



# GRID VIRTUAL GPU FOR CITRIX XENSERVER Version 312.53 / 312.56

RN-06927-001 | October 1<sup>st</sup>, 2013

## Release Notes



# CONTENTS

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

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

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

# RELEASE NOTES

This edition of Release Notes summarizes current status, information on validated platforms, and known issues with NVIDIA GRID Virtual GPU R312 drivers and hardware on Citrix XenServer.

Included in this release is NVIDIA GRID Virtual GPU Manager for Citrix XenServer, version 312.53, and NVIDIA Windows 7 32/64-bit drivers for vGPU, version 312.56.

# VALIDATED PLATFORMS

This chapter describes server hardware platforms and card configurations NVIDIA has currently validated with GRID vGPU.

Server Platform	SBIOS version	Card Configurations
Dell PowerEdge R720	1.5.1 (1/7/2013)	1x GRID K1 1x GRID K2
Cisco UCS C240 M3	C240M3.1.5.1c.0 (1/31/2013)	2x GRID K1 2x GRID K2
Supermicro SYS-1027GR-TRF Supermicro SYS-2027GR-TRF	1.86 (11/14/2012)	2x GRID K1 2x GRID K2



**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".

Also supported, validation not yet completed:

Server Platform	SBIOS version	Card Configurations
IBM iDataPlex dx360 M4		2x GRID K1 2x GRID K2

# KNOWN ISSUES

<b>#2</b>	<b>VMs may fail to start when running pass-through GPUs together with vGPUs</b>	
	Description	When a mix of virtual GPUs and passthrough GPUs are in use, VMs may fail to start. This happens because the pciback kernel driver that is used for passthrough remains loaded in dom0 on a physical GPU selected to host a vGPU, and the NVIDIA kernel driver required for vGPU is unable to dynamically reload.
	Workaround	Create one or more <code>gpu-group</code> objects to hold physical GPU(s) to be used exclusively for passthrough. Use <code>'xe gpu-group-create'</code> to create new GPU groups, as described in the GRID Virtual GPU User Guide.
	Status	Open
	Ref. #	1318957 / 1341238
<b>#4</b>	<b>Upgrading the GRID Virtual GPU Manager RPM doesn't uninstall the private beta version of the RPM</b>	
	Description	If the private beta GRID Virtual GPU Manager package is already installed on XenServer, installing the Tech Preview package does not cause the private beta package to be uninstalled. This may lead to incorrect operation.
	Workaround	Manually remove the existing, earlier version of the Virtual GPU Manager RPM before installing the updated one. This step is documented in the User Guide, section 2.3.  <pre>[root@xenserver ~]# rpm -qa   grep NVIDIA NVIDIA-vgx-312.38-xenserver-6-2 [root@xenserver ~]# rpm -ev NVIDIA-vgx-312.38-xenserver-6-2 [root@xenserver ~]#</pre>
	Status	Open.
	Ref. #	1315586

<b>#8</b>	<b>Mouse cursor disappears when dragged over Catia window</b>	
	Description	When connected to a GPU-enabled VM using XenDesktop, dragging the cursor over the active UI window of the Catia V6 application causes the cursor to disappear.
	Workaround	This is a known issue in the Citrix XenDesktop Tech Preview for virtual GPU, and can be worked around by installation of a modified HDX DLL and registry key. Refer to Citrix's XenDesktop release notes for instructions.
	Status	Open
	Ref. #	1311739
<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>#13</b>	<b>Console VGA display freezes if NVIDIA driver is uninstalled from a VM</b>	
	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
<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 released as part of the vGPU private beta program, will fail to initialize vGPU when booted on a XenServer platform running the Tech Preview 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)". XenServer's <code>/var/log/messages</code> contains the error message: <code>vmiop_log: error: Unable to fetch Guest NVIDIA driver information</code></p>
	Fix	Install the latest NVIDIA Tech Preview release drivers in the VM.
	Status	Open
	Ref. #	
<b>#18</b>	<b>VM VGA console may drop into low resolution / limited color mode prior to NVIDIA driver being installed</b>	
	Description	If a virtual machine is configured with virtual GPU, booted, shutdown and rebooted without the NVIDIA drivers being installed in the VM, the VM console may drop into a low resolution mode with limited colors.
	Workaround	Install the NVIDIA driver in the VM, then reboot the VM to return its VGA console to normal resolution and color depth.
	Status	Open
	Ref. #	1192791

<b>#19</b>	<b>Windows 8 VM fails on startup with inaccessible boot device bugcheck when vGPU is enabled</b>	
	Description	A Windows 8 VM with XenServer tools installed fails on startup when a vGPU is enabled for the VM, with an inaccessible boot device bugcheck (bluescreen).
	Workaround	Enable vGPU on the Windows 8 VM, and install NVIDIA drivers in the VM, before installing XenServer tools in the VM.
	Status	Open
	Ref. #	1362638
<b>#20</b>	<b>VMs with a virtual GPU cannot be powered on unless VT-d is enabled on the platform</b>	
	Description	<p>A VM with a virtual GPU cannot be powered on if VT-d is not enabled on the platform. If started from XenCenter, an error dialog reports “You attempted to run a VM on a host which doesn’t have I/O virtualization (IOMMU/VT-d) enabled, which is needed by the VM. If started from the command line using <code>xe vm-start</code>, the command reports:</p> <pre>[root@xenserver ~]# xe vm-start uuid=0fa61c2a-6748-9d6c-34b3-64a9fd9d4dd9 There are no suitable hosts to start this VM on. The following table provides per-host reasons for why the VM could not be started:</pre> <pre>xenserver-vgx (VM IPs 10.31.213.50-95, dom0 .98, OOB .99)      : Cannot start here [Api_errors.Server_error("VM_REQUIRES_IOMMU", _)]</pre> <p>There were no hosts available to complete the specified operation.</p>
	Workaround	Enable VT-d in the platform’s system BIOS settings.
	Status	Open
	Ref. #	
<b>#21</b>	<b>Forced reboot of a heavily loaded VM causes VM halt with error "Internal error : Xenopsd internal error: Domain.Domain_stuck_in_dying_state"</b>	
	Description	<p>When running a VM with a heavy workload using virtual GPU, forcing a reboot of the VM may cause it to halt and report the following error in <code>/var/log/messages</code>:</p> <pre>Internal error : Xenopsd internal error : Domain.Domain_stuck_in_dying_state(29)</pre>
	Workaround	Re-start the VM.

	Status	Open
	Ref. #	1362033
<b>#22</b>	<b>XenDesktop session fails to refresh when any application is launched and maximized at 2560x1600 resolution</b>	
	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 <code>DWORD</code> value, <code>HKLM_EnabledDirtyRect</code> , with a value of 0. Reboot the VM.
	Status	Open
	Ref. #	1369303
<b>#23</b>	<b>XenDesktop session falls back to GDI scraper mode when using multiple, misaligned monitors</b>	
	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
<b>#24</b>	<b>XenDesktop session displays blank image when using multiple monitors each with different custom resolutions</b>	
	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

<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	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> .
	Status	Open
	Ref. #	
<b>#26</b>	<b>GRID K1 GPU group is listed as "Group of NVIDIA Corporation None GPUs"</b>	
	Description	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".
	Workaround	This issue has no impact on vGPU creation and operation. The name label field of the GPU group can be manually corrected using: <pre>xe gpu-group-param-set uuid=&lt;group uuid&gt; name-label="Group of NVIDIA Corporation GRID K1 GPUs"</pre>
	Status	Open
	Ref. #	
<b>#27</b>	<b>Local VGA console is momentarily unblanked when XenDesktop changes resolution of the VM desktop</b>	
	Description	When XenDesktop establishes a remote connection to a VM using vGPU, the VM's local VGA console display in XenCenter is blanked (assuming the VM local console has not been disabled by setting <code>platform:vgpu_vnc_enabled=false</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.
	Workaround	Disable the VM's local VGA console:  <pre>xe vm-param-set uuid=&lt;vm-uuid&gt; platform:vgpu_vnc_enabled=false</pre>
	Status	Open

	Ref. #	1375164
<b>#28</b>	<b>Execution of nvidia-smi in dom0 causes stutter in XenDesktop streaming</b>	
	Description	Launching <code>nvidia-smi</code> in XenServer's dom0 may cause a visible stutter (a momentary reduction in framerate) in XenDesktop sessions connecting to VMs using vGPU.
	Workaround	
	Status	Open
	Ref. #	1332542
<b>#30</b>	<b>nvidia-smi fails to operate when passthrough and virtual GPUs are active</b>	
	Description	<p>If passthrough and vGPUs are configured and utilized simultaneously on a platform, <code>nvidia-smi</code> may report the following error:</p> <pre>NVIDIA: could not open the device file /dev/nvidial (Input/output error). NVIDIA-SMI has failed because it couldn't communicate with NVIDIA driver. Make sure that latest NVIDIA driver is installed and running.</pre>
	Workaround	
	Status	Open
	Ref. #	1320904

## **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.