



# GRID VIRTUAL GPU FOR CITRIX XENSERVER

Version 340.57, 341.08

RN-06927-001 | November 10th, 2014

## Release Notes



# CONTENTS

**Release Notes** ..... 1

**Validated Platforms** ..... 2

**Known Issues** ..... 3

# RELEASE NOTES

These Release Notes summarize current status, information on validated platforms, and known issues with NVIDIA GRID Virtual GPU software and hardware on Citrix XenServer.

Included in this release is NVIDIA GRID Virtual GPU Manager versions 340.57 for Citrix XenServer 6.2 SP1 with hotfixes XS62ESP1009 and XS62ESP1011, and NVIDIA Windows drivers for vGPU, version 341.08.



**Note:** If you install the wrong package for the version of XenServer you are using, GRID vGPU Manager will fail to load. See issue #43.

The GRID vGPU Manager and Windows guest VM drivers must be installed together. Older VM drivers will not function correctly with this release of GRID vGPU Manager. Similarly, older GRID vGPU Managers will not function correctly with this release of Windows guest drivers. See issue #17.

Updates in this release:

- ▶ K160Q, K180Q, and K280Q vGPUs added
- ▶ Miscellaneous bug fixes



**Note:** to use these additional vGPUs, a hotfix must be applied to XenServer 6.2 SP1. See issue #41.

# VALIDATED PLATFORMS

This release of virtual GPU provides support for NVIDIA GRID K1 and K2 GPUs on Citrix XenServer, running on validated server hardware platforms. For a list of validated server platforms, refer to <http://www.nvidia.com/buygrid>.

## Software versions

This release has been tested with the following software versions:

Software	Version tested
Citrix XenServer	Version 6.2 with XS62ESP1 and applicable hotfixes including XS62ESP1009 and XS62ESP1011. The GRID vGPU Manager included in this release will not install without these hotfixes: <ul style="list-style-type: none"><li>▶ XS62ESP1009 <a href="http://support.citrix.com/article/CTX141191">http://support.citrix.com/article/CTX141191</a></li><li>▶ XS62ESP1011 <a href="http://support.citrix.com/article/CTX141472">http://support.citrix.com/article/CTX141472</a></li></ul>
Citrix XenDesktop	Version 7.1, 7.5, 7.6 in HDX 3D Pro mode. Versions 7.1 and 7.5 have the following hotfix applied: <ul style="list-style-type: none"><li>▶ XD710ICAWSWX86004 (32-bit) <a href="http://support.citrix.com/article/CTX140262">http://support.citrix.com/article/CTX140262</a> or XD710ICAWSWX64004 (64-bit) <a href="http://support.citrix.com/article/CTX140263">http://support.citrix.com/article/CTX140263</a></li></ul>



**Note:** GRID vGPU on Citrix XenServer does not support operation with physical GPUs BARs mapped above the 4 Gigabyte boundary in the system address space.

Ensure that GPUs are mapped below the 4G boundary by disabling your server's SBIOS option that controls 64-bit memory-mapped I/O support. This option may be labeled "Enable >4G Decode" or "Enable 64-bit MMIO". See issue #33.

# KNOWN ISSUES

<b>#9</b>	<b>Video goes blank when run in loop in Windows Media Player</b>	
	Description	When connected to a vGPU-enabled VM using Citrix XenDesktop, a video played back in looping mode on Windows Media Player goes blank or freezes after a few iterations.
	Workaround	None
	Status	Open
	Ref. #	1306623

<b>#17</b>	<b>VM running older NVIDIA vGPU drivers fails to initialize vGPU when booted</b>	
	Description	<p>A VM running older NVIDIA drivers, such as those from a previous vGPU release, will fail to initialize vGPU when booted on a XenServer platform running the current release of GRID Virtual GPU Manager.</p> <p>In this scenario, the VM boots in standard VGA mode with reduced resolution and color depth. The NVIDIA GRID GPU is present in Windows Device Manager but displays a warning sign, and a device status of “Windows has stopped this device because it has reported problems. (Code 43)”.</p> <p>Depending on the versions of drivers in use, XenServer’s <code>/var/log/messages</code> may contain the error message:</p> <pre>vmiop_log: error: Unable to fetch Guest NVIDIA driver information</pre> <p>or may report a version mismatch between guest and host drivers:</p> <pre>vmiop_log: error: Guest VGX version(1.1) and Host VGX version(1.2) do not match</pre>
	Fix	Install the latest NVIDIA vGPU release drivers in the VM.
	Status	Open
	Ref. #	

<b>#25</b>	<b>Virtual GPU fails to start if ECC is enabled on GRID K2 card</b>	
Description	<p>If the ECC (error correcting code) feature is enabled on a GRID K2 card, virtual GPU fails to start. The following error is logged in <code>/var/log/messages</code>:</p> <pre>vmiop_log: error: Initialization: VGX not supported with ECC Enabled.</pre> <p>Virtual GPU is not currently supported with ECC active. GRID K2 cards ship with ECC disabled by default, but ECC may subsequently be enabled using <code>nvidia-smi</code>.</p>	
Workaround	<p>Use <code>nvidia-smi</code> to list status on all GPUs, and check for ECC noted as enabled on GRID K2 GPUs. Change the ECC status to off on a specific GPU by executing <code>'nvidia-smi -i &lt;id&gt; -e 0'</code>, where <code>&lt;id&gt;</code> is the index of the GPU as reported by <code>nvidia-smi</code>.</p>	
Status	Open	
Ref. #		

<b>#27</b>	<b>Local VGA console is momentarily unblanked when XenDesktop changes resolution of the VM desktop</b>	
Description	<p>When XenDesktop establishes a remote connection to a VM using vGPU, the VM's local VGA console display in XenCenter is blanked (assuming the VM local console has not been disabled by setting <code>platform:vgpu_extra_args="disable_vnc=1"</code>). If the XenDesktop session changes resolution of the VM's desktop, the local VGA console momentarily unblanks, allowing a XenCenter user to briefly view the desktop.</p>	
Workaround	<p>Disable the VM's local VGA console:</p> <pre>xe vm-param-set uuid=&lt;vm-uuid&gt; platform:vgpu_extra_args="disable_vnc=1"</pre>	
Status	Open	
Ref. #	NVIDIA-145 / 1375164	

<b>#31</b>	<b>Single vGPU benchmark scores are lower than passthrough GPU</b>	
Description	<p>A single vGPU configured on a physical GPU produces lower benchmark scores than the physical GPU run in passthrough mode.</p> <p>Aside from performance differences that may be attributed to a vGPU's smaller framebuffer size, vGPU incorporates a performance balancing feature known as Frame Rate Limiter (FRL), which is enabled on all vGPUs. FRL is used to ensure balanced performance across multiple vGPUs that are resident on the same physical GPU. The FRL setting is designed to give good interactive remote graphics experience but may reduce scores in benchmarks that depend on measuring frame rendering rates, as compared to the same benchmarks running on a passthrough GPU.</p>	
Workaround	<p>FRL is controlled by an internal vGPU setting. NVIDIA does not validate vGPU with FRL disabled, but for validation of benchmark performance, FRL can be temporarily disabled by specifying <code>frame_rate_limiter=0</code> in the VM's <code>platform:vgpu_extra_args</code> parameter:</p> <pre data-bbox="634 877 1471 982">[root@xenserver ~]# xe vm-param-set uuid=e71afda4-53f4-3alb-6c92-a364a7f619c2 platform:vgpu_extra_args="frame_rate_limiter=0" [root@xenserver ~]#</pre> <p>The setting takes effect the next time the VM is started or rebooted.</p> <p>With this setting in place, the VM's vGPU will run without any frame rate limit. The FRL can be reverted back to its default setting by removing the <code>vgpu_extra_args</code> key from the <code>platform</code> parameter, or by removing <code>frame_rate_limiter=0</code> from the <code>vgpu_extra_args</code> key, or by setting <code>frame_rate_limiter=1</code>. For example:</p> <pre data-bbox="634 1318 1471 1402">[root@xenserver ~]# xe vm-param-set uuid=e71afda4-53f4-3alb-6c92-a364a7f619c2 platform:vgpu_extra_args="frame_rate_limiter=1" [root@xenserver ~]#</pre>	
Status	Open	
Ref. #		



<b>#33</b>	<b>Virtual GPU fails to start when GPUs are mapped above 4G</b>	
	Description	<p>GRID vGPU 1.0 on Citrix XenServer does not support operation with GPUs mapped above the 4 gigabyte (4G) boundary in the system's physical address space.</p> <p>If GPUs are mapped above 4G, the GRID vGPU Manager rpm will warn at the time of installation:</p> <pre>Warning: vGPU does not support GPUs mapped in 64-bit address space. Please disable 64-bit MMIO from the system's BIOS. Refer to vGPU release notes for details."</pre> <p>Also, the NVIDIA kernel driver will fail to load in XenServer's dom0, so the <code>nvidia</code> module won't appear in the module listing produced by <code>lsmod</code>. Additionally, the following warning messages will be present in the output of <code>dmesg</code>:</p> <pre>NVRM: This PCI I/O region assigned to your NVIDIA device is invalid: NVRM: BAR1 is 128M @ 0xf800000000000000 (PCI:03ff:00:07.0) NVRM: This is a 64-bit BAR mapped above 4GB by the system NVRM: BIOS or the Linux kernel. The NVIDIA Linux/x86 NVRM: graphics driver and other system software components NVRM: do not support this configuration.</pre>
	Workaround	Ensure that GPUs are mapped below the 4G boundary by disabling your server's SBIOS option that controls 64-bit memory-mapped I/O support. This option may be labeled "Enable >4G Decode" or "Enable 64-bit MMIO".
	Status	Fixed in XenServer "Creedence".
	Ref. #	NVIDIA-184
<b>#34</b>	<b>VM bugchecks on shutdown/restart when XenDesktop is installed and NVIDIA driver is uninstalled or upgraded.</b>	
	Description	If the XenDesktop agent is installed in a VM before any NVIDIA GPU driver is installed, the VM will bugcheck (bluescreen) when the NVIDIA driver is subsequently upgraded or uninstalled. The bugcheck code is <code>0x7E, SYSTEM_THREAD_EXCEPTION_NOT_HANDLED</code> .
	Workaround	Do a force shutdown of the VM and restart it. Alternatively, install the NVIDIA driver in guest VMs before installing XenDesktop.
	Status	Open
	Ref. #	NVIDIA-295 / 200018125

<b>#35</b>	<b>Application frame rate may drop when running XenDesktop at 2560x1600 resolution.</b>	
	Description	An application's rendering frame rate may drop when running XenDesktop at 2560x1600 resolution, relative to the frame rate obtained at lower resolutions.
	Fix	Using the Windows <code>regedit</code> utility within the VM, open the <code>HKLM\SOFTWARE\Citrix\Graphics</code> registry key and create a new DWORD value, <code>EncodeSpeed</code> , 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
<b>#39</b>	<b>Windows VM operated with XenDesktop 7.1 with rotated screen resolution of 1600x2560 results in a VM crash.</b>	
	Description	Connecting via XenDesktop to a VM with a vGPU, from a client that has a 2560x1600 monitor in rotated mode, causes the VM to crash with bugcheck 0x7F.
	Fix	Apply XenServer 6.2 SP1 hotfix XS62ESP1008, available at <a href="http://support.citrix.com/article/CTX141036">http://support.citrix.com/article/CTX141036</a> .
	Status	Fixed
	Ref. #	NVIDIA-247 / 1482213

#40

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

Description

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

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

This is because GPUs operating in passthrough mode are not visible to nvidia-smi and the NVIDIA kernel driver operating in XenServer's dom0.

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



Fix N/A

Status Open

Ref. # N/A

<b>#41</b>	<b>K120Q, K220Q vGPU types aren't shown in XenCenter GUI</b>	
	Description	With host driver 340.57 included in this release, new vGPU types K120Q, K160Q, K180Q, K220Q, and K280Q are supported for GRID K1 and GRID K2 cards respectively, but these vGPU types are not listed in the XenCenter GPU selection interface.
	Fix	Install XenServer hotfix XS62ESP1004 from Citrix, available at <a href="http://support.citrix.com/article/CTX140417">http://support.citrix.com/article/CTX140417</a>
	Status	Fixed
	Ref. #	N/A
<b>#42</b>	<b>Windows Aero is disabled on XenDesktop session using 3 or 4 monitors in 2560x1600 resolution</b>	
	Description	Windows Aero may be disabled when XenDesktop is connected to a VM with a vGPU or passthrough GPU, with 3 or 4 monitors at 2560x1600 resolution.
	Workaround	
	Status	Open
	Ref. #	NVIDIA-226 / 1456343

<b>#43</b>	<b>VMs configured with large memory fail to initialize vGPU when booted</b>	
	Description	<p>When starting multiple VMs configured with large amounts of RAM (typically more than 32GB per VM), a VM may fail to initialize vGPU. In this scenario, the VM boots in standard VGA mode with reduced resolution and color depth. The NVIDIA GRID GPU is present in Windows Device Manager but displays a warning sign, and a device status of “Windows has stopped this device because it has reported problems. (Code 43)”.</p> <p>XenServer’s <code>/var/log/messages</code> contains these error messages:</p> <pre> vmiop_log: error: NVOS status 0x29 vmiop_log: error: Assertion Failed at 0x7620fd4b:179 vmiop_log: error: 8 frames returned by backtrace ... vmiop_log: error: VGPU message 12 failed, result code: 0x29 ... vmiop_log: error: NVOS status 0x8 vmiop_log: error: Assertion Failed at 0x7620c8df:280 vmiop_log: error: 8 frames returned by backtrace ... vmiop_log: error: VGPU message 26 failed, result code: 0x8 </pre>
Workaround	<p>vGPU reserves a portion of the VM’s framebuffer for use in GPU mapping of VM system memory. The reservation is sufficient to support up to 32GB of system memory, and may be increased to accommodate up to 64GB by specifying <code>enable_large_sys_mem=1</code> in the VM’s <code>platform:vgpu_extra_args</code> parameter:</p> <pre>[root@xenserver ~]# xe vm-param-set uuid=e71afda4-53f4-3a1b-6c92-a364a7f619c2 platform:vgpu_extra_args="enable_large_sys_mem=1"</pre> <p>The setting takes effect the next time the VM is started or rebooted. With this setting in place, less GPU FB is available to applications running in the VM. To accommodate system memory larger than 64GB, the reservation can be further increased by specifying <code>extra_fb_reservation</code> in the VM’s <code>platform:vgpu_extra_args</code> parameter, and settings its value to the desired reservation size in megabytes. The default value of 64M is sufficient to support 64GB of RAM. We recommend adding 2M of reservation for each additional 1GB of system memory. For example, to support 96GB of RAM, set <code>extra_fb_reservation</code> to 128:</p> <pre>platform:vgpu_extra_args="enable_large_sys_mem=1, extra_fb_reservation=128"</pre> <p>The reservation can be reverted back to its default setting by removing the <code>vgpu_extra_args</code> key from the platform parameter, or by removing <code>enable_large_sys_mem</code> from the <code>vgpu_extra_args</code> key, or by setting <code>enable_large_sys_mem=0</code>.</p>	
Status	Open	
Ref. #	1510886	

## Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

### **HDMI**

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

### **OpenCL**

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

### **Trademarks**

NVIDIA and the NVIDIA logo are trademarks and/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**

© 2012 NVIDIA Corporation. All rights reserved.