU is **not** supported.)
- ▶ Tesla V100 SXM2 (XenMotion with vGPU is **not** supported.)
- ▶ Tesla V100 SXM2 32GB (XenMotion with vGPU is **not** supported.)
- ▶ Tesla V100 PCIe (XenMotion with vGPU is **not** supported.)
- ▶ Tesla V100 PCIe 32GB (XenMotion with vGPU is **not** supported.)
- ▶ Tesla V100 FHHL (XenMotion with vGPU is **not** supported.)

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

### Changes to Citrix Product Names and Release Numbers

Citrix has changed the product names of its hypervisor and virtual desktop software.

Old Name	New Name
XenServer	Citrix Hypervisor
XenApp and XenDesktop	Citrix Virtual Apps and Desktops

The release number of Citrix Virtual Apps and Desktops releases **after 7.18** is in the form  $7\ yymm$ , where  $yymm$  denotes the year and month that the product was released. For example, for a product released in August 2018,  $yymm$  is 1808.

For more information, see these topics in the Citrix documentation:

- ▶ [New product names](#)
- ▶ [New product and component version numbers](#)

### Supported Citrix XenServer Releases

This release family of NVIDIA vGPU software is supported on the Citrix XenServer Hypervisor 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.

Cumulative update releases for a base release of Citrix XenServer 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	Notes
Since 6.3: Citrix XenServer 7.6	RTM build and compatible cumulative update releases	All NVIDIA GPUs that support NVIDIA vGPU software are supported.  This release supports XenMotion with vGPU on suitable GPUs as listed in <a href="#">Supported NVIDIA GPUs and Validated Server Platforms</a> .
Since 6.2: Citrix XenServer 7.5	RTM build and compatible cumulative update releases	All NVIDIA GPUs that support NVIDIA vGPU software are supported.  This release supports XenMotion with vGPU on suitable GPUs as listed in <a href="#">Supported NVIDIA GPUs and Validated Server Platforms</a> .
6.0-6.2 only: Citrix XenServer 7.4	RTM build and compatible cumulative update releases	All NVIDIA GPUs that support NVIDIA vGPU software are supported.  This release supports XenMotion with vGPU on suitable GPUs as listed in <a href="#">Supported NVIDIA GPUs and Validated Server Platforms</a> .
6.0, 6.1 only: Citrix XenServer 7.3	RTM build and compatible cumulative update releases	All NVIDIA GPUs that support NVIDIA vGPU software are supported.  XenMotion with vGPU is <b>not</b> supported.
6.0, 6.1 only: Citrix XenServer 7.2	RTM build and compatible cumulative update releases	All NVIDIA GPUs that support NVIDIA vGPU software are supported.  XenMotion with vGPU is <b>not</b> supported.
Citrix XenServer 7.1	RTM build and compatible cumulative update releases	All NVIDIA GPUs that support NVIDIA vGPU software are supported.  XenMotion with vGPU is <b>not</b> supported.

Software	Releases Supported	Notes
Citrix XenServer 7.0	RTM build 125380 and compatible cumulative update releases	All NVIDIA GPUs that support NVIDIA vGPU software are supported.  XenMotion with vGPU is <b>not</b> supported.

## 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 6.3:</b> Versions 7 1808, 7.18, 7.17, and 7.15 in HDX 3D Pro mode</p> <p><b>6.2 only:</b> Versions 7.18, 7.17, and 7.15 in HDX 3D Pro mode</p> <p><b>6.1 only:</b> Versions 7.17, and 7.15 in HDX 3D Pro mode</p> <p><b>6.0 only:</b> Versions 7.15, 7.11, 7.8, and 7.6 in HDX 3D Pro mode</p>

## 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 or Volta architecture support only 64-bit guest operating systems. No 32-bit guest operating systems are supported in pass-through mode for these GPUs.

### 2.3.1. Windows Guest OS Support

NVIDIA vGPU software supports **only** the Windows releases listed 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. When configured for NVIDIA vGPU, GPUs based on the Volta architecture support only

64-bit Windows guest operating systems. No 32-bit Windows guest operating systems are supported for these GPUs when configured for NVIDIA vGPU.

 If a specific release, even an update release, is not listed, it's not supported.  
XenMotion with vGPU is supported on supported Windows guest OS releases.

Guest OS	NVIDIA vGPU - Citrix XenServer Releases	Pass-Through GPU - Citrix XenServer Releases
Windows Server 2016 1607, 1709	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>
Windows Server 2012 R2	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>
Windows Server 2008 R2	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p> <p>Supported only on GPUs based on the Maxwell architecture</p>
Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703), Fall Creators Update (1709) (64-bit)	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>
Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703), Fall Creators Update (1709) (32-bit)	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p> <p>Supported only on GPUs based on the Maxwell architecture</p>

Guest OS	NVIDIA vGPU - Citrix XenServer Releases	Pass-Through GPU - Citrix XenServer Releases
Windows 8.1 Update (64-bit)	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>
Windows 8.1 Update (32-bit)	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p> <p>Supported only on GPUs based on the Maxwell architecture</p>
Windows 8.1 (64-bit)	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>	-
Windows 8.1 (32-bit)	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p> <p>Supported only on GPUs based on the Maxwell and Pascal architectures</p>	-
Windows 8 (64-bit)	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>	-
Windows 8 (32-bit)	<p>Since 6.3: 7.6, 7.5, 7.1, 7.0</p> <p>6.2 only: 7.5, 7.4, 7.1, 7.0</p> <p>6.0, 6.1 only: 7.4, 7.3, 7.2, 7.1, 7.0</p>	-

Guest OS	NVIDIA vGPU - Citrix XenServer Releases	Pass-Through GPU - Citrix XenServer Releases
	Supported only on GPUs based on the Maxwell and Pascal architectures	
Windows 7 (64-bit)	<p><b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0</p> <p><b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0</p> <p><b>6.0, 6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0</p>	<p><b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0</p> <p><b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0</p> <p><b>6.0, 6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0</p> <p>Supported only on GPUs based on the Maxwell architecture</p>
Windows 7 (32-bit)	<p><b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0</p> <p><b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0</p> <p><b>6.0, 6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0</p> <p>Supported only on GPUs based on the Maxwell and Pascal architectures</p>	<p><b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0</p> <p><b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0</p> <p><b>6.0, 6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0</p> <p>Supported only on GPUs based on the Maxwell architecture</p>

## 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.  
XenMotion with vGPU is **not** supported on any Linux guest OS release.

Guest OS	NVIDIA vGPU - Citrix XenServer Releases	Pass-Through GPU - Citrix XenServer Releases
Red Hat Enterprise Linux 7.0-7.4	<p><b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0</p> <p><b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0</p> <p><b>6.0, 6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0</p>	<p><b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0</p> <p><b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0</p> <p><b>6.0, 6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0</p>
Since 6.1: Red Hat Enterprise Linux 7.5	<b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0	<b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0

Guest OS	NVIDIA vGPU - Citrix XenServer Releases	Pass-Through GPU - Citrix XenServer Releases
	<b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0 <b>6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0	<b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0 <b>6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0
CentOS 7.0-7.4	<b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0 <b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0 <b>6.0, 6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0	<b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0 <b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0 <b>6.0, 6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0
Since 6.1: CentOS 7.5	<b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0 <b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0 <b>6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0	<b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0 <b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0 <b>6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0
Ubuntu 16.04 LTS	<b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0 <b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0 <b>6.0, 6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0	<b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0 <b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0 <b>6.0, 6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0
Ubuntu 14.04 LTS	<b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0 <b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0 <b>6.0, 6.1 only:</b> 7.4, 7.3, 7.2, 7.1, 7.0	<b>Since 6.3:</b> 7.6, 7.5, 7.1, 7.0 <b>6.2 only:</b> 7.5, 7.4, 7.1, 7.0 <b>6.0, 6.1 only:</b> 7.4, 7.3, 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 9.1.

For more information about NVIDIA CUDA Toolkit, see [CUDA Toolkit 9.1 Documentation](#).

## 2.5. XenMotion with vGPU Support

XenMotion with vGPU is supported only on a subset of supported GPUs, Citrix XenServer Hypervisor releases, and guest operating systems.



Supported GPUs:

- ▶ Tesla M6
- ▶ Tesla M10
- ▶ Tesla M60
- ▶ Tesla P4
- ▶ Tesla P6
- ▶ Tesla P40

Supported Citrix XenServer Hypervisor releases: 7.6, 7.5 and 7.4 only.

Supported guest OS releases: Windows only. XenMotion with vGPU is **not** supported on Linux.

# 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 host'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 pass-through mode

### Description

If all GPUs in the platform are assigned to VMs in pass-through 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 pass-through 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 pass-through mode, 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 for a GK104GL [GRID K2] GPU. The assignment is set to 'Maximum density: put as many VMs as possible on the same GPU'. The GPU is currently assigned to six VMs, all in 'passthrough' mode:

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

On the right side, the 'Allowed vGPU types' are listed with checkmarks:

- 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)

Buttons for 'Edit...', 'Edit Selected GPUs...', 'Select All', and 'Clear All' are visible.

## 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 6.0

Bug ID	Summary and Description
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>
200333574	<p><b>CILP is not working in Windows VMs</b></p> <p>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:</p> <ul style="list-style-type: none"><li>▶ The CUDA application times out.</li><li>▶ A TDR is triggered in the VM that is running graphics application, which may cause flickering or an application crash.</li></ul>

### Issues Resolved in Release 6.1


Bug ID	Summary and Description
-	<p><b>Migrating a VM Configured with vGPU Fails</b></p> <p>Migrating a VM configured with vGPU fails if vGPU migration is not configured in NVIDIA vGPU software.</p>

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

### Issues Resolved in Release 6.2

Bug ID	Summary and Description
2104867	<p><b>Screen resolution reverts to a lower value after a VM is rebooted</b></p> <p>When a VM is booted, the NVIDIA vGPU software graphics driver is initially unlicensed. Screen resolution is limited to a maximum of 1280×1024 until the VM requires a license for NVIDIA vGPU software.</p>
200407382	<p><b>Benign Calling <code>load_byte_array(tra)</code> messages are logged</b></p> <p>In Linux guest VMs, the <code>Calling load_byte_array(tra)</code> messages from the <code>nvidia-gridd</code> daemon are logged in <code>/var/log/syslog</code>.</p>

### Issues Resolved in Release 6.3

Bug ID	Summary and Description
2206368	<p><b>NVIDIA Notification Icon prevents log off of Citrix Published Application user sessions</b></p> <p>By default on Windows Server operating systems, the <b>NVIDIA Notification Icon</b> application is started with every Citrix Published Application user session. This application might prevent the Citrix Published Application user session from being logged off even after the user has quit all other applications.</p> <div style="background-color: #90EE90; padding: 5px; margin-top: 10px;">  <p>This issue is resolved by the provision of Windows registry keys for disabling the <b>NVIDIA Notification Icon</b> application</p> </div>

Bug ID	Summary and Description
	for Citrix Published Application user sessions. For more information, see <i>Virtual GPU Software User Guide</i> .
2194234	<p><b>The mouse cursor moves slowly when traversing the Windows 10 Start menu</b></p> <p>When a user traverses the Windows 10 Start menu with the mouse, highlighting of items on the menu is slow and lags behind the motion of the mouse.</p>
200346607	<p><b>Licensing pop-up windows contain the text <code>microsoft.explorer.notification</code></b></p> <p>On Windows 10 Creators Update (1703), licensing pop-up windows contain the text <code>microsoft.explorer.notification</code>.</p>

### Issues Resolved in Release 6.4

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

# Chapter 5.

## SECURITY UPDATES

### 5.1. Since 6.4: Restricting Access to GPU Performance Counters

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

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

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

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

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

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

1. Open **NVIDIA Control Panel**:
  - ▶ Right-click on the Windows desktop and select **NVIDIA Control Panel** from the menu.
  - ▶ Open **Windows Control Panel** and double-click the **NVIDIA Control Panel** icon.

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

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

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

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



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

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

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

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

2. Restart the VM.

## 5.1.3. Linux Guest VMs and Hypervisor Host: Restricting Access to GPU Performance Counters

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

Perform this task from the guest VM to which the GPU is passed through or from your hypervisor host machine.

In Linux guest VMs, this task requires `sudo` privileges. On your hypervisor host machine, this task must be performed as the root user on the machine.

1. Log in to the guest VM or open a command shell on your hypervisor host machine.
2. Set the kernel module parameter `NVreg_RestrictProfilingToAdminUsers` to 1 by adding this parameter to the `/etc/modprobe.d/nvidia.conf` file.

- ▶ If you are setting only this parameter, add an entry for it to the `/etc/modprobe.d/nvidia.conf` file as follows:

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

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

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

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

3. Restart the VM or reboot your hypervisor host machine.

# Chapter 6.

## KNOWN ISSUES

### 6.1. NVIDIA Notification Icon prevents log off of Citrix Published Application user sessions

#### Description

By default on Windows Server operating systems, the **NVIDIA Notification Icon** application is started with every Citrix Published Application user session. This application might prevent the Citrix Published Application user session from being logged off even after the user has quit all other applications.

#### Since 6.3: Resolution

Disable the **NVIDIA Notification Icon** application for Citrix Published Application user sessions as explained in *Virtual GPU Software User Guide*.

#### Status

Resolved in NVIDIA vGPU software release 6.3 by the provision of Windows registry keys for disabling the **NVIDIA Notification Icon** application for Citrix Published Application user sessions.

#### Ref. #

2206368



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

### Description

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

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

### Workaround

See [Black Screen at Logon with Nvidia HDX 3D Pro enabled](#) in the Citrix knowledge base.

### Status

Not a bug

### Ref. #

2115733

## 6.3. 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

## 6.4. Migrating a VM configured with vGPU can cause errors

**Description**

If the guest VM graphics driver comes from an older release of NVIDIA vGPU software than the Virtual GPU Manager, migrating a VM configured with vGPU causes XID error 43, a TDR, and a VM crash.

For example, this issue occurs if the guest VM graphics driver comes from NVIDIA vGPU software 6.1 and the Virtual GPU Manager comes from NVIDIA vGPU software 6.2.

**Workaround**

Ensure that the guest VM graphics driver and the Virtual GPU Manager come from the same release of NVIDIA vGPU software.

**Status**

Open

**Ref. #**

200426078

## 6.5. Even when the scheduling policy is equal share, unequal GPU utilization is reported

**Description**

When the scheduling policy is equal share, unequal GPU engine utilization can be reported for the vGPUs on the same physical GPU.

For example, GPU engine usage for three P40-8Q vGPUs on a Tesla P40 GPU might be reported as follows:

```
[root@localhost:~] nvidia-smi vgpu
Wed Jun 27 10:33:18 2018
+-----+
| NVIDIA-SMI 390.59                Driver Version: 390.59                |
+-----+-----+-----+
| GPU  Name          | Bus-Id          | GPU-Util  |
|   vGPU ID      Name          | VM ID      VM Name          | vGPU-Util  |
+-----+-----+-----+
| 0   Tesla P40     | 00000000:81:00.0 | 52%       |
|   2122661      GRID P40-8Q     | 2122682   centos7.4-xmpl-211... | 19%       |
|   2122663      GRID P40-8Q     | 2122692   centos7.4-xmpl-211... | 0%        |
|   2122659      GRID P40-8Q     | 2122664   centos7.4-xmpl-211... | 25%       |
+-----+-----+-----+
| 1   Tesla P40     | 00000000:85:00.0 | 58%       |
|   2122662      GRID P40-8Q     | 2122689   centos7.4-xmpl-211... | 0%        |
|   2122658     GRID P40-8Q   | 2122667 centos7.4-xmpl-211... | 59%     |
|   2122660      GRID P40-8Q     | 2122670   centos7.4-xmpl-211... | 0%        |
+-----+-----+-----+
```

The vGPU utilization of the vGPU 2122658 is reported as 59%. However, the expected vGPU utilization should not exceed 33%.

This behavior is a result of the mechanism that is used to measure GPU engine utilization.

## Status

Open

## Ref. #

2175888

## 6.6. When the scheduling policy is fixed share, GPU utilization is reported as higher than expected

### Description

When the scheduling policy is fixed share, GPU engine utilization can be reported as higher than expected for a vGPU.

For example, GPU engine usage for six P40-4Q vGPUs on a Tesla P40 GPU might be reported as follows:

```
[root@localhost:~] nvidia-smi vgpu
Mon Aug 20 10:33:18 2018
+-----+
| NVIDIA-SMI 390.42                Driver Version: 390.42                |
+-----+-----+-----+
| GPU  Name          | Bus-Id          | GPU-Util  |
|   vGPU ID      Name          | VM ID      VM Name          | vGPU-Util  |
+-----+-----+-----+
| 0   Tesla P40     | 00000000:81:00.0 | 99%       |
+-----+-----+-----+
```

	85109	GRID P40-4Q		85110	win7-xmpl-146048-1		32%	
	87195	GRID P40-4Q		87196	win7-xmpl-146048-2		39%	
	88095	GRID P40-4Q		88096	win7-xmpl-146048-3		26%	
	89170	GRID P40-4Q		89171	win7-xmpl-146048-4		0%	
	90475	GRID P40-4Q		90476	win7-xmpl-146048-5		0%	
	93363	GRID P40-4Q		93364	win7-xmpl-146048-6		0%	
+-----+-----+-----+-----+-----+-----+-----+-----+-----+								
	1	Tesla P40		00000000:85:00.0			0%	
+-----+-----+-----+-----+-----+-----+-----+-----+-----+								

The vGPU utilization of vGPU 85109 is reported as 32%. For vGPU 87195, vGPU utilization is reported as 39%. And for 88095, it is reported as 26%. However, the expected vGPU utilization of any vGPU should not exceed approximately 16.7%.

This behavior is a result of the mechanism that is used to measure GPU engine utilization.

### Status

Open

### Ref. #

2227591

## 6.7. Since 6.1: `nvidia-smi` reports that vGPU migration is supported on all hypervisors

### Description

The command `nvidia-smi vgpu -m` shows that vGPU migration is supported on all hypervisors, even hypervisors or hypervisor versions that do not support vGPU migration.

### Status

Closed

### Ref. #

200407230

## 6.8. 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

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

### Description

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

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

**Workaround**

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

**Status**

Resolved in NVIDIA vGPU software release 6.2.

**Ref. #**

2104867

## 6.10. 6.1 Only: Benign Calling `load_byte_array(tra)` messages are logged

**Description**

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

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

**Workaround**

Ignore these messages as they are benign.

**Status**

Resolved in NVIDIA vGPU software release 6.2.

**Ref. #**

200407382

## 6.11. 6.0 Only: Migrating a VM Configured with vGPU Fails

### Description

Migrating a VM configured with vGPU fails if vGPU migration is not configured in NVIDIA vGPU software.

### Resolution

Ensure that vGPU migration is configured in NVIDIA vGPU software to enable support for XenMotion. For instructions, see *Virtual GPU Software User Guide*.

### Status

Resolved in NVIDIA vGPU software release 6.1, in which vGPU migration is enabled by default.

## 6.12. 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

## 6.13. 6.0 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 6.1.

### Ref. #

200391532



## 6.14. 6.0-6.2 Only: Licensing pop-up windows contain the text `microsoft.explorer.notification`

### Description

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

### Version

Windows 10 Creators Update (1703)

### Status

Resolved in NVIDIA vGPU software release 6.3

### Ref. #

200346607

## 6.15. 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

## 6.16. 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

## 6.17. 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

## 6.18. 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

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

### Description

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

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

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

### Version

Red Hat Enterprise Linux 6.8 and 6.9

CentOS 6.8 and 6.9

### Status

Open

### Ref. #

- ▶ 200358191
- ▶ 200319854
- ▶ 1895945

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

### Description

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

### Workaround

This workaround requires `sudo` privileges.



**Do not** use this workaround with Red Hat Enterprise Linux 6.8 and 6.9 or CentOS 6.8 and 6.9. To prevent a segmentation fault in Dbus code from causing the `nvidia-gridd` service from exiting, the GUI for licensing must be disabled with these OS versions.

1. If **NVIDIA X Server Settings** is running, shut it down.

2. If the `/etc/nvidia/gridd.conf` file does not already exist, create it by copying the supplied template file `/etc/nvidia/gridd.conf.template`.
3. As root, edit the `/etc/nvidia/gridd.conf` file to set the `EnableUI` option to `TRUE`.
4. Start the `nvidia-gridd` service.

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

When **NVIDIA X Server Settings** is restarted, the **Manage License** option is now available.

### Status

Open

## 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 `xe` 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.