



Virtual GPU Software R595

Known Issues

Table of Contents

Chapter 1. Wayland with Ubuntu 24.04 does not work with multi-vGPU VMs.....	1
Chapter 2. Wayland sessions on Red Hat Enterprise Linux 10 with vGPUs on Blackwell-architecture GPUs are corrupted and unusable.....	2
Chapter 3. Display resolutions can't be changed in GNOME Settings with Wayland.....	3
Chapter 4. Application windows disappear after display resolution change in Wayland.....	4
Chapter 5. Remote desktop connections fail for vGPUs with 1 GB of frame buffer when used with Wayland.....	5
Chapter 6. With Red Hat Enterprise Linux 10 on VMware vSphere, the VM crashes when graphical.target is set.....	6
Chapter 7. The dmesg command reports non-fatal asserts and warnings.....	7
Chapter 8. vGPUs created with a vendor-specific VFIO framework on Ubuntu do not support Console VNC.....	8
Chapter 9. NVIDIA vGPU software does not support Ubuntu 22.04 with an upgraded HWE kernel.....	9
Chapter 10. Linux graphics driver installation from a .run file fails.....	10
Chapter 11. GPU device is unavailable on Windows 11 VMs with more than 1 TB of memory.....	11
Chapter 12. Default client configuration token folder missing on Windows client VMs.....	12
Chapter 13. Omnissa IDD, not the NVIDIA vGPU software graphics driver is driving the display in Windows 11 24H2 VMs.....	14
Chapter 14. vGPU VM fails to boot with error vmiop-display unable to reserve vgpu.....	15
Chapter 15. NVIDIA Control Panel is not available in multiuser environments.....	16
Chapter 16. NVIDIA Control Panel crashes if a user session is disconnected and reconnected.....	19
Chapter 17. VM assigned multiple fractional vGPUs from the same GPU hangs.....	20
Chapter 18. CUDA profilers cannot gather hardware metrics on NVIDIA vGPU.....	21
Chapter 19. NVIDIA vGPU software graphics driver for Windows sends a remote call to ngx.download.nvidia.com.....	22
Chapter 20. Multiple RDP session reconnections on Windows Server 2022 can consume all frame buffer.....	23
Chapter 21. VM with multiple legacy fractional vGPUs on the same GPU fails to boot.....	24

Chapter 22. On NVIDIA H100, creation of multiple compute instances after deletion of existing compute instances fails.....	25
Chapter 23. NLS client fails to acquire a license with the error The allowed time to process response has expired.....	26
Chapter 24. NVIDIA vGPU software graphics driver fails to load on KVM-based hypervisors.....	28
Chapter 25. With multiple active sessions, NVIDIA Control Panel incorrectly shows that the system is unlicensed.....	30
Chapter 26. VP9 and AV1 decoding with web browsers are not supported on Microsoft Windows Server 2019.....	31
Chapter 27. Ubuntu guest driver initialization fails with vGPUs and GPUs that support SR-IOV.....	32
Chapter 28. nvidia-smi ignores the second NVIDIA vGPU device added to a Microsoft Windows Server 2016 VM.....	33
Chapter 29. After an upgrade of the Linux graphics driver from an RPM package in a licensed VM, licensing fails.....	35
Chapter 30. After an upgrade of the Linux graphics driver from a Debian package, the driver is not loaded into the VM.....	36
Chapter 31. Desktop session freezes when a VM is migrated to or from a host running an NVIDIA vGPU software 14 release.....	37
Chapter 32. Application or vGPU VM crashes when multiple application instances are launched.....	38
Chapter 33. The reported NVENC frame rate is double the actual frame rate.....	39
Chapter 34. VM fails after a second vGPU is assigned to it.....	40
Chapter 35. Hypervisor host reboots when multiple cloned VMs are simultaneously powered on or migrated.....	41
Chapter 36. NVENC does not work with Teradici Cloud Access Software on Windows.....	42
Chapter 37. When a licensed client deployed by using VMware instant clone technology is destroyed, it does not return the license.....	43
Chapter 38. A licensed client might fail to acquire a license if a proxy is set.....	45
Chapter 39. Session connection fails with four 4K displays and NVENC enabled on a 2Q, 3Q, or 4Q vGPU.....	46
Chapter 40. Disconnected sessions cannot be reconnected or might be reconnected very slowly with NVWMI installed.....	48
Chapter 41. For some accounted processes, nvidia-smi incorrectly reports maximum memory use of 0 MiB.....	49

Chapter 42. Windows VM crashes during Custom (Advanced) driver upgrade..... 50

Chapter 43. VMs with vGPUs on GPUs based on the NVIDIA Ampere architecture fail to power on..... 51

Chapter 44. Linux VM hangs after vGPU migration to a host running a newer vGPU manager version..... 52

Chapter 45. NVIDIA A100 HGX 80GB vGPU names shown as Graphics Device by nvidia-smi..... 54

Chapter 46. Idle Teradici Cloud Access Software session disconnects from Linux VM..... 56

Chapter 47. GPU Operator doesn't support vGPU on GPUs based on architectures before NVIDIA Turing..... 57

Chapter 48. No virtual GPU types are listed in Citrix XenCenter..... 58

Chapter 49. Idle NVIDIA A100, NVIDIA A40, and NVIDIA A10 GPUs show 100% GPU utilization..... 60

Chapter 50. NVIDIA vGPU software graphics driver fails to load after upgrade on XenServer..... 62

Chapter 51. Windows guest VMs with vGPUs or GPUs with large BAR memory settings fail to boot to the desktop in UEFI mode..... 64

Chapter 52. Guest VM frame buffer listed by nvidia-smi for vGPUs on GPUs that support SRIOV is incorrect..... 66

Chapter 53. Migrated VM with more than two vGPUs crashes on destination host..... 67

Chapter 54. VMs fail to boot on RHV 4.4..... 68

Chapter 55. Driver upgrade in a Linux guest VM with multiple vGPUs might fail..... 70

Chapter 56. NVIDIA Control Panel fails to start if launched too soon from a VM without licensing information..... 71

Chapter 57. VNC client session goes blank and console VNC is corrupted when the guest driver is uninstalled..... 72

Chapter 59. Omnissa Horizon clients cannot connect to a Windows 10 2004 VM with multiple displays..... 74

Chapter 60. Suspend and resume between hosts running different versions of the vGPU manager fails..... 75

Chapter 61. On Linux, a VMware Horizon 7.12 session freezes after a switch to full screen 76

Chapter 62. On Linux, a VMware Horizon 7.12 session with two 4K displays freezes..... 77

Chapter 63. On Linux, the frame rate might drop to 1 after several minutes..... 78

Chapter 64. Frame buffer consumption grows with Omnissa Horizon over Blast Extreme... 80

Chapter 65. Citrix XenCenter displays misleading information about vGPU types.....	81
Chapter 66. Citrix Virtual Apps and Desktops connection freezes initially.....	83
Chapter 67. Microsoft DDA fails with some GPUs.....	84
Chapter 68. DWM crashes randomly occur in Windows VMs.....	86
Chapter 69. NVIDIA Control Panel fails to launch in a platform layer or published image....	87
Chapter 70. Remote desktop session freezes with assertion failure and XID error 43 after migration.....	88
Chapter 71. Citrix Virtual Apps and Desktops session freezes when the desktop is unlocked.....	89
Chapter 72. NVIDIA vGPU software graphics driver fails after Linux kernel upgrade with DKMS enabled.....	90
Chapter 73. Red Hat Enterprise Linux and CentOS 6 VMs hang during driver installation..	92
Chapter 74. On XenServer, all vGPUs in a VM must be of the same type.....	93
Chapter 75. Console VGA cannot be disabled.....	94
Chapter 76. Console VNC is unusable with Xorg on multiple vGPUs in a VM.....	95
Chapter 77. Blue screen crash occurs or no devices are found after VM reset.....	96
Chapter 78. Tesla T4 is enumerated as 32 separate GPUs by VMware vSphere ESXi.....	97
Chapter 79. Users' sessions may freeze during vMotion migration of VMs configured with vGPU.....	98
Chapter 80. Migration of VMs configured with vGPU stops before the migration is complete.....	99
Chapter 81. ECC memory settings for a vGPU cannot be changed by using NVIDIA X Server Settings.....	100
Chapter 82. Changes to ECC memory settings for a Linux vGPU VM by nvidia-smi might be ignored.....	101
Chapter 83. Incorrect GPU type shown for Quadro RTX 8000 GPUs in Citrix XenCenter....	103
Chapter 84. NVIDIA Notification Icon prevents log off of Citrix Published Application user sessions.....	104
Chapter 85. Host core CPU utilization is higher than expected for moderate workloads...	105
Chapter 86. H.264 encoder falls back to software encoding on 1Q vGPUs with a 4K display.....	106
Chapter 87. H.264 encoder falls back to software encoding on 2Q vGPUs with 3 or more 4K displays.....	107

Chapter 88. Frame capture while the interactive logon message is displayed returns blank screen.....	108
Chapter 89. RDS sessions do not use the GPU with Microsoft Windows Server as guest OS.....	109
Chapter 90. VMware vMotion fails gracefully under heavy load.....	110
Chapter 91. View session freezes intermittently after a Linux VM acquires a license.....	111
Chapter 92. When the scheduling policy is fixed share, GPU utilization is reported as higher than expected.....	112
Chapter 93. Benign warnings during Virtual GPU Manager installation or uninstallation after hypervisor upgrade.....	114
Chapter 94. Benign not in gzip format messages during Virtual GPU Manager installation or uninstallation.....	116
Chapter 95. License is not acquired in Windows VMs.....	117
Chapter 96. nvidia-smi reports that vGPU migration is supported on all hypervisors.....	118
Chapter 97. NVIDIA Control Panel Crashes in a VM connected to two 4K displays.....	119
Chapter 98. Hot plugging and unplugging vCPUs causes a blue-screen crash in Windows VMs.....	120
Chapter 99. GPU resources not available error during VMware instant clone provisioning..	121
Chapter 100. Module load failed during VIB downgrade from R390 to R384.....	123
Chapter 101. vGPU guest VM driver not properly loaded on servers with more than 512 GB or 1 TB or more of system memory.....	124
Chapter 102. Luxmark causes a segmentation fault on an unlicensed Linux client.....	126
Chapter 103. On Linux, 3D applications run slowly when windows are dragged.....	127
Chapter 104. A segmentation fault in DBus code causes nvidia-gridd to exit on Red Hat Enterprise Linux and CentOS.....	128
Chapter 105. No Manage License option available in NVIDIA X Server Settings by default.	129
Chapter 106. Licenses remain checked out when VMs are forcibly powered off.....	130
Chapter 107. Memory exhaustion can occur with vGPU profiles that have 512 Mbytes or less of frame buffer.....	131
Chapter 108. VM bug checks after the guest VM driver for Windows 10 RS2 is installed..	133
Chapter 109. On XenServer 7.0, VMs unexpectedly reboot and XenServer crashes or freezes.....	135
Chapter 110. With no NVIDIA driver installed, XenServer misidentifies Tesla M10 cards....	136
Chapter 111. vGPU VM fails to boot in ESXi if the graphics type is Shared.....	137

Chapter 112. GNOME Display Manager (GDM) fails to start on Red Hat Enterprise Linux 7.2 and CentOS 7.0.....	138
Chapter 113. NVIDIA Control Panel fails to start and reports that “you are not currently using a display that is attached to an Nvidia GPU”.....	139
Chapter 114. VM configured with more than one vGPU fails to initialize vGPU when booted.....	140
Chapter 115. A VM configured with both a vGPU and a passthrough GPU fails to start the passthrough GPU.....	141
Chapter 116. vGPU allocation policy fails when multiple VMs are started simultaneously..	142
Chapter 117. Before Horizon agent is installed inside a VM, the Start menu’s sleep option is available.....	143
Chapter 118. vGPU-enabled VMs fail to start, nvidia-smi fails when VMs are configured with too high a proportion of the server’s memory.....	144
Chapter 119. On reset or restart VMs fail to start with the error VMIOP: no graphics device is available for vGPU.....	145
Chapter 120. nvidia-smi shows high GPU utilization for vGPU VMs with active Horizon sessions.....	146
Chapter 121. Video goes blank when run in loop in Windows Media Player.....	147
Chapter 122. Local VGA console is momentarily unblanked when Citrix Virtual Apps and Desktops changes resolution of the VM desktop.....	148
Chapter 123. VM bugchecks on shutdown/restart when Citrix Virtual Apps and Desktops is installed and NVIDIA driver is uninstalled or upgraded.....	149
Chapter 124. Application frame rate may drop when running Citrix Virtual Apps and Desktops at 2560×1600 resolution.....	150
Chapter 125. Windows VM BSOD.....	151
Chapter 126. Windows VM BSOD when upgrading NVIDIA drivers over a Citrix Virtual Apps and Desktops session.....	152
Chapter 127. XenCenter does not allow vGPUs to be selected as a GPU type for Linux VMs.....	153
Chapter 128. If X server is killed on a RHEL7 VM running vGPU, XenCenter console may not automatically switch to text console.....	154
Chapter 129. Citrix Virtual Apps and Desktops shows only a black screen when connected to a vGPU VM.....	155

Chapter 1. Wayland with Ubuntu 24.04 does not work with multi-vGPU VMs

Description

If the Wayland display server protocol is used with the Ubuntu 24.04 guest OS in a VM to which multiple vGPUs are assigned, remote users cannot connect to the VM over RDP. When this issue occurs, the `nvidia-smi` command does not list any processes, such as the `gnome-shell`, `Xwayland`, and `gnome-remote-desktop-daemon` processes.

This issue is related to the version of the Mutter compositor introduced in Ubuntu 24.04.3. It occurs only when Wayland is used with an affected Ubuntu 24.04 release in a VM to which multiple vGPUs are assigned. The following configurations behave as expected:

- > Wayland in a VM to which only one vGPU is assigned
- > The X11 window system in a VM to which multiple vGPUs are assigned

Version

This issue is introduced in Ubuntu 24.04.3 guest OS. It does not affect earlier Ubuntu 24.04 releases.

Status

Not an NVIDIA bug

Ref.

5742102

Chapter 2. Wayland sessions on Red Hat Enterprise Linux 10 with vGPUs on Blackwell-architecture GPUs are corrupted and unusable

Description

Wayland sessions on Red Hat Enterprise Linux 10 with vGPUs on Blackwell-architecture GPUs are corrupted and unusable.

Version

Red Hat Enterprise Linux guest OS releases 10.1 and 10.0.

Status

Resolved in the following GNOME Remote Desktop packages for the affected Red Hat Enterprise Linux guest OS releases:

- > **Red Hat Enterprise Linux 10.1:** `gnome-remote-desktop-47.3-3.e110_1.x86_64.rpm`
- > **Red Hat Enterprise Linux 10.0:** `gnome-remote-desktop-47.3-3.e110_0.x86_64.rpm`

Resolution

Upgrade the installed GNOME Remote Desktop packages to their latest versions.

Ref.

5830923

Chapter 3. Display resolutions can't be changed in GNOME Settings with Wayland

Description

When the Wayland display server protocol is used with NVIDIA vGPU, some changes to display resolutions can't be applied in **GNOME Settings**. When the changes can't be applied, the **Apply** button is grayed out and the **GNOME Settings** window displays the following error message:

```
Changes Cannot be Applied.  
This could be due to hardware limitations.
```

Workaround

Try one of the following workarounds:

- > In **GNOME Settings**, disable the monitor first, select the resolution that you want, and re-enable the monitor.
- > Use the [gdctl](#) - GNOME Display Controller command to change the display resolution.

Status

Not a bug

Ref.

5526137

Chapter 4. Application windows disappear after display resolution change in Wayland

Description

After the display resolution is changed in the Wayland display server protocol, some application windows might disappear.

Version

This issue affects only Red Hat Enterprise Linux 10 guest OS releases.

Workaround

Close each application by context-clicking its icon in the dock and restart the application.

Status

Open

Ref.

5897794

Chapter 5. Remote desktop connections fail for vGPUs with 1 GB of frame buffer when used with Wayland

Description

Insufficient video memory might cause connections over Remote Desktop Protocol (RDP) to fail for vGPUs with 1 GB of frame buffer when used with the Wayland display server protocol. When this issue occurs, either the RDP session is not connected at all or is connected but with a blank screen.

Workaround

Use a vGPU type with at least 2 GB of frame buffer.

Status

Open

Ref.

5913053

Chapter 6. With Red Hat Enterprise Linux 10 on VMware vSphere, the VM crashes when `graphical.target` is set

Description

When a Red Hat Enterprise Linux 10 guest OS release is used on VMware vSphere, the VM crashes when `graphical.target` is set. Specifically, this issue occurs when the VM is booted in `multiuser.target` mode and the command `sudo systemctl isolate graphical.target` is run.

Version

This issue affects only Red Hat Enterprise Linux 10 guest OS releases on VMware vSphere.

Workaround

In VMware vSphere advanced parameters, set `svga.present` to `FALSE`.

Status

Not an NVIDIA bug

Ref.

5882844

Chapter 7. The `dmesg` command reports non-fatal asserts and warnings

Description

During the loading of the NVIDIA vGPU software graphics driver the `dmesg` command reports non-fatal asserts and warnings. When this issue occurs, the following warning is written to the Linux kernel log file in the guest VM:

```
WARNING: at nvidia-drm/nvidia-drm-crtc.h:322 __nv_drm_handle_flip_event [nvidia_drm]
Modules linked in: nvidia_drm(OE) hid_generic nvidia_modeset(OE) usbhid hid
nvidia(OE)
video wmi ahci drm_ttm_helper psmouse libahci ttm virtio_rng
```

Status

Open

Ref.

5381783

Chapter 8. vGPUs created with a vendor-specific VFIO framework on Ubuntu do not support Console VNC

Description

vGPUs created with a vendor-specific VFIO framework on Ubuntu do not support Console VNC. This issue affects vGPUs that support SR-IOV on releases of Ubuntu that use a vendor-specific VFIO framework, such as Ubuntu 24.04. This issue occurs because `libvirt` on releases of Ubuntu that use a vendor-specific VFIO framework does not support Console VNC. Resolution of this issue would entail changes to `libvirt`, not the NVIDIA Virtual GPU Manager.

Status

Not an NVIDIA bug

Ref.

3856353

Chapter 9. NVIDIA vGPU software does not support Ubuntu 22.04 with an upgraded HWE kernel

Description

NVIDIA vGPU software does not support the use of the Ubuntu 22.04 distribution with the Ubuntu 24.04 hardware enablement (HWE) kernel (kernel version 6.8.x). If Ubuntu 22.04 is used with an upgraded HWE kernel, any attempt to create an `mdev` device file fails.

Workaround

- > If you want to use the Ubuntu 22.04 distribution, downgrade the HWE kernel to the Ubuntu 22.04 HWE kernel.
- > If you want to use the Ubuntu 24.04 HWE kernel, upgrade to the Ubuntu 24.04 distribution.

Status

Not an NVIDIA bug.

Ref.

4864447

Chapter 10. Linux graphics driver installation from a `.run` file fails

Description

Installation of the NVIDIA vGPU software graphics driver for Linux from a `.run` file fails while the installer is building the kernel modules. This issue occurs if the `/tmp` file system has insufficient space for the temporary files that are created when the kernel modules are being built.

Workaround

Set the `TMPDIR` environment variable to a disk location that has more than 2GB of free space and try again

Status

Not a bug

Ref.

5352892

Chapter 11. GPU device is unavailable on Windows 11 VMs with more than 1 TB of memory

Description

On Windows 11 VMs with more than 1 TB of system memory, GPU device unavailable errors (Error 43) occur. This issue affects NVIDIA vGPU and GPU pass through deployments.

Version

This issue affects Windows 11 guest VMS.

Workaround

Limit the amount of system memory assigned to the VM to less than 1 TB.

Status

Open

Ref.

5115698

Chapter 12. Default client configuration token folder missing on Windows client VMs

Description

In air-gapped environments where root certificates are not available on the host machine, timestamps cannot be verified. As a result, the NVIDIA vGPU software graphics driver fails to create the default client configuration token folder on Windows (%SystemDrive%\Program Files\NVIDIA Corporation\vGPU Licensing\ClientConfigToken). If the folder is created manually and the client configuration token is copied there, the client fails to obtain a license. Typically, root certificates are imported by Windows updates from the Microsoft Trusted Root Program.

Workaround

Determine whether the NVIDIA Authenticode signature certificate and the timestamp signature certificate are installed and, if not, install them.

To determine whether the root NVIDIA Authenticode signature certificate is installed:

1. Context-click the file and click the **Digital Signatures** tab.
2. In the **Signature list**, select the NVIDIA certificate and click **Details**.
3. Click **View Certificate**, then click **Certification Path**.

The root certificate that is needed appears at the top of the certification path.

4. Run the `certmgr.msc` command and in the **certmgr** window that opens, expand **Trusted Root Certification Authorities** and click **Certificates** to see whether the certificate that you identified in the previous step is installed.

To determine whether the root timestamp signature certificate is installed:

1. Context-click the file and click the **Digital Signatures** tab.
2. In the **Signature list**, select the NVIDIA certificate and click **Details**.

3. In the **Countersignatures** section, click the timestamp authority, for example, Digicert or Entrust, then click **Details** below the countersignature section.
4. Click **View Certificate**, then click **Certification Path**.

The root certificate that is needed appears at the top of the certification path.

5. Run the `certmgr.msc` command and in the **certmgr** window that opens, expand **Trusted Root Certification Authorities** and click **Certificates** to see whether the certificate that you identified in the previous step is installed.

Root certificates for both Digicert and Entrust are required for timestamping and can be downloaded from the following websites:

- > [DigiCert Trusted Root Authority Certificates](#)
- > [Entrust Téléchargements de certificats racine](#)

Status

Not an NVIDIA bug

Ref.

4684895

Chapter 13. Omnissa IDD, not the NVIDIA vGPU software graphics driver is driving the display in Windows 11 24H2 VMs

Description

In Windows 11 24H2 guest VMs, the display is driven in Omnissa Horizon sessions by the Omnissa Horizon Indirect Display Driver (IDD) instead of the NVIDIA vGPU software graphics driver. This issue does not cause any visual corruption. However, OpenGL applications run at 30 fps instead of 60 fps, and pages for controlling the settings of multiple displays are missing from **NVIDIA Control Panel**.

Version

This issue affects **only** Omnissa Horizon with Windows 11 24H2 guest VMs. To find out which Omnissa Horizon versions support Windows 11 24H2, refer to [Omnissa Knowledge Base Article: Supported Windows 10 and Windows 11 Guest Operating Systems for Horizon Agent and Remote Experience, for Omnissa Horizon 8.x \(2006 and Later\) \(78714\)](#).

Workaround

Status

Not an NVIDIA bug

Ref.

4923798

Chapter 14. vGPU VM fails to boot with error `vmiop-display unable to reserve vgpu`

Description

A VM that is configured with NVIDIA vGPU on any NVIDIA RTX Ada graphics card, such as the NVIDIA RTX 6000 Ada and NVIDIA RTX 5000 Ada, fails to boot. When this issue occurs, the error message `vmiop-display unable to reserve vgpu` is written to the log files for the VM on the hypervisor host. This issue occurs because an issue with VMware vSphere Hypervisor (ESXi) prevents the hypervisor software from parsing the names of the virtual GPU types for these cards.

Version

This issue affects **only** the VMware vSphere Hypervisor (ESXi) 8.0U2 General Availability (GA) release. Other VMware vSphere Hypervisor (ESXi) releases that NVIDIA vGPU software supports are **not** affected.

Status

Not an NVIDIA bug

Ref.

4293546

Chapter 15. NVIDIA Control Panel is not available in multiuser environments

Description

After the NVIDIA vGPU software graphics driver for Windows is installed, the **NVIDIA Control Panel** app might be missing from the system. This issue typically occurs in the following situations:

- > Multiple users connect to virtual machines by using remote desktop applications such as Microsoft RDP, Omnissa Horizon, and Citrix Virtual Apps and Desktops.
- > VM instances are created by using Citrix Machine Creation Services (MCS) or VMware Instant Clone technology.
- > Roaming user desktop profiles are deployed.

This issue occurs because the **NVIDIA Control Panel** app is now distributed through the **Microsoft Store**. The **NVIDIA Control Panel** app might fail to be installed when the NVIDIA vGPU software graphics driver for Windows is installed if the **Microsoft Store** app is disabled, the system is not connected to the Internet, or installation of apps from the **Microsoft Store** is blocked by your system settings.

To determine whether the **NVIDIA Control Panel** app is installed on your system, use the **Windows Settings** app or the `Get-AppxPackage` Windows PowerShell command.

- > To use the **Windows Settings** app:
 1. From the Windows **Start** menu, choose **Settings > Apps > Apps & features**.
 2. In the **Apps & features** window, type `nvidia control panel` in the search box and confirm that the **NVIDIA Control Panel** app is found.
- > To use the `Get-AppxPackage` Windows PowerShell command:
 1. Run **Windows PowerShell** as Administrator.
 2. Determine whether the **NVIDIA Control Panel** app is installed for the current user.

```
PS C:\> Get-AppxPackage -Name NVIDIACorp.NVIDIAControlPanel
```

3. Determine whether the **NVIDIA Control Panel** app is installed for all users.

```
PS C:\> Get-AppxPackage -AllUsers -Name NVIDIACorp.NVIDIAControlPanel
```

This example shows that the **NVIDIA Control Panel** app is installed for the users Administrator, pliny, and trajan.

```
PS C:\> Get-AppxPackage -AllUsers -Name NVIDIACorp.NVIDIAControlPanel

Name                : NVIDIACorp.NVIDIAControlPanel
Publisher           : CN=D6816951-877F-493B-B4EE-41AB9419C326
Architecture       : X64
ResourceId          :
Version            : 8.1.964.0
PackageFullName    :
                   : NVIDIACorp.NVIDIAControlPanel_8.1.964.0_x64__56jybvy8sckqj
InstallLocation    : C:\Program Files\WindowsApps
                   : \NVIDIACorp.NVIDIAControlPanel_8.1.964.0_x64__56jybvy8sckqj
IsFramework        : False
PackageFamilyName  : NVIDIACorp.NVIDIAControlPanel_56jybvy8sckqj
PublisherId        : 56jybvy8sckqj
PackageUserInformation :
                   : {S-1-12-1-530092550-1307989247-1105462437-500 [Administrator]: Installed,
                   : S-1-12-1-530092550-1307989247-1105462437-1002 [pliny]: Installed,
                   : S-1-12-1-530092550-1307989247-1105462437-1003 [trajan]: Installed}
IsResourcePackage  : False
IsBundle           : False
IsDevelopmentMode  : False
NonRemovable       : False
IsPartiallyStaged  : False
SignatureKind      : Store
Status             : Ok
```

Preventing this Issue

If your system does not allow the installation apps from the **Microsoft Store**, download and run the standalone **NVIDIA Control Panel** installer that is available from NVIDIA Licensing Portal. For instructions, refer to [Virtual GPU Software User Guide](#).

If your system can allow the installation apps from the **Microsoft Store**, ensure that:

- > The Microsoft Store app is enabled.
- > Installation of Microsoft Store apps is not blocked by your system settings.
- > No local or group policies are set to block Microsoft Store apps.

Workaround

If the **NVIDIA Control Panel** app is missing, install it separately from the graphics driver by downloading and running the standalone **NVIDIA Control Panel** installer that is available from NVIDIA Licensing Portal. For instructions, refer to [Virtual GPU Software User Guide](#).

If the issue persists, contact NVIDIA Enterprise Support for further assistance.

Status

Open

Ref. #

3999308

Chapter 16. **NVIDIA Control Panel** crashes if a user session is disconnected and reconnected

Description

On all supported Windows Server guest OS releases, **NVIDIA Control Panel** crashes if a user session is disconnected and then reconnected while **NVIDIA Control Panel** is open.

Version

This issue affects all supported Windows Server guest OS releases.

Status

Open

Ref.

4086605

Chapter 17. VM assigned multiple fractional vGPUs from the same GPU hangs

Description

A VM that has been assigned multiple fractional vGPUs from the same physical GPU hangs or becomes inaccessible during installation of the NVIDIA vGPU software graphics driver in the VM. This issue affects only GPUs based on the NVIDIA Turing and NVIDIA Volta GPU architectures. This issue does not occur if the VM has been assigned multiple fractional vGPUs from different physical GPUs.

Version

This issue affects only GPUs based on the NVIDIA Turing and NVIDIA Volta GPU architectures.

Status

Open

Ref.

4020171

Chapter 18. CUDA profilers cannot gather hardware metrics on NVIDIA vGPU

Description

NVIDIA CUDA Toolkit profilers cannot gather hardware metrics on NVIDIA vGPU. This issue affects only traces that gather hardware metrics. Other traces are not affected by this issue and work normally.

Version

This issue affects NVIDIA vGPU software releases starting with 15.2.

Status

Open

Ref.

4041169

Chapter 19. NVIDIA vGPU software graphics driver for Windows sends a remote call to `ngx.download.nvidia.com`

Description

After the NVIDIA vGPU software graphics for windows has been installed in the guest VM, the driver sends a remote call to `ngx.download.nvidia.com` to download and install additional components. Such a remote call might be a security issue.

Workaround

Before running the NVIDIA vGPU software graphics driver installer, disable the remote call to `ngx.download.nvidia.com` by setting the following Windows registry key:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\NVIDIA Corporation\Global\NGXCore]
"EnableOTA"=dword:00000000
```



Note: If this Windows registry key is set to 1 or deleted, the remote call to `ngx.download.nvidia.com` is enabled again.

Status

Open

Ref.

4031840

Chapter 20. Multiple RDP session reconnections on Windows Server 2022 can consume all frame buffer

Description

Multiple RDP session reconnections in a Windows Server 2022 guest VM can consume all the frame buffer of a vGPU or physical GPU. When this issue occurs, users' screens becomes black, their sessions are disconnected but left intact, and they cannot log on again. The following error message is written to the event log on the hypervisor host:

```
The Desktop Window Manager process has exited.  
(Process exit code: 0xe0464645, Restart count: 1, Primary display device ID: )
```

Version

This issue affects only the Windows Server 2022 guest OS.

Workaround

Periodically restart the Windows Server 2022 guest VM to prevent all frame buffer from being consumed.

Status

Open

Ref.

3583766

Chapter 21. VM with multiple legacy fractional vGPUs on the same GPU fails to boot

Description

A VM to which multiple legacy fractional vGPUs on the same physical GPU are assigned fails to boot. A fractional vGPU is assigned only a fraction of the physical GPU's frame buffer. A legacy NVIDIA vGPU does not support single root I/O virtualization (SR-IOV). When this issue occurs, error messages similar to the following examples are written to the `vmware.log` file on the hypervisor host:

```
2022-11-23T09:01:06.643Z In(05) vmx - VMIOP: Registered device 0000:da:00.0
...
2022-11-23T09:01:06.715Z In(05) vmx - VMIOP: Failed to register device 0000:da:00.0
error = Failure
```

Status

Not an NVIDIA bug

Ref.

3879209

Chapter 22. On NVIDIA H100, creation of multiple compute instances after deletion of existing compute instances fails

Description

After compute instances are created and deleted on an NVIDIA H100 GPU, creation of multiple instances in a single `nvidia-smi` command fails. For example, the command `nvidia-smi mig -cci 0,1,2` fails with the following error message:

```
Unable to create a compute instance on GPU 0 GPU instance ID 0 using profile 0:  
Invalid Argument  
Failed to create compute instances: Invalid Argument
```

Workaround

Create each compute instance in a separate `nvidia-smi` command, for example:

```
$ nvidia-smi mig -cci 0  
$ nvidia-smi mig -cci 1  
$ nvidia-smi mig cci 2
```

Status

Open

Ref.

3829786

Chapter 23. NLS client fails to acquire a license with the error

The allowed time to process response has expired

Description

A licensed client of NVIDIA License System (NLS) fails to acquire a license with the error `The allowed time to process response has expired`. This error can affect clients of a Cloud License Service (CLS) instance or a Delegated License Service (DLS) instance.

This error occurs when the time difference between the system clocks on the client and the server that hosts the CLS or DLS instance is greater than 10 minutes. A common cause of this error is the failure of either the client or the server to adjust its system clock when daylight savings time begins or ends. The failure to acquire a license is expected to prevent clock windback from causing licensing errors.

Workaround

Ensure that system clock time of the client and any server that hosts a DLS instance match the current time in the time zone where they are located.

To prevent this error from occurring when daylight savings time begins or ends, enable the option to automatically adjust the system clock for daylight savings time:

- > **Windows:** Set the **Adjust for daylight saving time automatically** option.
- > **Linux:** Use the `hwclock` command.

Status

Not a bug

Ref. #

3859889

Chapter 24. NVIDIA vGPU software graphics driver fails to load on KVM-based hypervisors

Description

The NVIDIA vGPU software graphics driver fails to load on hypervisors based on Linux with KVM. This issue affects UEFI VMs configured with a vGPU or pass-through GPU that requires a large BAR address space. This issue does not affect VMs that are booted in legacy BIOS mode. The issue occurs because BAR resources are not mapped into the VM.

On a Windows VM, error code 12 is reported in **Device Manager** for the vGPU or pass-through GPU.

Workaround

1. In `virsh`, open for editing the XML document of the VM to which the vGPU or GPU is assigned.

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

The name of the VM to which the vGPU or GPU is assigned.

2. Declare the custom `libvirt` XML namespace that supports command-line pass through of QEMU arguments.

Declare this namespace by modifying the start tag of the top-level `domain` element in the first line of the XML document.

```
<domain type='kvm' xmlns:qemu='http://libvirt.org/schemas/domain/qemu/1.0'>
```

3. At the end of the XML document, between the `</devices>` end tag and the `</domain>` end tag, add the highlighted `qemu` elements.

These elements pass the QEMU arguments for mapping the required BAR resources into the VM, setting the MMIO aperture size to 262144. If necessary, replace the value of 262144 with the MMIO aperture size that your VM requires.

```
</devices>
  <qemu:commandline>
    <qemu:arg value='-fw_cfg' />
    <qemu:arg value='opt/ovmf/X-PciMmio64Mb,string=262144' />
  </qemu:commandline>
</domain>
```

4. Start the VM to which the vGPU or GPU is assigned.

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

vm-name

The name of the VM to which the vGPU or GPU is assigned.

Status

Not an NVIDIA bug

Ref.

200719557

Chapter 25. With multiple active sessions, **NVIDIA Control Panel** incorrectly shows that the system is unlicensed

Description

In an environment with multiple active desktop sessions, the **Manage License** page of **NVIDIA Control Panel** shows that a licensed system is unlicensed. However, the `nvidia-smi` command and the management interface of the NVIDIA vGPU software license server correctly show that the system is licensed. When an active session is disconnected and reconnected, the **NVIDIA Display Container** service crashes.

The **Manage License** page incorrectly shows that the system is unlicensed because of stale data in **NVIDIA Control Panel** in an environment with multiple sessions. The data is stale because **NVIDIA Control Panel** fails to get and update the settings for remote sessions when multiple sessions or no sessions are active in the VM. The **NVIDIA Display Container** service crashes when a session is reconnected because the session is not active at the moment of reconnection.

Status

Open

Ref.

3761243

Chapter 26. VP9 and AV1 decoding with web browsers are not supported on Microsoft Windows Server 2019

Description

VP9 and AV1 decoding with web browsers are not supported on Microsoft Windows Server 2019 and later supported releases. This issue occurs because starting with Windows Server 2019, the required codecs are not included with the OS and are not available through the **Microsoft Store** app. As a result, hardware decoding is not available for viewing YouTube videos or using collaboration tools such as Google Meet in a web browser.

Version

This issue affects Microsoft Windows Server releases starting with Windows Server 2019.

Status

Not an NVIDIA bug

Ref.

200756564

Chapter 27. Ubuntu guest driver initialization fails with vGPUs and GPUs that support SR-IOV

Description

Initialization of the NVIDIA vGPU software graphics driver fails for Ubuntu guest VMs configured with NVIDIA vGPUs or pass-through GPUs that support SR-IOV. An NVIDIA vGPU that supports SR-IOV resides on a physical GPU that supports SR-IOV with NVIDIA vGPU software, such as a GPU based on the NVIDIA Ampere architecture. This issue affects only guest VMs running a version of the Ubuntu OS with kernel version 5.13 or later, such as Ubuntu 20.04.

The failure of the driver to initialize is caused by an issue with extended Message Signaled Interrupts (MSI-X). When this issue occurs, the following error message is written to the kernel log file `/var/log/kern.log` in the guest VM:

```
RmInitAdapter failed!
```

Status

Open

Ref.

3660894

Chapter 28. `nvidia-smi` ignores the second NVIDIA vGPU device added to a Microsoft Windows Server 2016 VM

Description

After a second NVIDIA vGPU device is added to a Microsoft Windows Server 2016 VM, the device does not appear in the output from the `nvidia-smi` command. This issue occurs only if the VM is already running NVIDIA vGPU software for the existing NVIDIA vGPU device when the second device is added to the VM.

The `nvidia-smi` command cannot retrieve the guest driver version, license status, and accounting mode of the second NVIDIA vGPU device.

```
nvidia-smi vgpu --query
GPU 00000000:37:00.0
  Active vGPUs           : 1
  vGPU ID                : 3251695793
  VM ID                  : 3575923
  VM Name                 : SVR-Reg-W(P)-KuIn
  vGPU Name              : GRID V100D-32Q
  vGPU Type              : 185
  vGPU UUID              : 29097249-2359-11b2-8a5b-8e896866496b
  Guest Driver Version : 596.36
  License Status     : Licensed
  Accounting Mode    : Disabled
...
GPU 00000000:86:00.0
  Active vGPUs           : 1
  vGPU ID                : 3251695797
  VM ID                  : 3575923
  VM Name                 : SVR-Reg-W(P)-KuIn
  vGPU Name              : GRID V100D-32Q
  vGPU Type              : 185
  vGPU UUID              : 2926dd83-2359-11b2-8b13-5f22f0f74801
  Guest Driver Version : Not Available
  License Status     : N/A
  Accounting Mode    : N/A
```

Version

This issue affects only VMs that are running Microsoft Windows Server 2016 as a guest OS.

Workaround

To avoid this issue, configure the guest VM with both NVIDIA vGPU devices **before** installing the NVIDIA vGPU software graphics driver.

If you encounter this issue after the VM is configured, use one of the following workarounds:

- > Reinstall the NVIDIA vGPU software graphics driver.
- > Forcibly uninstall the Microsoft Basic Display Adapter and reboot the VM.
- > Upgrade the guest OS on the VM to Microsoft Windows Server 2019.

Status

Not an NVIDIA bug

Ref.

3562801

Chapter 29. After an upgrade of the Linux graphics driver from an RPM package in a licensed VM, licensing fails

Description

After the NVIDIA vGPU software graphics driver for Linux is upgraded from an RPM package in a licensed VM, licensing fails. The `nvidia-smi vgpu -q` command shows the driver version and license status as N/A. Restarting the `nvidia-gridd` service fails with a `Unit not found error`.

Workaround

Perform a clean installation of the NVIDIA vGPU software graphics driver for Linux from an RPM package.

1. Remove the currently installed driver.
2. Install the new version of the driver.

```
$ rpm -iv nvidia-linux-grid-595_595.71.05_amd64.rpm
```

Status

Open

Ref.

3512766

Chapter 30. After an upgrade of the Linux graphics driver from a Debian package, the driver is not loaded into the VM

Description

After the NVIDIA vGPU software graphics driver for Linux is upgraded from a Debian package, the driver is not loaded into the VM.

Workaround

Use one of the following workarounds to load the driver into the VM:

- > Reboot the VM.
- > Remove the `nvidia` module from the Linux kernel and reinsert it into the kernel.

1. Remove the `nvidia` module from the Linux kernel.

```
$ sudo rmmmod nvidia
```

2. Reinsert the `nvidia` module into the Linux kernel.

```
$ sudo modprobe nvidia
```

Status

Not a bug

Ref.

200748806

Chapter 31. Desktop session freezes when a VM is migrated to or from a host running an NVIDIA vGPU software 14 release

Description

When a VM configured with a Tesla V100 or Tesla T4 vGPU is migrated between a host running an NVIDIA vGPU software 14 release and a host running an NVIDIA vGPU software 13 release, the remote desktop session freezes. After the session freezes, the VM must be rebooted to recover the session. This issue occurs only when the NVIDIA hardware-based H.264/HEVC video encoder (NVENC) is enabled.

Version

The issue affects migrations between a host running an NVIDIA vGPU software 14 release and a host running an NVIDIA vGPU software 13 release.

Workaround

Disable NVENC.

Status

Open

Ref.

3512790

Chapter 32. Application or vGPU VM crashes when multiple application instances are launched

Description

When multiple application instances are launched on a legacy vGPU that is allocated only a fraction of the physical GPU's frame buffer, the application or VM to which the vGPU is assigned crashes but the guest VM remains accessible. A legacy NVIDIA vGPU does not support single root I/O virtualization (SR-IOV). This issue does **not** affect NVIDIA vGPUs that support SR-IOV.

When this issue occurs, the following error message is written to the `vmware.log` file:

```
vmiop_log: (0x0): vGPU message 7 failed
```

This issue occurs when the plugin for legacy NVIDIA vGPUs creates more BAR1 mappings than the hypervisor allows a VM to create. These mappings depend on the number and type of applications running in the VM.

Status

Open

Ref.

200680865

Chapter 33. The reported NVENC frame rate is double the actual frame rate

Description

The frame rate in frames per second (FPS) for the NVIDIA hardware-based H.264/HEVC video encoder (NVENC) reported by the `nvidia-smi encodersessions` command and NVWMI is double the actual frame rate. Only the reported frame rate is incorrect. The actual encoding of frames is **not** affected.

This issue affects only Windows VMs that are configured with NVIDIA vGPU.

Status

Open

Ref.

2997564

Chapter 34. VM fails after a second vGPU is assigned to it

Description

After a second vGPU is added to a VM and the VM is restarted, the VM fails. NVIDIA vGPU software supports up to a maximum of 16 vGPUs per VM on VMware vSphere Hypervisor (ESXi).

When this issue occurs, the following messages are written to the log file on the hypervisor host:

```
2021-09-27T17:11:42.303Z| vthread-2105551| | I005: vmiop_log: (0x0): Start restoring vGPU state ...
2021-09-27T17:11:43.465Z| vcpu-0| | E002: vmiop_log: (0x0): Deferred restore for RPCs cannot continue, since restore data was not saved
2021-09-27T17:11:43.465Z| vcpu-0| | E002: vmiop_log: (0x0): Deferred call for vmiopd restore_rpc_data failed at un-stun!
2021-09-27T17:11:43.465Z| vcpu-0| | E002: vmiop_log: (0x0): Failed to complete restore for deferred functions.
2021-09-27T18:44:27.034Z| vthread-2105550| | E002: vmiop_log: (0x0): vGPU message 1 failed, guest VGX version is already initialized...
2021-09-27T18:44:27.034Z| vthread-2105550| | E002: vmiop_log: (0x0): vGPU message 1 failed, result code: 0x40
...
2021-09-27T18:44:35.359Z| vthread-2105550| | I005: vmiop_log: (0x0): Guest driver unloaded!
```

Workaround

To avoid this issue, create your VMs in EFI mode.

If you encounter this issue with a VM that was created in legacy BIOS mode, shut down and restart the VM or power off the VM and power it on again.

Status

Not an NVIDIA bug

Ref.

3386681

Chapter 35. Hypervisor host reboots when multiple cloned VMs are simultaneously powered on or migrated

Description

When multiple cloned VMs are simultaneously powered on or migrated, the hypervisor host reboots. No crash dump is generated and no error messages related to the reboot are written to the log files `daemon.log` or `kern.log`.

Version

This issue affects only Citrix Hypervisor 8.2.

Status

Open

Ref.

200726850

Chapter 36. NVENC does not work with Teradici Cloud Access Software on Windows

Description

The NVIDIA hardware-based H.264/HEVC video encoder (NVENC) does not work with Teradici Cloud Access Software on Windows. This issue affects NVIDIA vGPU and GPU pass through deployments.

This issue occurs because the check that Teradici Cloud Access Software performs on the DLL signer name is case sensitive and NVIDIA recently changed the case of the company name in the signature certificate.

Status

Not an NVIDIA bug

This issue is resolved in the latest 21.07 and 21.03 Teradici Cloud Access Software releases.

Ref.

200749065

Chapter 37. When a licensed client deployed by using VMware instant clone technology is destroyed, it does not return the license

Description

When a user logs out of a VM deployed by using Omnisca Horizon instant clone technology, the VM is deleted and OS is not shut down cleanly. The NVIDIA vGPU software license that was being used by the VM is not returned to the license server, which could cause the license server to run out of licenses.

Workaround

Deploy the instant-clone desktop pool with the following options:

- > **Floating** user assignment
- > **All Machines Up-Front** provisioning

This configuration will allow the MAC address to be reused on the newly cloned VMs.

For more information, refer to the documentation for the version of Omnisca Horizon or VMware Horizon that you are using:

- > Omnisca Horizon 8: [Worksheet for Creating an Instant-Clone Desktop Pool](#)
- > VMware Horizon 7: [Worksheet for Creating an Instant-Clone Desktop Pool in Horizon Console](#)

Status

Not an NVIDIA bug

When a licensed client deployed by using VMware instant clone technology is destroyed, it does not return the license

Ref. #

200744338

Chapter 38. A licensed client might fail to acquire a license if a proxy is set

Description

If a proxy is set with a system environment variable such as `HTTP_PROXY` or `HTTPS_PROXY`, a licensed client might fail to acquire a license.

Workaround

Perform this workaround on each affected licensed client.

1. Add the address of the NVIDIA vGPU software license server to the system environment variable `NO_PROXY`.

The address must be specified exactly as it is specified in the client's license server settings either as a fully-qualified domain name or an IP address. If the `NO_PROXY` environment variable contains multiple entries, separate the entries with a comma (,).

If high availability is configured for the license server, add the addresses of the primary license server and the secondary license server to the system environment variable `NO_PROXY`.

2. Restart the NVIDIA driver service that runs the core NVIDIA vGPU software logic.
 - > On Windows, restart the **NVIDIA Display Container** service.
 - > On Linux, restart the `nvidia-gridd` service.

Status

Closed

Ref.

200704733

Chapter 39. Session connection fails with four 4K displays and NVENC enabled on a 2Q, 3Q, or 4Q vGPU

Description

Desktop session connections fail for a 2Q, 3Q, or 4Q vGPU that is configured with four 4K displays and for which the NVIDIA hardware-based H.264/HEVC video encoder (NVENC) is enabled. This issue affects only Teradici Cloud Access Software sessions on Linux guest VMs.

This issue is accompanied by the following error message:

```
This Desktop has no resources available or it has timed out
```

This issue is caused by insufficient frame buffer.

Workaround

Ensure that sufficient frame buffer is available for all the virtual displays that are connected to a vGPU by changing the configuration in one of the following ways:

- > Reducing the number of virtual displays. The number of 4K displays supported with NVENC enabled depends on the vGPU.

vGPU	4K Displays Supported with NVENC Enabled
2Q	1
3Q	2
4Q	3

- > Disabling NVENC. The number of 4K displays supported with NVENC disabled depends on the vGPU.

vGPU	4K Displays Supported with NVENC Disabled
2Q	2

vGPU	4K Displays Supported with NVENC Disabled
3Q	2
4Q	4

- Using a vGPU type with more frame buffer. Four 4K displays with NVENC enabled on any Q-series vGPU with at least 6144 MB of frame buffer are supported.

Status

Not an NVIDIA bug

Ref.

200701959

Chapter 40. Disconnected sessions cannot be reconnected or might be reconnected very slowly with NVWMI installed

Description

Disconnected sessions cannot be reconnected or might be reconnected very slowly when the NVIDIA Enterprise Management Toolkit (NVWMI) is installed. This issue affects Citrix Virtual Apps and Desktops and Omnicast Horizon sessions on Windows guest VMs.

Workaround

Ensure that the NVWMI service is disabled.



Note: By default, NVWMI is disabled in the NVIDIA vGPU software graphics driver.

Status

Not a bug

Ref.

3262923

Chapter 41. For some accounted processes, `nvidia-smi` incorrectly reports maximum memory use of 0 MiB

Description

When the `--query-accounted-apps` option is used to list accounted compute processes, the `nvidia-smi` command incorrectly reports that the maximum memory use of some processes is 0 MiB. With the `--query-accounted-apps` option, `nvidia-smi` reports the maximum memory used by an application during its lifetime, and other information such as which GPU the application is running on and GR engine utilization.

Status

Open

Ref.

200647764

Chapter 42. Windows VM crashes during **Custom (Advanced)** driver upgrade

Description

When the NVIDIA vGPU software graphics driver in a Windows VM is upgraded with the **Custom (Advanced)** option selected, the VM crashes.



Status

Open

Ref. #

200700291

Chapter 43. VMs with vGPUs on GPUs based on the NVIDIA Ampere architecture fail to power on

Description

An otherwise correctly configured VMware vSphere ESXi 7.0 Update 2 server fails to boot VMs with vGPUs on GPUs based on the NVIDIA Ampere if the server being managed by a version of VMware vCenter Server older than 7.0.2. This version of VMware vCenter is released with ESXi 7.0 VMware vSphere Update 2.

When this issue occurs, the following error message is seen:

```
Insufficient resources. One or more devices (pciPassthru0) required by VM vm-name are not available on host host-name
```

Workaround

Use VMware vCenter Server 7.0.2 or a later compatible update

Status

Open

Chapter 44. Linux VM hangs after vGPU migration to a host running a newer vGPU manager version

Description

When a Linux VM configured with a Tesla V100 or Tesla T4 vGPU is migrated from a host that is running a vGPU manager 11 release before 11.6 to a host that is running a vGPU manager 13 release, the VM hangs. After the migration, the destination host and VM become unstable. When this issue occurs, XID error 31 is written to the log files on the destination hypervisor host.

Version

This issue affects migration from a host that is running a vGPU manager 11 release before 11.6 to a host that is running a vGPU manager 13 release.

Workaround

If the VM is configured with a Tesla T4 vGPU, perform the following sequence of steps before attempting the migration:

1. Upgrade the host that is running a vGPU manager 11 release to release 11.6 or a later vGPU manager 11 release.
2. Disconnect any remoting tool that is using NVENC.



Note: You cannot use this workaround for a VM that is configured with a Tesla V100 vGPU.

Status

Open

Ref. #

200691445

Chapter 45. NVIDIA A100 HGX 80GB vGPU names shown as Graphics Device by nvidia-smi

Description

The names of vGPUs that reside on the NVIDIA A100 80GB GPU are incorrectly shown as Graphics Device by the nvidia-smi command. The correct names indicate the vGPU type, for example, A100DX-40C.

```
$ nvidia-smi
Mon Jan 25 02:52:57 2021
+-----+
| NVIDIA-SMI 460.32.04      Driver Version: 460.32.04      CUDA Version: 11.2      |
+-----+
| GPU   Name                Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|=====+-----+=====+
|   0   Graphics Device      On         | 00000000:07:00.0 Off  |      0%      Default  |
| N/A   N/A    P0     N/A /  N/A | 6053MiB / 81915MiB |           Disabled   |
+-----+-----+-----+
|   1   Graphics Device      On         | 00000000:08:00.0 Off  |      0%      Default  |
| N/A   N/A    P0     N/A /  N/A | 6053MiB / 81915MiB |           Disabled   |
+-----+-----+-----+
+-----+
| Processes:
| GPU   GI    CI          PID    Type    Process name                      GPU Memory
|      ID    ID                                 |              Usage
|=====+=====+
| No running processes found
+-----+
```

Status

Open

Ref. #

200691204

Chapter 46. Idle Teradici Cloud Access Software session disconnects from Linux VM

Description

After a Teradici Cloud Access Software session has been idle for a short period of time, the session disconnects from the VM. When this issue occurs, the error messages `NVOS status 0x19` and `vGPU Message 21 failed` are written to the log files on the hypervisor host. This issue affects only Linux guest VMs.

Status

Open

Ref.

200689126

Chapter 47. GPU Operator doesn't support vGPU on GPUs based on architectures before NVIDIA Turing

Description

NVIDIA GPU Operator doesn't support vGPU deployments on GPUs based on architectures before the NVIDIA Turing™ architecture. This issue is caused by the omission of version information for the vGPU manager from the configuration information that GPU Operator requires. Without this information, GPU Operator does not deploy the NVIDIA driver container because the container cannot determine if the driver is compatible with the vGPU manager.

Status

Open

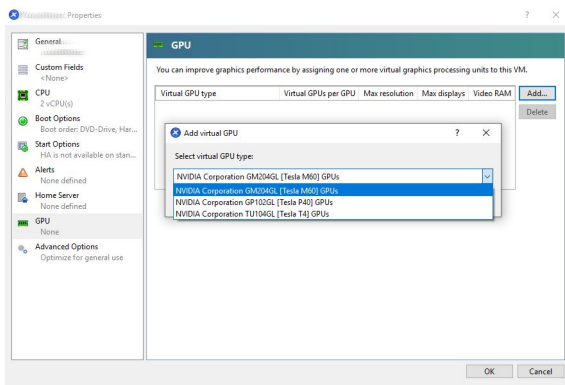
Ref.

3227576

Chapter 48. No virtual GPU types are listed in Citrix XenCenter

Description

No virtual GPU (vGPU) types are listed on the **Select virtual GPU type** drop-down list in **Citrix XenCenter**. However, the command `nvidia-smi vgpu -s` lists all supported vGPU types. This issue occurs because the driver version number of the Virtual GPU Manager contains multiple decimal points and is misinterpreted by Citrix XenCenter.



Version

Citrix Hypervisor 8.2

Resolution

Apply [Hotfix XS82E002 - For Citrix Hypervisor 8.2](#) from Citrix.

Status

Not an NVIDIA bug. Resolved by Citrix in hotfix XS82E002 for Citrix Hypervisor 8.2

Ref. #

200653755

Chapter 49. Idle NVIDIA A100, NVIDIA A40, and NVIDIA A10 GPUs show 100% GPU utilization

Description

The `nvidia-smi` command shows 100% GPU utilization for NVIDIA A100, NVIDIA A40, and NVIDIA A10 GPUs even if no vGPUs have been configured or no VMs are running. On Linux with KVM hypervisors, GPU is affected by this issue only if the `sriov-manage` script has **not** been run to enable the virtual function for the GPU in the `sysfs` file system.

```
[root@host ~]# nvidia-smi
Fri May 15 11:45:28 2026
+-----+
| NVIDIA-SMI 595.71.03      Driver Version: 595.71.03      CUDA Version:  13.2      |
+-----+-----+
| GPU   Name                Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|                                           MIG M.         |
+-----+-----+
| 0    A100-PCIE-40GB       On          | 00000000:5E:00.0 Off |             0        |
| N/A   50C    P0      97W / 250W | 0MiB / 40537MiB |    100%    Default  |
|                                           Disabled      |
+-----+-----+

Processes:
+-----+-----+
| GPU   GI   CI           PID   Type   Process name                      GPU Memory |
|   ID  ID  ID                |                 |           Usage |
+-----+-----+
| No running processes found |
+-----+-----+
```

Workaround

On Linux with KVM hypervisors, run the `sriov-manage` script to enable the virtual function for the GPU in the `sysfs` file system as explained in [Virtual GPU Software User Guide](#).

On VMware vSphere, boot any VMs that are configured with a vGPU that resides on the GPU.

After this workaround has been completed, the `nvidia-smi` command shows 0% GPU utilization for affected GPUs when they are idle.

```
root@host ~]# nvidia-smi
Fri May 15 11:47:38 2026

+-----+
| NVIDIA-SMI 595.71.03      Driver Version: 595.71.03      CUDA Version:  13.2   |
+-----+-----+
| GPU   Name                   Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|=====+=====+
|    0   A100-PCIE-40GB         On          | 00000000:5E:00:0 |      0%      Default |
| N/A   50C    P0     97W / 250W |  0MiB / 40537MiB |             Disabled |
+-----+-----+

Processes:
+-----+-----+
| GPU   GI   CI          PID   Type   Process name                      GPU Memory |
|      ID   ID                                     |            Usage |
+-----+-----+
| No running processes found |
+-----+-----+
```

Status

Open

Ref.

200605527

Chapter 50. NVIDIA vGPU software graphics driver fails to load after upgrade on XenServer

Description

NVIDIA vGPU software graphics driver fails to load after an upgrade of NVIDIA vGPU software from an 11.x release to a 12.x release on XenServer. This issue occurs because the hypervisor name in the RPM package name has changed from `xenserver` to `CitrixHypervisor`.

The attempted upgrade fails its dependency checks as follows:

```
[root@localhost ~]# rpm -Uvh NVIDIA-vGPU-CitrixHypervisor-8.2-595.71.03 .x86_64.rpm
error: Failed dependencies:
        NVIDIA-vGPU-xenserver conflicts with NVIDIA-vGPU-
CitrixHypervisor-1:8.2-595.71.03 .x86_64
```

Workaround

1. Determine the full name of the installed RPM package for the Virtual GPU Manager.

```
[root@localhost ~]# rpm -qa | grep NV
NVIDIA-vGPU-xenserver-8.2-580.159.01.x86_64
```

2. Remove the installed RPM package for the Virtual GPU Manager.

```
[root@localhost ~]# rpm -ev NVIDIA-vGPU-xenserver-8.2-580.159.01.x86_64
Preparing packages...
NVIDIA-vGPU-xenserver-1:8.2-580.159.01.x86_64
```

3. Perform a fresh installation of the latest RPM package for the Virtual GPU Manager.

```
[root@localhost ~]# rpm -ivh NVIDIA-vGPU-CitrixHypervisor-8.2-595.71.03 .x86_64.rpm
Preparing... ##### [100%]
Updating / installing...
 1:NVIDIA-vGPU-CitrixHypervisor-1:8.##### [100%]
```

4. Reboot the hypervisor host.

```
[root@localhost ~]# reboot
```

Status

Open

Ref. #

200682984

Chapter 51. Windows guest VMs with vGPUs or GPUs with large BAR memory settings fail to boot to the desktop in UEFI mode

Description

Windows guest VMs configured with vGPUs or physical GPUs that have large BAR memory settings fail to boot to the desktop in UEFI mode. The VM can be reached through its IP address but the desktop session is blank after the connection to the VM is established.


By default, Citrix Hypervisor creates a VM in UEFI boot mode and the boot mode of a VM cannot be changed after the VM is created.

Version

Citrix Hypervisor 8.2

Workaround

Delete the VM and re-create it, selecting **BIOS Boot** mode when you specify the installation method for the OS software on the new VM.



Status

Open

Ref.

200676622

Chapter 52. Guest VM frame buffer listed by `nvidia-smi` for vGPUs on GPUs that support SRIOV is incorrect

Description

The amount of frame buffer listed in a guest VM by the `nvidia-smi` command for vGPUs on GPUs that support Single Root I/O Virtualization (SR-IOV) is incorrect. Specifically, the amount of frame buffer listed is the amount of frame buffer allocated for the vGPU type minus the size of the VMMU segment (`vmmu_page_size`). Examples of GPUs that support SRIOV are GPUs based on the NVIDIA Ampere architecture, such as NVIDIA A100 PCIe 40GB or NVIDIA A100 HGX 40GB.

For example, frame buffer for -4C and -20C vGPU types is listed as follows:

- > For -4C vGPU types, frame buffer is listed as 3963 MB instead of 4096 MB.
- > For -20C vGPU types, frame buffer is listed as 20347 MB instead of 20480 MB.

Status

Open

Ref.

200524749

Chapter 53. Migrated VM with more than two vGPUs crashes on destination host

Description

When a VM that is configured with more than two vGPUs is migrated, the VM crashes on the destination host. When this issue occurs, the log file on the hypervisor host is flooded with NVOS status messages 0x23, 0x33, 0x5, and 0x21 and vGPU failure messages 4, 6, 7, 9, 12, 14, 19, 21, 23, 26, 32, 43, 54, and 56.

Status

Open

Ref.

200660221

Chapter 54. VMs fail to boot on RHV

4.4

Description

On RHV 4.4, VMs fail to boot with the error `Host doesn't support passthru of host PCI device`. This issue affects GPU pass through deployments with all supported GPUs and NVIDIA vGPU deployments with GPUs based on the NVIDIA Ampere architecture. This issue occurs because the `intel_iommu` parameter and the `nouveau.modeset` parameter are not set correctly.

Version

This issue affects RHV 4.4.

Workaround

Perform this workaround on the hypervisor host. This workaround requires root user privileges on the hypervisor host.

1. In a plain-text editor, edit the file `/boot/loader/entries/rhvh-4.4.1.1-0.20200722.0+1-4.18.0-193.13.2.el8_2.x86_64.conf` to add the following options to the boot options.

```
> nouveau.modeset=0
> intel_iommu=on
```



Note: Line breaks have been added to this example to enhance readability.

```
title rhvh-4.4.1.1-0.20200722.0 (4.18.0-193.13.2.el8_2.x86_64)
version 4.18.0-193.13.2.el8_2.x86_64
linux //rhvh-4.4.1.1-0.20200722.0+1/vmlinuz-4.18.0-193.13.2.el8_2.x86_64
initrd //rhvh-4.4.1.1-0.20200722.0+1/initramfs-4.18.0-193.13.2.el8_2.x86_64.img
options crashkernel=auto resume=/dev/mapper/rhvh00-swap \
rd.lvm.lv=rhvh00/rhvh-4.4.1.1-0.20200722.0+1 rd.lvm.lv=rhvh00/swap \
root=/dev/rhvh00/rhvh-4.4.1.1-0.20200722.0+1 \
boot=UUID=38ff2175-b761-403d-8a91-d7ec9f7ec2f7 rootflags=discard \
img.bootid=rhvh-4.4.1.1-0.20200722.0+1 intel_iommu=on nouveau.modeset=0
id rhel-20200825140238-4.18.0-193.13.2.el8_2.x86_64
grub_users $grub_users
```

```
grub_arg --unrestricted  
grub_class kernel
```

2. Reboot the hypervisor host machine.

Status

Not an NVIDIA bug

Ref.

200653675

Chapter 55. Driver upgrade in a Linux guest VM with multiple vGPUs might fail

Description

Upgrading the NVIDIA vGPU software graphics driver in a Linux guest VM with multiple vGPUs might fail. This issue occurs if the driver is upgraded by overinstalling the new release of the driver on the current release of the driver while the `nvidia-gridd` service is running in the VM.

Workaround

1. Stop the `nvidia-gridd` service.
2. Try again to upgrade the driver.

Status

Open

Ref.

200633548

Chapter 56. NVIDIA Control Panel fails to start if launched too soon from a VM without licensing information

Description

If NVIDIA licensing information is not configured on the system, any attempt to start **NVIDIA Control Panel** by right-clicking on the desktop within 30 seconds of the VM being started fails.

Workaround

Restart the VM and wait at least 30 seconds before trying to launch **NVIDIA Control Panel**.

Status

Open

Ref.

200623179

Chapter 57. VNC client session goes blank and console VNC is corrupted when the guest driver is uninstalled

Description

When the NVIDIA vGPU software graphics driver is uninstalled from a Windows 10 guest VM for which the boot mode is UEFI, the VNC client session goes blank and console VNC is corrupted.

This issue affects only Windows 10 guest VMs for which the boot mode is UEFI. It does **not** affect Windows 10 guest VMs for which the boot mode is BIOS.

Workaround

Reboot the VM by using Citrix XenCenter.

Status

Not an NVIDIA bug

Ref.

200609669

Chapter 58. Citrix Virtual Apps and Desktops session corruption occurs in the form of residual window borders

Description

When a window is dragged across the desktop in a Citrix Virtual Apps and Desktops session, corruption of the session in the form of residual window borders occurs.

Version

This issue affects only Citrix Virtual Apps and Desktops version 7 2003

Workaround

Use Citrix Virtual Apps and Desktops version 7 1912 or 2006.

Status

Not an NVIDIA bug

Ref.

200608675

Chapter 59. Omnissa Horizon clients cannot connect to a Windows 10 2004 VM with multiple displays

Description

Some Omnissa Horizon clients cannot connect to a Windows 10 2004 VM with multiple displays. When this issue occurs, the VM becomes unusable and clients cannot connect to the VM even if only a single display is connected to it.

This issue occurs because the desktop capture mechanism for the affected Omnissa Horizon clients is provided by NVIDIA® Frame Buffer Capture (NVFBC) and NVFBC is deprecated on Windows 10 starting with Windows 10 October 2019 Update. For more information, see [NVFBC Windows 10 Support Deprecation Technical Bulletin \(PDF\)](#).

Version

This issue affects only Windows 10 May 2020 Update (2004) guest VMs.

Workaround

Obtain a version of Omnissa Horizon for which the desktop capture mechanism is **not** provided by NVFBC.

Status

Not an NVIDIA bug

Ref.

200607827

Chapter 60. Suspend and resume between hosts running different versions of the vGPU manager fails

Description

Suspending a VM configured with vGPU on a host running one version of the vGPU manager and resuming the VM on a host running a version from an older main release branch fails. For example, suspending a VM on a host that is running the vGPU manager from release 20.1 and resuming the VM on a host running the vGPU manager from release 19.5 fails. When this issue occurs, an error is reported by your hypervisor management software.

- > Citrix XenCenter: `There were no servers available to complete the specified operation`
- > VMware vCenter Server: `One or more devices (pciPassthru0) required by VM vm-name are not available on host host-name`

Status

Not an NVIDIA bug

Ref.

200602087

Chapter 61. On Linux, a VMware Horizon 7.12 session freezes after a switch to full screen

Description

On a Linux VM configured with a -1Q vGPU, one 4K display, and VMware Horizon 7.12, the VMware Horizon session might become unresponsive after a switch from large screen (windowed) to full screen. When this issue occurs, the VM's log file contains the error message `Unable to set requested topology`.

Version

This issue affects deployments that use VMware Horizon 7.12.

Workaround

Use VMware Horizon 7.11.

Status

Open

Ref.

200617112

Chapter 62. On Linux, a VMware Horizon 7.12 session with two 4K displays freezes

Description

On a Linux VM configured with a -1Q vGPU, two 4K displays, and VMware Horizon 7.12, the VMware Horizon session might become unresponsive. When this issue occurs, the VM's log file contains the error message `Failed to setup capture session (error 8). Unable to allocate video memory.`

Version

This issue affects deployments that use VMware Horizon 7.12.

Workaround

Use VMware Horizon 7.11 or a vGPU with more frame buffer.

Status

Open

Ref.

200617081

Chapter 63. On Linux, the frame rate might drop to 1 after several minutes

Description

On Linux, the frame rate might drop to 1 frame per second (FPS) after NVIDIA vGPU software has been running for several minutes. Only some applications are affected, for example, `glxgears`. Other applications, such as Unigine Heaven, are not affected. This behavior occurs because Display Power Management Signaling (DPMS) for the Xorg server is enabled by default and the display is detected to be inactive even when the application is running. When DPMS is enabled, it enables power saving behavior of the display after several minutes of inactivity by setting the frame rate to 1 FPS.

Workaround

1. If necessary, stop the Xorg server.

```
# /etc/init.d/xorg stop
```

2. In a plain text editor, edit the `/etc/X11/xorg.conf` file to set the options to disable DPMS and disable the screen saver.

- a). In the `Monitor` section, set the `DPMS` option to `false`.

```
Option "DPMS" "false"
```

- b). At the end of the file, add a `ServerFlags` section that contains option to disable the screen saver.

```
Section "ServerFlags"  
    Option "BlankTime" "0"  
EndSection
```

- c). Save your changes to `/etc/X11/xorg.conf` file and quit the editor.

3. Start the Xorg server.

```
# etc/init.d/xorg start
```

Status

Open

Ref. #

200605900

Chapter 64. Frame buffer consumption grows with Omnissa Horizon over Blast Extreme

Description

When Omnissa Horizon is used with the Blast Extreme display protocol, frame buffer consumption increases over time after multiple disconnections from and reconnections to a VM. This issue occurs even if the VM is in an idle state and no graphics applications are running.

Workaround

Reboot the VM.

Status

Not an NVIDIA bug

Ref.

200602520

Workaround

For information about display configurations supported by Q-series and B-series vGPU types, see [Virtual GPU Software User Guide](#).

Status

Not an NVIDIA bug

Ref.

200556224

Chapter 66. Citrix Virtual Apps and Desktops connection freezes initially

Description

When -0B and -0Q vGPU types are used with Citrix Virtual Apps and Desktops version 7 1903 and later versions, the session freezes or a black screen is seen when the connection is first made. When this issue occurs, the error message `DXGI_ERROR_DEVICE_REMOVED` is displayed. The affected versions of Citrix Virtual Apps and Desktops use Microsoft DDAPI.

Version

Citrix Virtual Apps and Desktops version 7 1903 and later versions

Workaround

Disable the **use hardware encoding for video codec** Citrix policy, which is enabled by default.

Status

Open

Ref.

200494400

Chapter 67. Microsoft DDA fails with some GPUs

Description

Microsoft Discrete Device Assignment (DDA) fails with GPUs that have more than 16 GB of GPU memory. After the NVIDIA vGPU software graphics driver is installed in the guest VM, a second display device appears on the GPU and the driver prompts for a reboot. After the reboot, the device disappears and the Microsoft Hyper-V Video device appears.

This issue occurs because less memory-mapped input/output (MMIO) space is configured for the operating system than the device requires.

Workaround

Perform this workaround in a **Windows Power Shell** window on the hypervisor host.

Set the upper MMIO space to the amount that the device requires to allow all of the MMIO to be mapped. Upper MMIO space starts at approximately 64 GB in address space.

```
Set-VM -HighMemoryMappedIoSpace mmio-space -VMName vm-name
```

mmio-space

The amount of MMIO space that the device requires, appended with the appropriate unit of measurement, for example, **64GB** for 64 GB of MMIO space.

The required amount of MMIO space depends on the amount of BAR1 memory on the installed GPUs and the number of GPUs assigned to the VM as follows:

$$mmio-space = 2 \# gpu-bar1-memory \# assigned-gpus$$

gpu-bar1-memory

The amount of BAR1 memory on one of the installed GPUs. For example, in a server in which eight GPUs are installed and each GPU has 32 GB of BAR1 memory, *gpu-bar1-memory* is 32 GB.

assigned-gpus

The number of GPUs assigned to the VM.

vm-name

The name of the VM to which the GPU is assigned.

The following example sets the upper MMIO space to 64 GB for the VM named `mygpvm`, to which one GPU with 32 GB of BAR1 memory is assigned.

```
Set-VM -HighMemoryMappedIoSpace 64GB -VMName mygpvm
```

For more information, see [Deploy graphics devices using Discrete Device Assignment](#) on the Microsoft technical documentation site.

Status

Not an NVIDIA bug

Ref.

2812853

Chapter 68. DWM crashes randomly occur in Windows VMs

Description

Desktop Windows Manager (DWM) crashes randomly occur in Windows VMs, causing a blue-screen crash and the bug check `CRITICAL_PROCESS_DIED`. Computer Management shows problems with the primary display device.

Version

This issue affects Windows 10 1809, 1903 and 1909 VMs.

Status

Not an NVIDIA bug

Ref.

2730037

Chapter 69. NVIDIA Control Panel fails to launch in a platform layer or published image

Description

When NVIDIA vGPU software is used with Citrix App Layering and a platform layer for a vGPU is being configured, **NVIDIA Control Panel** might fail to launch. The driver might be working normally or it might fail with code 31 or code 43.

For more information and a workaround for this issue, see [NVIDIA Control Panel fails to launch in platform layer or published image](#) in the Citrix Support Knowledge Center.

Chapter 70. Remote desktop session freezes with assertion failure and XID error 43 after migration

Description

After multiple VMs configured with vGPU on a single hypervisor host are migrated simultaneously, the remote desktop session freezes with an assertion failure and XID error 43. This issue affects only GPUs that are based on the Volta GPU architecture. It does not occur if only a single VM is migrated.

When this error occurs, the following error messages are logged to the log file on the hypervisor host:

```
Jan  3 14:35:48 ch81-m1 vgpu-12[8050]: error: vmiop_log: NVOS status 0x1f
Jan  3 14:35:48 ch81-m1 vgpu-12[8050]: error: vmiop_log: Assertion Failed at
0x4b8cacf6:286
...
Jan  3 14:35:59 ch81-m1 vgpu-12[8050]: error: vmiop_log: (0x0): XID 43 detected on
physical_chid:0x174, guest_chid:0x14
```

Status

Open

Ref.

200581703

Chapter 71. Citrix Virtual Apps and Desktops session freezes when the desktop is unlocked

Description

When a Citrix Virtual Apps and Desktops session that is locked is unlocked by pressing **Ctrl+Alt+Del**, the session freezes. This issue affects only VMs that are running Microsoft Windows 10 1809 as a guest OS.

Version

Microsoft Windows 10 1809 guest OS

Workaround

Restart the VM.

Status

Not an NVIDIA bug

Ref.

2767012

Chapter 72. NVIDIA vGPU software graphics driver fails after Linux kernel upgrade with DKMS enabled

Description

After the Linux kernel is upgraded (for example by running `sudo apt full-upgrade`) with Dynamic Kernel Module Support (DKMS) enabled, the `nvidia-smi` command fails to run. If DKMS is enabled, an upgrade to the Linux kernel triggers a rebuild of the NVIDIA vGPU software graphics driver. The rebuild of the driver fails because the compiler version is incorrect. Any attempt to reinstall the driver fails because the kernel fails to build.

When the failure occurs, the following messages are displayed:

```
-> Installing DKMS kernel module:
  ERROR: Failed to run `/usr/sbin/dkms build -m nvidia -v 595.58.03 -k
5.3.0-28-generic`:
  Kernel preparation unnecessary for this kernel. Skipping...
  Building module:
  cleaning build area...
  'make' -j8 NV_EXCLUDE_BUILD_MODULES='' KERNEL_UNAME=5.3.0-28-generic
IGNORE_CC_MISMATCH='' modules... (bad exit status: 2)
  ERROR (dkms apport): binary package for nvidia: 595.58.03 not found
Error! Bad return status for module build on kernel: 5.3.0-28-generic
(x86_64)
  Consult /var/lib/dkms/nvidia/595.58.03/build/make.log for more information.
-> error.
  ERROR: Failed to install the kernel module through DKMS. No kernel module
was installed;
  please try installing again without DKMS, or check the DKMS logs for more
information.
  ERROR: Installation has failed. Please see the file '/var/log/nvidia-
installer.log' for details.
  You may find suggestions on fixing installation problems in the README
available on the Linux driver download page at www.nvidia.com.
```

Workaround

When installing the NVIDIA vGPU software graphics driver with DKMS enabled, use one of the following workarounds:

- > Before running the driver installer, install the `dkms` package, then run the driver installer with the `-dkms` option.
- > Run the driver installer with the `--no-cc-version-check` option.

Status

Not a bug.

Ref.

2836271

Chapter 73. Red Hat Enterprise Linux and CentOS 6 VMs hang during driver installation

Description

During installation of the NVIDIA vGPU software graphics driver in a Red Hat Enterprise Linux or CentOS 6 guest VM, a kernel panic occurs, and the VM hangs and cannot be rebooted. This issue is observed on older Linux kernels when the NVIDIA device is using message-signaled interrupts (MSIs).

Version

This issue affects the following guest OS releases:

- > Red Hat Enterprise Linux 6.6 and later compatible 6.x versions
- > CentOS 6.6 and later compatible 6.x versions

Workaround

1. Disable MSI in the guest VM to fall back to INTx interrupts by adding the following line to the file `/etc/modprobe.d/nvidia.conf`:

```
options nvidia NVreg_EnableMSI=0
```

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

2. Install the NVIDIA vGPU Software graphics driver in the guest VM.

Status

Closed

Ref.

200556896

Chapter 74. On XenServer, all vGPUs in a VM must be of the same type

Description

When a VM on XenServer is configured with multiple vGPUs, all vGPUs must be of the same type. XenServer does **not** support the assignment multiple vGPUs of different types to a single VM.

What happens when you try to assign a different type of vGPU than the vGPUs already assigned to a VM depends on the Citrix utility that you are using.

- > If you are using Citrix XenCenter, only vGPUs of the same type as the vGPUs already assigned are available for selection.
- > If you are using the `xe` command, the command fails and the following error message is displayed:

```
Cannot create a virtual GPU that is incompatible with the existing types on the VM.
```

Status

Not an NVIDIA bug.

Ref.

200568154

Chapter 75. Console VGA cannot be disabled

Description

In Citrix Hypervisor releases since 8.1, a VM's console VGA interface cannot be disabled by using the command **`xe vm-param-set uuid=vm-uuid platform:vgpu_extra_args="disable_vnc=1"`**.

This behavior is the result of the following changes that Citrix introduced in XenServer release 8.1:

- > The command `xe vm-param-set` for assigning plugin configuration parameters has been withdrawn.
- > Plugin parameters can no longer be set for vGPUs that were created by using Citrix XenCenter. They can be set only for vGPUs that were created from the command line by using the `xe` command.

Version

Since Citrix Hypervisor release 8.1

Workaround

Create the vGPU by using the `xe` command and specify plugin parameters for the group to which the vGPU belongs as explained in [Virtual GPU Software User Guide](#).

Status

Not a bug

Ref.

200569785

Chapter 76. Console VNC is unusable with Xorg on multiple vGPUs in a VM

Description

If a Linux VM is configured with multiple vGPUs and the Xorg display server is running on more than one vGPU in the VM, severe corruption in console VNC is observed.

Workaround

Ensure that the Xorg display server is running on only one vGPU in the VM.

Status

Not an NVIDIA bug.

Ref.

200568261

Chapter 77. Blue screen crash occurs or no devices are found after VM reset

Description

If a VM on Microsoft Windows Server with Hyper-V role is reset from the hypervisor host, a blue screen crash (BSOD) occurs on Windows VMs and the `nvidia-smi` command reports `No devices were found` on Linux VMs. This issue occurs only on Windows Server 2019 with Tesla T4 GPUs with SRIOV enabled, Quadro RTX 8000 passive GPUs, and Quadro RTX 6000 passive GPUs.

Workaround

Contact NVIDIA Enterprise Support for a workaround for this issue, referencing the knowledge base article *Workaround for Blue Screen Crashes On Hyper-V DDA With SRIOV-Enabled GPUs*. This article is available only to NVIDIA Enterprise Support personnel.

Status

Not an NVIDIA bug

Ref.

200567935

Chapter 78. Tesla T4 is enumerated as 32 separate GPUs by VMware vSphere ESXi

Description

Some servers, for example, the Dell R740, do not configure SR-IOV capability if the SR-IOV SBIOS setting is disabled on the server. If the SR-IOV SBIOS setting is disabled on such a server that is being used with the Tesla T4 GPU, VMware vSphere ESXi enumerates the Tesla T4 as 32 separate GPUs. In this state, you cannot use the GPU to configure a VM with NVIDIA vGPU or for GPU pass through.

Workaround

Ensure that the SR-IOV SBIOS setting is enabled on the server.

Status

Not an NVIDIA bug

A fix is available from VMware in VMware vSphere ESXi 7.0 Update 2.

Ref.

2697051

Chapter 79. Users' sessions may freeze during vMotion migration of VMs configured with vGPU

Description

When vMotion is used to migrate a VM configured with vGPU to another host, users' sessions may freeze for up to several seconds during the migration.

These factors may increase the length of time for which a session freezes:

- > Continuous use of the frame buffer by the workload, which typically occurs with workloads such as video streaming
- > A large amount of vGPU frame buffer
- > A large amount of system memory
- > Limited network bandwidth

Workaround

Administrators can mitigate the effects on end users by avoiding migration of VMs configured with vGPU during business hours or warning end users that migration is about to start and that they may experience session freezes.

End users experiencing this issue must wait for their sessions to resume when the migration is complete.

Status

Open

Ref.

2569578

Chapter 80. Migration of VMs configured with vGPU stops before the migration is complete

Description

When a VM configured with vGPU is migrated to another host, the migration stops before it is complete.

With Citrix XenCenter, the message `The vGPU is not compatible with any pGPU in the destination.` appears in the Citrix XenCenter management GUI when this issue occurs.

This issue occurs if the ECC memory configuration (enabled or disabled) on the source and destination hosts are different. The ECC memory configuration on both the source and destination hosts must be identical.

Workaround

Before attempting to migrate the VM again, ensure that the ECC memory configuration on both the source and destination hosts are identical.

Status

Not an NVIDIA bug

Ref.

200520027

Chapter 81. ECC memory settings for a vGPU cannot be changed by using **NVIDIA X Server Settings**

Description

The ECC memory settings for a vGPU cannot be changed from a Linux guest VM by using **NVIDIA X Server Settings**. After the ECC memory state has been changed on the **ECC Settings** page and the VM has been rebooted, the ECC memory state remains unchanged.

Workaround

Use the `nvidia-smi` command in the guest VM to enable or disable ECC memory for the vGPU as explained in [Virtual GPU Software User Guide](#).

If the ECC memory state remains unchanged even after you use the `nvidia-smi` command to change it, use the workaround in [Changes to ECC memory settings for a Linux vGPU VM by nvidia-smi might be ignored](#).

Status

Open

Ref.

200523086

Chapter 82. Changes to ECC memory settings for a Linux vGPU VM by `nvidia-smi` might be ignored

Description

After the ECC memory state for a Linux vGPU VM has been changed by using the `nvidia-smi` command and the VM has been rebooted, the ECC memory state might remain unchanged.

This issue occurs when multiple NVIDIA configuration files in the system cause the kernel module option for setting the ECC memory state `RMGuestECCState` in `/etc/modprobe.d/nvidia.conf` to be ignored.

When the `nvidia-smi` command is used to enable ECC memory, the file `/etc/modprobe.d/nvidia.conf` is created or updated to set the kernel module option `RMGuestECCState`. Another configuration file in `/etc/modprobe.d/` that contains the keyword `NVreg_RegistryDwordsPerDevice` might cause the kernel module option `RMGuestECCState` to be ignored.

Workaround

This workaround requires administrator privileges.

1. Move the entry containing the keyword `NVreg_RegistryDwordsPerDevice` from the other configuration file to `/etc/modprobe.d/nvidia.conf`.
2. Reboot the VM.

Status

Open

Ref. #

200505777

Chapter 83. Incorrect GPU type shown for Quadro RTX 8000 GPUs in Citrix XenCenter

Description

On the **GPU** tab in **Citrix XenCenter**, Quadro RTX 8000 GPUs are incorrectly shown as [QUADRO RTX 6000].

However, the virtual GPU types for a Quadro RTX 8000 GPU are listed correctly.

NVIDIA vGPU software does not support systems with a mixture of GPU types. Therefore, you can assume that any GPU in a system that you know to contain Quadro RTX 8000 GPUs that is shown as [QUADRO RTX 6000] is, in fact, a Quadro RTX 8000 GPU.

Status

Not an NVIDIA bug

Ref.

200492607

Chapter 84. 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

Chapter 85. 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

Chapter 86. H.264 encoder falls back to software encoding on 1Q vGPUs with a 4K display

Description

On 1Q vGPUs with a 4K display, a shortage of frame buffer causes the H.264 encoder to fall back to software encoding.

Workaround

Use a 2Q or larger virtual GPU type to provide more frame buffer for each vGPU.

Status

Open

Ref.

2422580

Chapter 87. H.264 encoder falls back to software encoding on 2Q vGPUs with 3 or more 4K displays

Description

On 2Q vGPUs with three or more 4K displays, a shortage of frame buffer causes the H.264 encoder to fall back to software encoding.

This issue affects only vGPUs assigned to VMs that are running a Linux guest OS.

Workaround

Use a 4Q or larger virtual GPU type to provide more frame buffer for each vGPU.

Status

Open

Ref.

200457177

Chapter 88. 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

On supported hypervisors **except** XenServer, press **Enter** or wait for the screen to update for NvFBC to capture the frame.

On XenServer, refer to [Black Screen at Logon with Nvidia HDX 3D Pro enabled](#) in the Citrix knowledge base.

Status

Not a bug

Ref.

2115733

Chapter 89. RDS sessions do not use the GPU with Microsoft Windows Server as guest OS

Description

When Windows Server is used as a guest OS, Remote Desktop Services (RDS) sessions do not use the GPU. By default, the RDS sessions 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

This issue affects all Windows Server releases that are supported as a guest OS.

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.

Chapter 90. VMware vMotion fails gracefully under heavy load

Description

Migrating a VM configured with vGPU fails gracefully if the VM is running an intensive workload.

The error stack in the task details on the vSphere web client contains the following error message:

```
The migration has exceeded the maximum switchover time of 100 second(s).
ESX has preemptively failed the migration to allow the VM to continue running on the
source.
To avoid this failure, either increase the maximum allowable switchover time or wait
until
the VM is performing a less intensive workload.
```

Workaround

Increase the maximum switchover time by increasing the `vmotion.maxSwitchoverSeconds` option from the default value of 100 seconds.

For more information, see [Broadcom Knowledge Base Article: vMotion or Storage vMotion of a VM fails with the error: The migration has exceeded the maximum switchover time of 100 second\(s\).](#)

Status

Not an NVIDIA bug

Ref.

200416700

Chapter 91. View session freezes intermittently after a Linux VM acquires a license

Description

In a Linux VM, the view session can sometimes freeze after the VM acquires a license.

Workaround

Resize the view session.

Status

Not an NVIDIA bug

Ref.

200426961

Chapter 92. 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

Chapter 93. Benign warnings during Virtual GPU Manager installation or uninstallation after hypervisor upgrade

Description

If the Virtual GPU Manager is installed or uninstalled after the hypervisor is upgraded, several warning messages about missing files are displayed.

The messages are similar to the following examples:

```
WARNING: Can't read module /lib/modules/4.4.77-1.el6.nutanix.20170830.124.x86_64/weak-updates/ixgbevf.ko: No such file or directory
WARNING: Can't read module /lib/modules/4.4.77-1.el6.nutanix.20170830.124.x86_64/weak-updates/i40evf.ko: No such file or directory
...
egrep: /lib/modules/4.4.77-1.el6.nutanix.20170830.124.x86_64//weak-updates/ixgbevf.ko: No such file or directory
egrep: /lib/modules/4.4.77-1.el6.nutanix.20170830.124.x86_64//weak-updates/i40evf.ko: No such file or directory
...
gzip: /boot/initramfs-4.4.77-1.el6.nutanix.20170830.124.x86_64.tmp: not in gzip format
WARNING: Can't read module /lib/modules/4.4.77-1.el6.nutanix.20170830.124.x86_64/weak-updates/ixgbevf.ko: No such file or directory
WARNING: Can't read module /lib/modules/4.4.77-1.el6.nutanix.20170830.124.x86_64/weak-updates/i40evf.ko: No such file or directory
...
```

After a hypervisor upgrade, files that are the targets of some symbolic links no longer exist. Although these missing files cause warning messages to be displayed during the installation or uninstallation of the Virtual GPU Manager, the installation or uninstallation is completed without errors.

Workaround

Ignore these messages as they are benign.

Status

Open

Ref. #

200423757

Chapter 94. Benign `not in gzip format` messages during Virtual GPU Manager installation or uninstallation

Description

During the installation or uninstallation of the Virtual GPU Manager, the warning messages are displayed:

```
gzip: /boot/initramfs-4.4.77-1.el6.nutanix.20170830.100726.x86_64.img: not in gzip
format
gzip: /boot/initramfs-4.4.77-1.el6.nutanix.20170830.100726.x86_64.tmp: not in gzip
format
```

Workaround

Ignore these messages as they are benign.

Status

Open

Ref.

200405700

Chapter 95. License is not acquired in Windows VMs

Description

When a windows VM configured with a licensed vGPU is started, the VM fails to acquire a license.

Error messages in the following format are written to the NVIDIA service logs:

```
[000000020.860152600 sec] - [Logging.lib] ERROR: [nvGridLicensing.FlexUtility]
353@FlexUtility::LogFneError : Error: Failed to add trusted storage. Server
URL : license-server-url -
[1,7E2,2,1[7000003F,0,9B00A7]]
```

```
System machine type does not match expected machine type..
```

Workaround

This workaround requires administrator privileges.

1. Stop the **NVIDIA Display Container LS** service.
2. Delete the contents of the folder %SystemDrive%\Program Files\NVIDIA Corporation\Grid Licensing.
3. Start the **NVIDIA Display Container LS** service.

Status

Closed

Ref.

200407287

Chapter 96. `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

Chapter 97. 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

Chapter 98. Hot plugging and unplugging vCPUs causes a blue-screen crash in Windows VMs

Description

Hot plugging or unplugging vCPUs causes a blue-screen crash in Windows VMs that are running NVIDIA vGPU software graphics drivers.

When the blue-screen crash occurs, one of the following error messages may also be seen:

- > `SYSTEM_SERVICE_EXCEPTION (nvlddmkm.sys)`
- > `DRIVER_IRQL_NOT_LESS_OR_EQUAL (nvlddmkm.sys)`

NVIDIA vGPU software graphics drivers do not support hot plugging and unplugging of vCPUs.

Status

Closed

Ref.

2101499

Chapter 99. GPU resources not available error during VMware instant clone provisioning

Description

A GPU resources not available error might occur during VMware instant clone provisioning. On Windows VMs, a Video TDR failure - NVLDDMKM.sys error causes a blue screen crash.

This error occurs when options for VMware Virtual Shared Graphics Acceleration (vSGA) are set for a VM that is configured with NVIDIA vGPU. VMware vSGA is a feature of VMware vSphere that enables multiple virtual machines to share the physical GPUs on ESXi hosts. NVIDIA vGPU software does **not** support VMware vSGA.

Depending on the combination of options set, one of the following error messages is seen when the VM is powered on:

> Module 'MKS' power on failed.

This message is seen when the following options are set:

- > **Enable 3D support** is selected.
- > **3D Renderer** is set to **Hardware**
- > The graphics type of all GPUs on the ESXi host is Shared Direct.
- > Hardware GPU resources are not available. The virtual machine will use software rendering.

This message is seen when the following options are set:

- > **Enable 3D support** is selected.
- > **3D Renderer** is set to **Automatic**.
- > The graphics type of all GPUs on the ESXi host is Shared Direct.

Resolution

If you want to use NVIDIA vGPU, unset any options for VMware vSGA that are set for the VM.

1. Ensure that the VM is powered off.
2. Open the vCenter Web UI.
3. In the vCenter Web UI, right-click the VM and choose **Edit Settings**.
4. Click the **Virtual Hardware** tab.
5. In the device list, expand the **Video card** node and de-select the **Enable 3D support** option.
6. Start the VM.

Status

Not a bug

Ref.

2369683

Chapter 100Module load failed during VIB downgrade from R390 to R384

Description

Some registry keys are available only with the R390 Virtual GPU Manager, for example, `NVreg_IgnoreMMIOCheck`. If any keys that are available only with the R390 Virtual GPU Manager are set, the NVIDIA module fails to load after a downgrade from R390 to R384.

When `nvidia-smi` is run without any arguments to verify the installation, the following error message is displayed:

```
NVIDIA-SMI has failed because it couldn't communicate with the NVIDIA driver. Make sure that the latest NVIDIA driver is installed and running.
```

Workaround

Before uninstalling the R390 VIB, clear all parameters of the `nvidia` module to remove any registry keys that are available only for the R390 Virtual GPU Manager.

```
# esxcli system module parameters set -p "" -m nvidia
```

Status

Not an NVIDIA bug

Ref.

200366884

Chapter 101vGPU 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, namely, Tesla M10. 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

Chapter 102Luxmark 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

Chapter 1030 On Linux, 3D applications run slowly when windows are dragged

Description

When windows for 3D applications on Linux are dragged, the frame rate drops substantially and the application runs slowly.

This issue does not affect 2D applications.

Status

Open

Ref.

1949482

Chapter 104A 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

Chapter 105 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.



Note: 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.

If you are licensing a physical GPU for NVIDIA vGPU for Compute, you **must** use the configuration file `/etc/nvidia/gridd.conf`.

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

Chapter 106 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

Chapter 107 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 Omnisson 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,

- > On XenServer, these errors are reported in the `/var/log/messages` file.
- > On VMware vSphere, these errors are reported in the log file `vmware.log` in the guest VM's storage directory.

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

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](#).

On VMware vSphere, you can also use the [Windows OS Optimization Tool for Omnicast Horizon](#) to make and apply optimization recommendations for Windows 10 and other operating systems.

If you are using XenServer, also refer to [Windows 10 Optimization Guide](#) on the Citrix blog.

Status

Open

Ref.

- > 200130864
- > 1803861

Chapter 108 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 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

Chapter 109 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 the XenServer `/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

Chapter 110 With no NVIDIA driver installed, XenServer misidentifies Tesla M10 cards

Description

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

Version

XenServer 6.5 and 7.0

Workaround

None

Status

Not an NVIDIA bug

Ref.

NVIDIA-420/1792341

Chapter 11 1vGPU VM fails to boot in ESXi if the graphics type is Shared

Description

On VMware vSphere Hypervisor (ESXi), after vGPU is configured, VMs to which a vGPU is assigned may fail to start and the following error message may be displayed:

```
The amount of graphics resource available in the parent resource pool is insufficient for the operation.
```

The vGPU Manager VIB provides vSGA and vGPU functionality in a single VIB. After this VIB is installed, the default graphics type is Shared, which provides vSGA functionality. To enable vGPU support for VMs in VMware vSphere, you must change the default graphics type to Shared Direct. If you do not change the default graphics type you will encounter this issue.



Note: NVIDIA vGPU software does **not** support VMware vSGA.

Workaround

Change the default graphics type to Shared Direct as explained in [Virtual GPU Software User Guide](#).

Status

Open

Ref.

200256224

Chapter 112GNOME 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

Chapter 113 NVIDIA Control Panel fails to start and reports that “you are not currently using a display that is attached to an Nvidia GPU”

Description

When you launch NVIDIA Control Panel on a VM configured with vGPU, it fails to start and reports that you are not using a display attached to an NVIDIA GPU. This happens because Windows is using VMware’s SVGA device instead of NVIDIA vGPU.

Fix

Make NVIDIA vGPU the primary display adapter.

Use Windows screen resolution control panel to make the second display, identified as “2” and corresponding to NVIDIA vGPU, to be the active display and select the Show desktop only on 2 option. Click Apply to accept the configuration.

You may need to click on the Detect button for Windows to recognize the display connected to NVIDIA vGPU.



Note: If the Omnissa Horizon agent is installed in the VM, the NVIDIA GPU is automatically selected in preference to the SVGA device.

Status

Open

Ref.

Chapter 114 VM configured with more than one vGPU fails to initialize vGPU when booted

Description

Using the current VMware vCenter user interface, it is possible to configure a VM with more than one vGPU device. When booted, the VM boots in VMware SVGA mode and doesn't load the NVIDIA driver. The additional vGPU devices are present in Windows Device Manager but display a warning sign, and the following device status:

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

Workaround

NVIDIA vGPU currently supports a single virtual GPU device per VM. Remove any additional vGPUs from the VM configuration before booting the VM.

Status

Open

Ref.

Chapter 115A VM configured with both a vGPU and a passthrough GPU fails to start the passthrough GPU

Description

Using the current VMware vCenter user interface, it is possible to configure a VM with a vGPU device and a passthrough (direct path) GPU device. This is not a currently supported configuration for vGPU. The passthrough GPU appears in Windows Device Manager with a warning sign, and the following device status:

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

Workaround

Do not assign vGPU and passthrough GPUs to a VM simultaneously.

Status

Open

Ref.

1735002

Chapter 116 vGPU allocation policy fails when multiple VMs are started simultaneously

Description

If multiple VMs are started simultaneously, vSphere may not adhere to the placement policy currently in effect. For example, if the default placement policy (breadth-first) is in effect, and 4 physical GPUs are available with no resident vGPUs, then starting 4 VMs simultaneously should result in one vGPU on each GPU. In practice, more than one vGPU may end up resident on a GPU.

Workaround

Start VMs individually.

Status

Not an NVIDIA bug

Ref.

200042690

Chapter 117 Before Horizon agent is installed inside a VM, the Start menu's sleep option is available

Description

When a VM is configured with a vGPU, the **Sleep** option remains available in the **Windows Start** menu. Sleep is not supported on vGPU and attempts to use it will lead to undefined behavior.

Workaround

Do not use Sleep with vGPU.

Installing the Omnissa Horizon agent will disable the **Sleep** option.

Status

Closed

Ref.

200043405

Chapter 118 vGPU-enabled VMs fail to start, `nvidia-smi` fails when VMs are configured with too high a proportion of the server's memory.

Description

If vGPU-enabled VMs are assigned too high a proportion of the server's total memory, the following errors occur:

- > One or more of the VMs may fail to start with the following error:

```
The available Memory resources in the parent resource pool are insufficient for the operation
```

- > When run in the host shell, the `nvidia-smi` utility returns this error:

```
-sh: can't fork
```

For example, on a server configured with 256G of memory, these errors may occur if vGPU-enabled VMs are assigned more than 243G of memory.

Workaround

Reduce the total amount of system memory assigned to the VMs.

Status

Closed

Ref.

200060499

Chapter 119 On reset or restart VMs fail to start with the error

```
VMIOOP: no graphics device is available for vGPU...
```

Description

On a system running a maximal configuration, that is, with the maximum number of vGPU VMs the server can support, some VMs might fail to start post a reset or restart operation.

Fix

Upgrade to ESXi 6.0 Update 1.

Status

Closed

Ref.

200097546

Chapter 120 `nvidia-smi` shows high GPU utilization for vGPU VMs with active Horizon sessions

Description

vGPU VMs with an active Horizon connection utilize a high percentage of the GPU on the ESXi host. The GPU utilization remains high for the duration of the Horizon session even if there are no active applications running on the VM.

Workaround

None

Status

Open

Partially resolved for Horizon 7.0.1:

- > For Blast connections, GPU utilization is no longer high.
- > For PCoIP connections, utilization remains high.

Ref.

1735009

Chapter 121 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

Chapter 122 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 `disable_vnc=1` in the vGPU plugin's parameters). 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.

> Citrix Hypervisor 8.1 or later:

```
[root@xenserver ~]# xe vgpu-param-set uuid=vgpu-uuid extra_args=disable_vnc=1
```

> Citrix Hypervisor earlier than 8.1:

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

Status

Open

Ref.

NVIDIA-145/1375164

Chapter 123 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

Chapter 124 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

Chapter 125 Windows VM BSOD

Description

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

The XenServer `/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

Chapter 126 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

Chapter 127XenCenter 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 not 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

Chapter 128 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

Chapter 129 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

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND 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 NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

VESA DisplayPort

DisplayPort and DisplayPort Compliance Logo, DisplayPort Compliance Logo for Dual-mode Sources, and DisplayPort Compliance Logo for Active Cables are trademarks owned by the Video Electronics Standards Association in the United States and other countries.

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, NVIDIA GRID vGPU, NVIDIA Maxwell, NVIDIA Pascal, NVIDIA Turing, NVIDIA Volta, GPUDirect, 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

© 2013-2026 NVIDIA Corporation & affiliates. All rights reserved.

