



# GRID VIRTUAL GPU FOR CITRIX XENSERVER

Version 352.83, 354.80

RN-06927-001 | February 29<sup>th</sup>, 2016

**Release Notes**



# CONTENTS

- Release Notes** ..... 1
- Validated Platforms** ..... 2
  - Software versions ..... 2
  - Linux support ..... 3
  - Hardware configuration ..... 3
- Known Issues** ..... 4

# RELEASE NOTES

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

Included in this release is NVIDIA GRID Virtual GPU Manager versions 352.83 for Citrix XenServer 6.5 SP1, Citrix XenServer 6.2 SP1 with hotfixes XS62ESP1009 and XS62ESP1011, NVIDIA Windows drivers for vGPU version 354.80, and NVIDIA Linux drivers for vGPU version 352.83.



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

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 “VM running older NVIDIA vGPU drivers fails to initialize vGPU when booted,” on page 5.

Updates in this release:

- ▶ Miscellaneous bug fixes
- ▶ Windows 10 support

# VALIDATED PLATFORMS

This release of virtual GPU provides support for NVIDIA GRID K1, K2, and Tesla M60, M6 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 6.2	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 XenServer 6.5	Version 6.5 with XS65ESP1. The GRID vGPU Manager included in this release will not install without XenServer 6.5 SP1.
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</li><li>• 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 6.2 does not support operation with physical GPUs BARs mapped above the 4 Gigabyte boundary in the system address space.

For XenServer 6.2, 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 “Virtual GPU fails to start when GPUs are mapped above 4G.” on page 5.

## Linux support

GRID vGPU with Red Hat Enterprise Linux 7 / CentOS 7 guest VMs is supported on Tesla M60 and M6, as a technical preview feature on Citrix XenServer 6.5.



Note: Linux vGPU is not supported on XenServer 6.2.

## HARDWARE CONFIGURATION

Tesla M60 and M6 GPUs support compute and graphics modes, which can be configured using the *gpumodeswitch* tool provided with GRID software releases. GRID vGPU requires that M60 / M6 are configured in graphics mode.

# KNOWN ISSUES

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

VM running older NVIDIA vGPU drivers fails to initialize vGPU when booted	
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> <p>or report a signature mismatch:</p> <pre>vmiop_log: error: vGPU message signature mismatch.</pre>
Fix	Install the latest NVIDIA vGPU release drivers in the VM.
Status	Open
Ref. #	

Virtual GPU fails to start if ECC is enabled	
Description	<p>GRID K2, Tesla M60, and Tesla M6 support ECC (error correcting code) for improved data integrity. If ECC is enabled, 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 and Tesla M60, M6 cards in graphics mode 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 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. #	

Local VGA console is momentarily unblanked when XenDesktop changes resolution of the VM desktop	
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_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.
Workaround	Disable the VM's local VGA console:  <pre>xe vm-param-set uuid=&lt;vm-uuid&gt; platform:vgpu_extra_args="disable_vnc=1"</pre>
Status	Open
Ref. #	NVIDIA-145 / 1375164

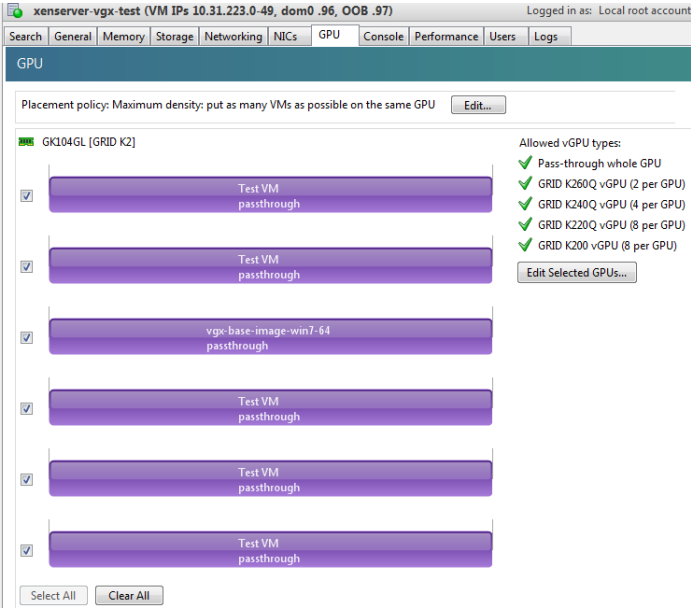


Single vGPU benchmark scores are lower than passthrough GPU	
Description	<p>A single vGPU configured on a physical GPU produces lower benchmark scores than the physical GPU run in passthrough mode.</p> <p>Aside from performance differences that may be attributed to a vGPU's smaller framebuffer size, 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>[root@xenserver ~]# xe vm-param-set uuid=e71afda4-53f4-3a1b-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>[root@xenserver ~]# xe vm-param-set uuid=e71afda4-53f4-3a1b-6c92-a364a7f619c2 platform:vgpu_extra_args="frame_rate_limiter=1" [root@xenserver ~]#</pre>
Status	Open
Ref. #	

Virtual GPU fails to start when GPUs are mapped above 4G	
Description	<p>GRID vGPU on Citrix XenServer 6.2 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>
Version	XenServer 6.2
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 6.5
Ref. #	NVIDIA-184

VM bugchecks on shutdown/restart when XenDesktop is installed and NVIDIA driver is uninstalled or upgraded.	
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

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

<h2 style="margin: 0;">nvidia-smi fails to operate when all GPUs are assigned to GPU passthrough mode.</h2>	
Description	<p>If all GPUs in the platform are assigned to VMs in passthrough mode, nvidia-smi will return an error:</p> <pre>[root@xenserver-vgx-test ~]# nvidia-smi Failed to initialize NVML: Unknown Error</pre> <p>This is because GPUs operating in passthrough mode are not visible to nvidia-smi and the NVIDIA kernel driver operating in XenServer's dom0.</p> <p>To confirm that all GPUs are operating in passthrough, use XenCenter's GPU tab to review current GPU assignment:</p> 
Fix	N/A
Status	Open
Ref. #	N/A

Windows Aero is disabled on XenDesktop session using 3 or 4 monitors in 2560x1600 resolution	
Description	<p>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.</p> <p>This is a limitation of Windows 7, refer Microsoft's knowledge base article at <a href="https://support.microsoft.com/en-us/kb/2724530">https://support.microsoft.com/en-us/kb/2724530</a>.</p>
Workaround	
Status	Closed
Ref. #	NVIDIA-226 / 1456343

VMs configured with large memory fail to initialize vGPU when booted	
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	Fixed
Ref. #	1510886

Windows VM BSOD	
Description	<p>Windows VM bugchecks on XenServer when running a large number of vGPU based VMs.</p> <p>XenServer's <code>/var/log/messages</code> contains these error messages:</p> <pre> 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. </pre>
Version	XenServer 6.2
Fix	Please ensure that you are running the latest OEM firmware for your GRID boards.
Status	Closed
Ref. #	NVIDIA-327 / 1632120

Desktop session is unable to scale to resolutions higher than 1280x1024 on VMs running vGPU on Tesla M6 or Tesla M60	
Description	GRID vGPU is a licensed feature on Tesla M6 and M60. A software license is required to enable full vGPU features within the guest VM.
Version	
Fix	Please refer to GRID Licensing User Guide on how to license your vGPU VM.
Status	Open
Ref. #	

Desktop sessions run at a low frame rate on VMs running vGPU on Tesla M6 or Tesla M60	
Description	GRID vGPU is a licensed feature on Tesla M6 and M60. A software license is required to enable full vGPU features within the guest VM.
Version	
Fix	Please refer to <a href="#">GRID Licensing User Guide</a> on how to license your vGPU VM.
Status	Open
Ref. #	

Applications fail to start on VMs running vGPU on Tesla M6 or Tesla M60	
Description	<p>Applications fail to start within the VM, and the following error is logged in <code>/var/log/messages</code>:</p> <pre>vmiop_log: error: Guest is unlicensed. Cannot allocate more than 0x20 channels!</pre> <p>GRID vGPU is a licensed feature on Tesla M6 and M60. A software license is required to enable full vGPU features within the guest VM.</p>
Version	
Fix	Please refer to <a href="#">GRID Licensing User Guide</a> on how to license your vGPU VM.
Status	Open
Ref. #	

MPlayer, or other video players, fail to start when using hardware acceleration on Linux VMs running vGPU	
Description	VDPAU is currently not supported on Linux VMs running vGPU.
Version	
Workaround	Configure the video player to not use VDPAU or hardware acceleration.
Status	Open
Ref. #	200144667



Windows VM BSOD when upgrading NVIDIA drivers over a XenDesktop session	
Description	Windows VM bugchecks when NVIDIA guest drivers are upgraded over a XenDesktop session.  If the VM is restarted after the bugcheck, the upgraded driver loads correctly and full functionality is available.
Version	
Fix	Upgrade XenDesktop to 7.6 Feature Pack 3
Status	Closed
Ref. #	NVIDIA-370 / 200130780

vGPU host driver RPM upgrade fails	
Description	Upgrading vGPU host driver RPM fails with the following message on the console:  <pre>[root@xenserver ~]# rpm -U NVIDIA-vGPU-xenserver-6.5-352.46.x86_64.rpm error: Failed dependencies:         NVIDIA-vgx-xenserver conflicts with NVIDIA-vGPU-xenserver-6.5-352.46.x86_64 [root@xenserver ~]#</pre>
Version	
Workaround	Uninstall the older vGPU RPM before installing the latest driver.  Use the following command to uninstall the older vGPU RPM: <pre>[root@xenserver ~]# rpm -e NVIDIA-vgx-xenserver</pre>
Status	Closed
Ref. #	200133635

XenCenter does not allow vGPUs to be selected as a GPU type for Linux VMs	
Description	<p>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.</p> <p>vGPU on Linux VMs is supported as a technical preview on XenServer 6.5, and does include XenCenter integration.</p>
Version	
Workaround	Please refer to <a href="#">Chapter 5 - XenServer vGPU Management in the GRID vGPU User Guide</a> on how to configure vGPU using <code>xe CLI</code> .
Status	Closed
Ref. #	NVIDIA-360

If X server is killed on a RHEL7 VM running vGPU, XenCenter console may not automatically switch to text console	
Description	<p>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.</p> <p>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.</p>
Version	
Workaround	Use <code>CTRL + ALT + F[1 2 3 ...]</code> to switch between Linux terminals
Status	Closed
Ref. #	NVIDIA-350 / 200123378

PowerPoint intermittently crashes in slideshow mode on XenDesktop in a dual-monitor setup	
Description	Running the PowerPoint application in slideshow mode over a XenDesktop session in a dual-monitor setup sometimes causes the application to crash.
Version	
Workaround	
Status	Fixed
Ref. #	200172195

Multiple WebGL tabs in Microsoft Internet Explorer may trigger TDR on Windows VMs	
Description	Running intensive WebGL applications in multiple IE tabs may trigger a TDR on Windows VMs.
Version	
Workaround	Disable hardware acceleration in IE.  To enable software rendering in IE, refer Microsoft's knowledge base article at <a href="https://support.microsoft.com/en-us/kb/2528233">https://support.microsoft.com/en-us/kb/2528233</a> .
Status	Open
Ref. #	200148377

CATIA intermittently crashes during quick object rotation	
Description	Using the mouse to quickly rotate large models in the CATIA application sometimes causes the application to crash.
Fix	
Status	Fixed
Ref. #	200155765

Loading large models in CATIA causes an OpenGL failure	
Description	Loading very large models in CATIA exhausts the frame buffer, causing Open GL to fail with the error OPENGL 8.
Fix	
Status	Fixed
Ref. #	1717876

Black screen on initial connection to a VM using XenDesktop from a client with multiple displays	
Description	The initial connection to a VM that uses XenDesktop from a client with multiple displays results in a black screen. Disconnecting and reconnecting to the VM resolves the issue.
Fix	
Status	Fixed
Ref. #	1720471

XenDesktop shows only a black screen when connected to a vGPU VM	
Description	XenDesktop 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 the Microsoft power plans frequently asked questions at <a href="http://windows.microsoft.com/en-us/windows/power-plans-faq">http://windows.microsoft.com/en-us/windows/power-plans-faq</a> and from the list, select your OS version.
Status	Not an NVIDIA bug
Ref. #	1719877

## Notice

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

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

### HDMI

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

### OpenCL

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

### Trademarks

NVIDIA 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

© 2016 NVIDIA Corporation. All rights reserved.