



GRID VIRTUAL GPU FOR CITRIX XENSERVER Version 331.30 / 332.07

RN-06927-001 | December 19th, 2013

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 for Citrix XenServer, version 331.30, and NVIDIA Windows 7 32/64-bit drivers for vGPU, version 332.07. These packages form the initial production release of GRID vGPU software and supersede all previous vGPU releases.



Note: the GRID vGPU Manager and Windows guest VM drivers must be installed together. Older VM drivers, such as those released for vGPU Tech Preview, 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.

When upgrading a XenDesktop-enabled VM from a vGPU Tech Preview driver to this driver, a VM may crash during reboot after driver installation. See issue #34.

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>.



Note: GRID vGPU 1.0 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

#9	Video goes blank when run in loop in Windows Media Player	
	Description	When connected to a vGPU-enabled VM using Citrix XenDesktop, a video played back in looping mode on Windows Media Player goes blank or freezes after a few iterations.
	Workaround	None
	Status	Open
	Ref. #	1306623
#13	Console VGA display freezes if NVIDIA driver is uninstalled from a VM	
	Description	The console VGA display shown in XenCenter freezes if the NVIDIA graphics driver is uninstalled in the VM.
	Workaround	Uninstallation of the NVIDIA driver is not required for normal operation of the VM, or during driver upgrade.
	Status	Open
	Ref. #	1289674

<p>#17</p>	<p>VM running older NVIDIA vGPU drivers fails to initialize vGPU when booted</p>	
	<p>Description</p>	<p>A VM running older NVIDIA drivers, such as those released for vGPU Tech Preview or as part of the vGPU private beta program, will fail to initialize vGPU when booted on a XenServer platform running the production 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(0.5) and Host VGX version(1.0) do not match</pre>
	<p>Fix</p>	<p>Install the latest NVIDIA vGPU release drivers in the VM.</p>
	<p>Status</p>	<p>Open</p>
	<p>Ref. #</p>	

#22	XenDesktop session fails to refresh when any application is launched and maximized at 2560x1600 resolution	
Description	When connecting via XenDesktop to a Windows VM running vGPU such that the remoted desktop resolution is 2560x1600, opening an application and maximizing its window results in image corruption; the desktop fails to refresh correctly, and when windows are closed, remnants remain visible.	
Workaround	Using the Windows <code>regedit</code> utility within the VM, open the <code>HKLM\SOFTWARE\Citrix\HDX3D\BitmapRemotingConfig</code> registry key and create a new DWORD value, <code>HKLM_EnableDirtyRect</code> , with a value of 0. Reboot the VM.	
Status	Open	
Ref. #	1369303	
#23	XenDesktop session falls back to GDI scraper mode when using multiple, misaligned monitors	
Description	If a XenDesktop session is established in full-screen mode from a client having multiple monitors, and their logical alignment is configured as unaligned in the Windows display control panel in the VM, the XenDesktop session may drop back to GDI scraper mode instead of using VGX optimized mode. This is detectable by Windows Aero being disabled and lower overall session performance.	
Workaround	Disconnect and reconnect the session. Alternatively, use Windows display control panel in the VM to logically align the monitors.	
Status	Open	
Ref. #	1359581	
#24	XenDesktop session displays blank image when using multiple monitors each with different custom resolutions	
Description	If a Citrix XenDesktop session is initiated in full-screen mode from a multi-monitor client that has each monitor set to different custom resolutions (for example, 1440x1196 on one monitor and 1440x900 on the other), a blank, gray window is seen in place of the expected desktop image.	
Workaround		
Status	Open	
Ref. #	NVIDIA-119	
#25	Virtual GPU fails to start if ECC is enabled on GRID K2 card	

	<p>Description</p>	<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 /var/log/messages:</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>
	<p>Workaround</p>	<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 <id> -e 0</code>', where <code><id></code> is the index of the GPU as reported by <code>nvidia-smi</code>.</p>
	<p>Status</p>	<p>Open</p>
	<p>Ref. #</p>	
<p>#26</p>	<p>GRID K1 GPU group is listed as "Group of NVIDIA Corporation None GPUs"</p>	
	<p>Description</p>	<p>The name label of the default GPU group created for GRID K1 GPUs may be reported as "Group of NVIDIA Corporation None GPUs" instead of "Group of NVIDIA Corporation GRID K1 GPUs".</p>
	<p>Workaround</p>	<p>This issue has no impact on vGPU creation and operation. The name label field of the GPU group can be manually corrected using:</p> <pre>xe gpu-group-param-set uuid=<group uuid> name-label="Group of NVIDIA Corporation GRID K1 GPUs"</pre>
	<p>Status</p>	<p>Open</p>
	<p>Ref. #</p>	
<p>#27</p>	<p>Local VGA console is momentarily unblanked when XenDesktop changes resolution of the VM desktop</p>	
	<p>Description</p>	<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>
	<p>Workaround</p>	<p>Disable the VM's local VGA console:</p> <pre>xe vm-param-set uuid=<vm-uuid> platform:vgpu_extra_args="disable_vnc=1"</pre>
	<p>Status</p>	<p>Open</p>
	<p>Ref. #</p>	<p>1375164</p>

#31	Single vGPU benchmark scores are lower than passthrough GPU	
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="609 871 1458 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="609 1312 1458 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. #		

#33	Virtual GPU fails to start when GPUs are mapped above 4G	
	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 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>
	Fix	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	
	Ref. #	NVIDIA-184
#34	VM bugchecks on shutdown/restart when XenDesktop is installed and NVIDIA tech preview driver is uninstalled or upgraded.	
	Description	If an NVIDIA vGPU Tech Preview driver in a Windows VM is upgraded or uninstalled on a platform running the 331.30 release of GRID vGPU Manager, and the XenDesktop 7.1 agent is present in the VM, the VM may bugcheck (bluescreen) while shutting down, with a 0x7E bugcheck code, <code>SYSTEM_THREAD_EXCEPTION_NOT_HANDLED</code> .
	Fix	Do a force shutdown of the VM and restart it.
	Status	Open
	Ref. #	NVIDIA-199 / 1419049

#35	Application frame rate may drop when running XenDesktop at 2560x1600 resolution.	
	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
#36	XenDesktop session fails to establish in dual-monitor mode with monitors at 1600x900 and 1920x1200 resolution	
	Description	When connecting in full-screen mode from a client with dual monitors at 1600x900 and 1920x1200 resolutions, the XenDesktop session may fail to establish in multi-monitor mode; only one virtual monitor is connected to the virtual machine's desktop.
	Fix	
	Status	Open
	Ref. #	NVIDIA-191 / 1411864
#37	VM with 1 vCPU may hang during recovery from a TDR	
	Description	A VM configured with a single vCPU may hang during recovery from a Windows Timeout Detection and Recovery (TDR) condition. (A TDR is a graphics timeout caused by an application, driver, or system error, and is normally recovered to allow continued operation of the VM.)
	Fix	A workaround is to assign two or more vCPUs to the VM.
	Status	Open
	Ref. #	1407959

#38	Windows 8 VM on heavily loaded GPU may TDR when XenDesktop session is changed between windowed and fullscreen mode.	
	Description	A Windows 8 VM running alongside other VMs on the same physical GPU, each running heavy GPU workloads, may TDR when a connected XenDesktop session is switched between windowed and fullscreen mode. TDR is a graphics timeout condition, which can result in the VM hanging or rebooting.
	Fix	
	Status	Open
	Ref. #	1389326

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.