



VIRTUAL GPU SOFTWARE R410 FOR CITRIX HYPERVISOR

RN-06927-001 _v7.0 through 7.5 Revision 02 | May 2020

Release Notes



TABLE OF CONTENTS

Chapter 1. Release Notes	1
1.1. Updates in Release 7.0.....	2
1.2. Updates in Release 7.1.....	2
1.3. Updates in Release 7.2.....	3
1.4. Updates in Release 7.3.....	3
1.5. Updates in Release 7.4.....	3
1.6. Updates in Release 7.5.....	3
Chapter 2. Validated Platforms	4
2.1. Supported NVIDIA GPUs and Validated Server Platforms.....	4
2.2. Hypervisor Software Releases.....	5
2.3. Guest OS Support.....	7
2.3.1. Windows Guest OS Support.....	7
2.3.2. Linux Guest OS Support.....	8
2.4. NVIDIA CUDA Toolkit Version Support.....	8
Chapter 3. Known Product Limitations	10
3.1. Issues may occur with graphics-intensive OpenCL applications on vGPU types with limited frame buffer.....	10
3.2. vGPU profiles with 512 Mbytes or less of frame buffer support only 1 virtual display head on Windows 10.....	11
3.3. NVENC requires at least 1 Gbyte of frame buffer.....	11
3.4. VM running older NVIDIA vGPU drivers fails to initialize vGPU when booted.....	12
3.5. Virtual GPU fails to start if ECC is enabled.....	12
3.6. Single vGPU benchmark scores are lower than pass-through GPU.....	14
3.7. nvidia-smi fails to operate when all GPUs are assigned to GPU pass-through mode.....	15
3.8. Windows Aero is disabled on Citrix Virtual Apps and Desktops session using 3 or 4 monitors in 2560×1600 resolution.....	16
3.9. VMs configured with large memory fail to initialize vGPU when booted.....	17
3.10. vGPU host driver RPM upgrade fails.....	18
Chapter 4. Resolved Issues	19
Chapter 5. Security Updates	21
5.1. Since 7.2: Restricting Access to GPU Performance Counters.....	21
5.1.1. Windows: Restricting Access to GPU Performance Counters for One User by Using NVIDIA Control Panel.....	21
5.1.2. Windows: Restricting Access to GPU Performance Counters Across an Enterprise by Using a Registry Key.....	22
5.1.3. Linux Guest VMs and Hypervisor Host: Restricting Access to GPU Performance Counters.....	22
Chapter 6. Known Issues	24
6.1. 7.4 Only: In GPU pass-through mode, the NVIDIA graphics drivers fail to load.....	24
6.2. 7.0-7.3 Only: Users' view sessions may become corrupted after migration.....	24

6.3. NVIDIA Notification Icon prevents log off of Citrix Published Application user sessions.....	25
6.4. 7.0 Only: After a long stress test, a blue screen crash (BSOD) occurs.....	25
6.5. Vulkan applications crash in Windows 7 guest VMs configured with NVIDIA vGPU.....	26
6.6. Host core CPU utilization is higher than expected for moderate workloads.....	27
6.7. Frame capture while the interactive logon message is displayed returns blank screen.....	27
6.8. RDS sessions do not use the GPU with some Microsoft Windows Server releases.....	28
6.9. Cloned VMs configured with a vGPU type different than the type in the master image fail to start.....	29
6.10. 7.0 Only: Blue Screen crash during guest driver installation.....	29
6.11. 7.0 Only: Migrating a VM configured with vGPU can cause errors.....	30
6.12. Even when the scheduling policy is equal share, unequal GPU utilization is reported.....	31
6.13. When the scheduling policy is fixed share, GPU utilization is reported as higher than expected.....	32
6.14. nvidia-smi reports that vGPU migration is supported on all hypervisors.....	33
6.15. NVIDIA Control Panel Crashes in a VM connected to two 4K displays.....	33
6.16. vGPU guest VM driver not properly loaded on servers with more than 512 GB or 1 TB or more of system memory.....	34
6.17. Luxmark causes a segmentation fault on an unlicensed Linux client.....	35
6.18. Resolution is not updated after a VM acquires a license and is restarted.....	35
6.19. On Tesla P40, P6, and P4 GPUs, the default ECC setting prevents NVIDIA vGPU from starting.....	36
6.20. NVIDIA vGPU encoder and process utilization counters don't work with Windows Performance Counters.....	36
6.21. A segmentation fault in DBus code causes nvidia-gridd to exit on Red Hat Enterprise Linux and CentOS.....	37
6.22. No Manage License option available in NVIDIA X Server Settings by default.....	38
6.23. Licenses remain checked out when VMs are forcibly powered off.....	38
6.24. Memory exhaustion can occur with vGPU profiles that have 512 Mbytes or less of frame buffer.....	39
6.25. VM bug checks after the guest VM driver for Windows 10 RS2 is installed.....	41
6.26. On Citrix Hypervisor 7.0, VMs unexpectedly reboot and Citrix Hypervisor crashes or freezes.....	41
6.27. With no NVIDIA driver installed, Citrix Hypervisor misidentifies Tesla M10 cards.....	42
6.28. GNOME Display Manager (GDM) fails to start on Red Hat Enterprise Linux 7.2 and CentOS 7.0.....	43
6.29. Video goes blank when run in loop in Windows Media Player.....	43
6.30. Local VGA console is momentarily unblanked when Citrix Virtual Apps and Desktops changes resolution of the VM desktop.....	44
6.31. VM bugchecks on shutdown/restart when Citrix Virtual Apps and Desktops is installed and NVIDIA driver is uninstalled or upgraded.....	44
6.32. Application frame rate may drop when running Citrix Virtual Apps and Desktops at 2560×1600 resolution.....	45
6.33. Windows VM BSOD.....	46

6.34. Windows VM BSOD when upgrading NVIDIA drivers over a Citrix Virtual Apps and Desktops session.....	47
6.35. XenCenter does not allow vGPUs to be selected as a GPU type for Linux VMs.....	47
6.36. If X server is killed on a RHEL7 VM running vGPU, XenCenter console may not automatically switch to text console.....	48
6.37. Citrix Virtual Apps and Desktops shows only a black screen when connected to a vGPU VM49	

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 Hypervisor.



The most current version of the documentation for this release of NVIDIA vGPU software can be found online at [NVIDIA Virtual GPU Software Documentation](#).



Release 7.5 of NVIDIA vGPU software is the last release in the NVIDIA vGPU software 7 release branch.

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

Software	7.0	7.1	7.2	7.3	7.4	7.5
NVIDIA Virtual GPU Manager for the Citrix Hypervisor releases listed in Hypervisor Software Releases	410.68	410.91	410.107	410.122	410.137	410.137
NVIDIA Windows driver	411.81	412.16	412.31	412.38	412.45	412.47
NVIDIA Linux driver	410.71	410.92	410.107	410.122	410.137	410.141



Caution

If you install the wrong NVIDIA vGPU software packages for the version of Citrix Hypervisor 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 7.1 with guest VM drivers from release 7.0. 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 7.1 with guest VM drivers from release 6.2.

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 7.0

New Features in Release 7.0

- ▶ vGPU support for NVIDIA frame buffer capture metrics
- ▶ vGPU support for render capture metrics from the hypervisor and guest VMs
- ▶ Support for NVIDIA GPU Cloud (NGC) containers with NVIDIA vGPU software
- ▶ Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 7.0

- ▶ Support for Ubuntu 18.04 LTS as a guest OS
- ▶ Support for Red Hat Enterprise Linux 7.6 as a guest OS
- ▶ Support for CentOS 7.6 as a guest OS
- ▶ Support for Windows 10 Spring Creators Update (1803) as a guest OS

Feature Support Withdrawn in Release 7.0

- ▶ 32-bit Windows guest operating systems are no longer supported.

1.2. Updates in Release 7.1

New Features in Release 7.1

- ▶ Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 7.1

- ▶ Support for the Tesla T4 GPU

1.3. Updates in Release 7.2

New Features in Release 7.2

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

Hardware and Software Support Introduced in Release 7.2

- ▶ Support for Citrix Virtual Apps and Desktops version 7 1903

1.4. Updates in Release 7.3

New Features in Release 7.3

- ▶ Security updates
- ▶ Miscellaneous bug fixes

1.5. Updates in Release 7.4

New Features in Release 7.4

- ▶ Security updates
- ▶ Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 7.4

- ▶ Support for the following releases of Citrix Virtual Apps and Desktops:
 - ▶ 7 1909
 - ▶ 7 1906

1.6. Updates in Release 7.5

New Features in Release 7.5

- ▶ Fix for bug 2708778: In GPU pass-through mode, the NVIDIA graphics drivers fail to load with error code 43.

Chapter 2.

VALIDATED PLATFORMS

This release family of NVIDIA vGPU software provides support for several NVIDIA GPUs on validated server hardware platforms, Citrix Hypervisor hypervisor software versions, and guest operating systems. It also supports the version of NVIDIA CUDA Toolkit that is compatible with R410 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 Hypervisor, running on validated server hardware platforms:

- ▶ GPUs based on the NVIDIA Maxwell™ graphic architecture:
 - ▶ Tesla M6
 - ▶ Tesla M10
 - ▶ Tesla M60
- ▶ GPUs based on the NVIDIA Pascal™ architecture:
 - ▶ 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.)
- ▶ GPUs based on the NVIDIA Volta architecture:
 - ▶ 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.)
- ▶ GPUs based on the NVIDIA Turing architecture:
 - ▶ **Since 7.1:** Tesla T4

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
Citrix Receiver	Citrix Workspace App

The release number of Citrix Virtual Apps and Desktops releases **after 7.18** is in the form $7_{yy}mm$, where yy and mm denote the year and month that the product was released. For example, for a product released in August 2018, yy and mm are 18 and 08, respectively.

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

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

Supported Citrix Hypervisor Releases

This release family of NVIDIA vGPU software is supported on the Citrix Hypervisor releases listed in the table.



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

Cumulative update releases for a base release of Citrix Hypervisor 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
Citrix Hypervisor (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 Supported NVIDIA GPUs and Validated Server Platforms .
Citrix Hypervisor (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 Supported NVIDIA GPUs and Validated Server Platforms .
Citrix Hypervisor (XenServer) 7.1	RTM build and compatible cumulative update releases	Not supported on Tesla T4. XenMotion with vGPU is not supported.
Citrix Hypervisor (XenServer) 7.0	RTM build 125380 and compatible cumulative update releases	Not supported on Tesla T4. XenMotion with vGPU is not supported.

Supported Virtual Desktop Software Releases

This release supports **only** the virtual desktop software releases listed in the table. Except where otherwise stated, HDX 3D Pro mode is supported but not required.

Software	Releases Supported
Citrix Virtual Apps and Desktops	<p>Since 7.4: Versions 7 1909, 7 1906, 7 1903, and 7 1808</p> <p>7.2, 7.3 only: Versions 7 1903 and 7 1808</p> <p>7.0, 7.1 only: Version 7 1808</p>
XenApp/XenDesktop	Versions 7.18 and 7.17

Software	Releases Supported
	Version 7.15 in HDX 3D Pro mode (required to ensure that the NVIDIA GPU can be used unrestricted)

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.

NVIDIA vGPU software supports **only** 64-bit guest operating systems. No 32-bit guest operating systems are supported.

2.3.1. Windows Guest OS Support

NVIDIA vGPU software supports **only** the 64-bit Windows releases listed in the table as a guest OS on Citrix Hypervisor. The releases of Citrix Hypervisor 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.

XenMotion with vGPU is supported on supported Windows guest OS releases.

Guest OS	NVIDIA vGPU - Citrix Hypervisor Releases	Pass-Through GPU - Citrix Hypervisor Releases
Windows Server 2016 1607, 1709	7.6, 7.5, 7.1, 7.0	7.6, 7.5, 7.1, 7.0
Windows Server 2012 R2	7.6, 7.5, 7.1, 7.0	7.6, 7.5, 7.1, 7.0
Windows Server 2008 R2	7.6, 7.5, 7.1, 7.0	7.6, 7.5, 7.1, 7.0 Supported only on GPUs based on the Maxwell architecture
Windows 10 RTM (1507), November Update (1511), Anniversary Update (1607), Creators Update (1703), Fall Creators Update (1709), Spring Creators Update (1803) (64-bit)	7.6, 7.5, 7.1, 7.0	7.6, 7.5, 7.1, 7.0
Windows 8.1 Update (64-bit)	7.6, 7.5, 7.1, 7.0	7.6, 7.5, 7.1, 7.0

Guest OS	NVIDIA vGPU - Citrix Hypervisor Releases	Pass-Through GPU - Citrix Hypervisor Releases
Windows 8.1 (64-bit)	7.6, 7.5, 7.1, 7.0	-
Windows 8 (64-bit)	7.6, 7.5, 7.1, 7.0	-
Windows 7 (64-bit)	7.6, 7.5, 7.1, 7.0	7.6, 7.5, 7.1, 7.0 Supported only on GPUs based on the Maxwell architecture

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 Hypervisor. The releases of Citrix Hypervisor 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 Hypervisor Releases	Pass-Through GPU - Citrix Hypervisor Releases
Red Hat Enterprise Linux 7.0-7.6	7.6, 7.5, 7.1, 7.0	7.6, 7.5, 7.1, 7.0
CentOS 7.0-7.6	7.6, 7.5, 7.1, 7.0	7.6, 7.5, 7.1, 7.0
Ubuntu 18.04 LTS	7.6	7.6
Ubuntu 16.04 LTS	7.6, 7.5, 7.1, 7.0	7.6, 7.5, 7.1, 7.0
Ubuntu 14.04 LTS	7.6, 7.5, 7.1, 7.0	7.6, 7.5, 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 10.0.

For more information about NVIDIA CUDA Toolkit, see [CUDA Toolkit 10.0 Documentation](#).

2.5. XenMotion with vGPU Support

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

Supported GPUs:

- ▶ Tesla M6
- ▶ Tesla M10
- ▶ Tesla M60
- ▶ Tesla P4
- ▶ Tesla P6
- ▶ Tesla P40
- ▶ **Since 7.1:** Tesla T4

Supported Citrix 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. Issues may occur with graphics-intensive OpenCL applications on vGPU types with limited frame buffer

Description

Issues may occur when graphics-intensive OpenCL applications are used with vGPU types that have limited frame buffer. These issues occur when the applications demand more frame buffer than is allocated to the vGPU.

For example, these issues may occur with the Adobe Photoshop and LuxMark OpenCL Benchmark applications:

- ▶ When the image resolution and size are changed in Adobe Photoshop, a program error may occur or Photoshop may display a message about a problem with the graphics hardware and a suggestion to disable OpenCL.
- ▶ When the LuxMark OpenCL Benchmark application is run, XID error 31 may occur.

Workaround

For graphics-intensive OpenCL applications, use a vGPU type with more frame buffer.

3.2. 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.

3.3. 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.

3.4. 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 Hypervisor 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 Hypervisor 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.

3.5. 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 Hypervisor 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.

3.6. 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 ~]#
```

3.7. `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 Hypervisor dom0.

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

The screenshot displays the GPU configuration page for a VM. At the top, the VM name is 'xenserver-vgx-test (VM IPs 10.31.223.0-49, dom0 .96, OOB .97)'. The 'GPU' tab is selected. Below the tabs, the placement policy is 'Maximum density: put as many VMs as possible on the same GPU'. A list of six GPU passthrough entries is shown, each with a checked checkbox and a label: 'Test VM passthrough', 'Test VM passthrough', 'vgx-base-image-win7-64 passthrough', 'Test VM passthrough', 'Test VM passthrough', and 'Test VM passthrough'. To the right, the 'Allowed vGPU types' section lists five options, all with green checkmarks: 'Pass-through whole GPU', 'GRID K260Q vGPU (2 per GPU)', 'GRID K240Q vGPU (4 per GPU)', 'GRID K220Q vGPU (8 per GPU)', and 'GRID K200 vGPU (8 per GPU)'. Buttons for 'Select All', 'Clear All', and 'Edit Selected GPUs...' are located at the bottom of the list.

Resolution

N/A

3.8. Windows Aero is disabled on Citrix Virtual Apps and Desktops session using 3 or 4 monitors in 2560×1600 resolution

Description

Windows Aero may be disabled when Citrix Virtual Apps and Desktops 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](#).

3.9. 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 Hypervisor 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`

3.10. 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 7.0

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

Issues Resolved in Release 7.1

Bug ID	Summary and Description
<ul style="list-style-type: none">▶ 2425350▶ 200492702	<p>After a long stress test, a blue screen crash (BSOD) occurs</p> <p>After a long stress test, the GPU hangs, causing severe performance issues or a blue screen crash in the VM.</p>
200461544	<p>Blue Screen crash during guest driver installation</p> <p>During installation of the NVIDIA vGPU software graphics driver in guest VMs running some Microsoft Windows 10 releases, a blue screen crash occurs with the error VIDEO_TDR_FAILURE 116.</p>
200426078	<p>Migrating a VM configured with vGPU can cause errors</p> <p>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.</p>

Issues Resolved in Release 7.2

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

Issues Resolved in Release 7.3

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

Issues Resolved in Release 7.4

Bug ID	Summary and Description
200434909	7.0-7.3 Only: Users' view sessions may become corrupted after migration When a VM configured with vGPU under heavy load is migrated to another host, users' view sessions may become corrupted after the migration.

Issues Resolved in Release 7.5

Bug ID	Summary and Description
2708778	7.4 Only: In GPU pass-through mode, the NVIDIA graphics drivers fail to load In GPU pass-through mode with some hardware configurations, the NVIDIA graphics drivers fail to load with error code 43 and multiple assertion failures.

Chapter 5.

SECURITY UPDATES

5.1. Since 7.2: 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. 7.4 Only: In GPU pass-through mode, the NVIDIA graphics drivers fail to load

Description

In GPU pass-through mode with some hardware configurations, the NVIDIA graphics drivers fail to load with error code 43 and multiple assertion failures.

Status

Resolved in NVIDIA vGPU software 7.5

Ref.

2708778

6.2. 7.0-7.3 Only: Users' view sessions may become corrupted after migration

Description

When a VM configured with vGPU under heavy load is migrated to another host, users' view sessions may become corrupted after the migration.

Workaround

Restart the VM.

Status

Resolved in NVIDIA vGPU software 7.4

Ref. #

200434909

6.3. 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.

Resolution

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

Status

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

Ref. #

2206368

6.4. 7.0 Only: After a long stress test, a blue screen crash (BSOD) occurs

Description

After a long stress test, the GPU hangs, causing severe performance issues or a blue screen crash in the VM.

This issue affects only supported GPUs based on the NVIDIA Pascal architecture or the NVIDIA Volta architecture.

The `nvidia-smi -q` command returns the following information, in which power readings and clocks data are shown as `Unknown Error`.

```

Power Readings
  Power Management      : Supported
  Power Draw            : Unknown Error
  Power Limit           : 250.00 W
  Default Power Limit   : 250.00 W
  Enforced Power Limit  : 250.00 W
  Min Power Limit       : 125.00 W
  Max Power Limit       : 250.00 W
Clocks
  Graphics              : Unknown Error
  SM                    : Unknown Error
  Memory                : 3615 MHz
  Video                 : Unknown Error

```

Status

Resolved in release 7.1

Ref. #

- ▶ 2425350
- ▶ 200492702

6.5. Vulkan applications crash in Windows 7 guest VMs configured with NVIDIA vGPU

Description

In Windows 7 guest VMs configured with NVIDIA vGPU, applications developed with Vulkan APIs crash or throw errors when they are launched. Vulkan APIs require sparse texture support, but in Windows 7 guest VMs configured with NVIDIA vGPU, sparse textures are not enabled.

In Windows 10 guest VMs configured with NVIDIA vGPU, sparse textures are enabled and applications developed with Vulkan APIs run correctly in these VMs.

Status

Open

Ref. #

200381348

6.6. Host core CPU utilization is higher than expected for moderate workloads

Description

When GPU performance is being monitored, host core CPU utilization is higher than expected for moderate workloads. For example, host CPU utilization when only a small number of VMs are running is as high as when several times as many VMs are running.

Workaround

Disable monitoring of the following GPU performance statistics:

- ▶ vGPU engine usage by applications across multiple vGPUs
- ▶ Encoder session statistics
- ▶ Frame buffer capture (FBC) session statistics
- ▶ Statistics gathered by performance counters in guest VMs

Status

Open

Ref.

2414897

6.7. 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

6.8. RDS sessions do not use the GPU with some Microsoft Windows Server releases

Description

When some releases of Windows Server are used as a guest OS, Remote Desktop Services (RDS) sessions do not use the GPU. With these releases, the RDS sessions by default use the Microsoft Basic Render Driver instead of the GPU. This default setting enables 2D DirectX applications such as Microsoft Office to use software rendering, which can be more efficient than using the GPU for rendering. However, as a result, 3D applications that use DirectX are prevented from using the GPU.

Version

- ▶ Windows Server 2016
- ▶ Windows Server 2012

Solution

Change the local computer policy to use the hardware graphics adapter for all RDS sessions.

1. Choose **Local Computer Policy > Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Remote Session Environment** .
2. Set the **Use the hardware default graphics adapter for all Remote Desktop Services sessions** option.

6.9. 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

6.10. 7.0 Only: Blue Screen crash during guest driver installation

Description

During installation of the NVIDIA vGPU software graphics driver in guest VMs running some Microsoft Windows 10 releases, a blue screen crash occurs with the error `VIDEO_TDR_FAILURE 116`.

Version

This issue affects the following Microsoft Windows 10 releases:

- ▶ RTM (1507)
- ▶ November Update (1511)

Workaround

1. Identify the NVIDIA GPU for which the **Hardware Ids** property contains values that start with `PCI\VEN_10DE`.
 - a. Open **Device Manager** and expand **Display adapters**.
 - b. For each NVIDIA GPU listed under **Display adapters**, double-click the GPU and in the **Properties** window that opens, click the **Details** tab and select **Hardware Ids** in the **Property** list.
2. For the device that you identified in the previous step, display the value of the **Class Guid** property.

The value of this property is a string, for example, `4d36e968-e325-11ce-bfc1-08002be10318`.

3. Open the **Registry Editor** and navigate to `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\class-guid\`, where *class-guid* is the value of the **Class Guid** property that you displayed in the previous step.

Under *class-guid*, multiple adapters numbered as four-digit numbers starting from **0000** are listed, for example, **0000** and **0001**.

4. For each adapter listed, create the `EnableVGXFlipQueue` Windows registry key with type `REG_DWORD` and a value of **0**.
5. Install the NVIDIA vGPU software graphics driver.

Status

Resolved in NVIDIA vGPU software release 7.1.

Ref.

200461544

6.11. 7.0 Only: 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.

Resolved in NVIDIA vGPU software release 7.1.

6.12. 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

6.13. 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	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

6.14. `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

6.15. 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 Virtual Apps and Desktops 7.16

Status

Not an NVIDIA bug

Ref.

200393766

6.16. 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

6.17. 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

6.18. 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

6.19. 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

6.20. 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

6.21. 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

6.22. 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

6.23. 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

6.24. 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 Virtual Apps and Desktops 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 the Citrix Hypervisor `/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

6.25. 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 Hypervisor, 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

6.26. On Citrix Hypervisor 7.0, VMs unexpectedly reboot and Citrix Hypervisor crashes or freezes

Description

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

The event log in the Citrix Hypervisor `/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 Citrix Hypervisor 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

6.27. With no NVIDIA driver installed, Citrix Hypervisor misidentifies Tesla M10 cards

Description

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

Version

Citrix Hypervisor 6.5 and 7.0

Workaround

None

Status

Not an NVIDIA bug

Ref.

NVIDIA-420/1792341

6.28. 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

6.29. Video goes blank when run in loop in Windows Media Player

Description

When connected to a vGPU-enabled VM using Citrix Virtual Apps and Desktops, 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

6.30. Local VGA console is momentarily unblanked when Citrix Virtual Apps and Desktops changes resolution of the VM desktop

Description

When Citrix Virtual Apps and Desktops 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 Citrix Virtual Apps and Desktops 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

6.31. VM bugchecks on shutdown/restart when Citrix Virtual Apps and Desktops is installed and NVIDIA driver is uninstalled or upgraded.

Description

If the Citrix Virtual Apps and Desktops 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 Citrix Virtual Apps and Desktops.

Status

Open

Ref.

NVIDIA-295/200018125

6.32. Application frame rate may drop when running Citrix Virtual Apps and Desktops at 2560×1600 resolution.

Description

An application's rendering frame rate may drop when running Citrix Virtual Apps and Desktops 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

6.33. Windows VM BSOD

Description

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

The Citrix Hypervisor `/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

Citrix Hypervisor 6.2

Fix

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

Status

Closed

Ref.

NVIDIA-327/1632120

6.34. Windows VM BSOD when upgrading NVIDIA drivers over a Citrix Virtual Apps and Desktops session

Description

Windows VM bugchecks when NVIDIA guest drivers are upgraded over a Citrix Virtual Apps and Desktops session.

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

Fix

Upgrade Citrix Virtual Apps and Desktops to 7.6 Feature Pack 3

Status

Closed

Ref.

NVIDIA-370/200130780

6.35. 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 Citrix Hypervisor 6.5, and does include XenCenter integration.

Version

Affects the XenCenter integration with Citrix Hypervisor 6.5 only.

Resolved in the XenCenter integration with Citrix Hypervisor 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

6.36. 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

6.37. Citrix Virtual Apps and Desktops shows only a black screen when connected to a vGPU VM

Description

Citrix Virtual Apps and Desktops 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.