



Release 331 Quadro, NVS, Tesla, & GRID Drivers for Windows - Version 331.40

RN-WQ33140-01_v01 | September 29, 2013
Windows Vista, Windows 7, Windows 8, & Windows 8.1

Release Notes



TABLE OF CONTENTS

1	Introduction to Release Notes	1
	Structure of the Document	1
	Changes in this Edition	1
2	Release 331 Driver Changes	2
	Version 331.40 Highlights	3
	Existing Support	3
	What's New in Release 331	3
	What's New in Version 331.40	4
	Discontinued and Unsupported Features in this Release	4
	Advanced Instructions for this Release	6
	Using Mosaic Technology	6
	Turning Off V-Sync to Boost Performance	7
	NVIDIA Application Configuration Engine (ACE)	7
	Using the WDDM Driver Model with Tesla GPU GOMs	8
	SLI Multi-OS - GPU Assignment in System Virtualization	10
	Open Issues in Version 331.40	13
	Windows Vista/Windows 7 32-bit Issues	13
	Windows Vista/Windows 7 64-bit Issues	13
	Windows 8 64-bit Issues	14
	Not NVIDIA Issues	15
	Windows Vista Considerations	15
	Windows 7 Considerations	15
	Unsupported Features	16
	OpenGL Application Issues	18
	Application Issues	18
	Other Issues	19
	Known Product Limitations	20
	System Should not Enter Sleep/Hibernate Mode With SDI Capture in Progress	20
	Once Installed, the SDI Capture Driver Release Must Always Match the Graphics Driver	21
	Some APIs do not Report Total Available Graphics Memory Correctly	21
	Using HDMI/DisplayPort Audio with Displays that have a High Native	

TABLE OF CONTENTS

Resolution	23
Using HDMI/DisplayPort Displays that do not Support Audio	23
Using HDMI/DisplayPort Audio in Dualview or Clone Mode Configurations	24
GPU Runs at a High Performance Level (full clock speeds) in Multi-display Modes	24
Aero Must be Enabled for Windowed SLI AFR Mode Under Vista	24
SLI Connector Requirement on NVIDIA Quadro SLI Cards	25
Applying Workstation Application Profiles	25
1280x1024 @ 60 Hz not Available on BenQ FP241W Monitors	25
Gigabyte GA-6BX Motherboard	25
3 The Release 331 Driver	26
Hardware and Software Support	26
Supported Operating Systems	26
Supported NVIDIA Products	27
Supported Languages	33
Driver Installation	34
Minimum Hard Disk Space	34
Before You Begin	34
Installation Instructions	34
Appendix A: Mode Support for Windows	35
General Mode Support Information	36
Default Modes Supported by GPU	37
Understanding the Mode Format	37
Tesla, Quadro, & NVS GPUs	38
Modes Supported by TV Encoders	40
Appendix B: NVIDIA Tesla Compute Cluster Mode	41
About Tesla Compute Cluster Mode	41
TCC Overview	41
Running CUDA Applications	42
Operating on Systems with non-TCC NVIDIA GPUs	43
Setting TCC Mode for Tesla Products	43

LIST OF TABLES

Table 2.1	Quadro Products Offering NVIDIA Mosaic	6
Table 1	GOM and Driver Model Compatibility	9
Table 3.1	Supported NVIDIA Quadro & NVS Products	27
Table 3.2	Supported NVIDIA Quadro Sync Products	29
Table 3.3	Supported NVIDIA Quadro G-Sync II Products	29
Table 3.4	Supported NVIDIA Quadro Plex Products.....	29
Table 3.5	Supported NVIDIA Quadro Blade/Embedded Graphics Board Series	29
Table 3.6	Supported NVIDIA Tesla Products	31
Table 3.7	Supported NVIDIA GRID Products.....	32
Table A.1	Modes Supported for High Resolution Displays	36
Table A.2	Non-standard Modes Supported	36
Table A.3	Mode Support for S-Video and Composite Out	40
Table A.4	Mode Support for Component YPrPb Out and DVI Out.....	40
Table B.1	Default TCC Mode for Tesla Products	43

01 INTRODUCTION TO RELEASE NOTES

This edition of *Release Notes* describes the Release 331 family of NVIDIA Quadro, NVS, Tesla, and GRID Drivers for Microsoft® Windows® Vista and later¹. NVIDIA provides these notes to describe performance improvements and bug fixes in each documented version of the driver.

Structure of the Document

This document is organized in the following sections:

- ▶ “Release 331 Driver Changes” on page 2 gives a summary of changes, and fixed and open issues in this version.
- ▶ “The Release 331 Driver” on page 26 describes the NVIDIA products and languages supported by this driver, the system requirements, and how to install the driver.
- ▶ “Mode Support for Windows” on page 35 lists the default resolutions supported by the driver.
- ▶ “NVIDIA Tesla Compute Cluster Mode” on page 41 describes the Tesla Compute Cluster mode.

Changes in this Edition

This edition of the *Release Notes* for Windows Vista and later includes information about NVIDIA graphics driver version 331.40, and lists changes made to the driver since version 327.23. These changes are discussed beginning with the chapter “Release 331 Driver Changes” on page 2.

1. Includes Windows Vista, Windows 7, Windows 8, and Windows 8.1.

02 RELEASE 331 DRIVER CHANGES

This chapter describes open issues for version 331.40, and resolved issues and driver enhancements for versions of the Release 331 driver up to version 331.40. The chapter contains these sections:

- ▶ “Version 331.40 Highlights” on page 3
- ▶ “Advanced Instructions for this Release” on page 6
- ▶ “Open Issues in Version 331.40” on page 13
- ▶ “Not NVIDIA Issues” on page 15
- ▶ “Known Product Limitations” on page 20

Version 331.40 Highlights

This section provides highlights of version 331.40 of the NVIDIA Release 331 Driver for Windows Vista and later.

- ▶ Existing Support
- ▶ What's New in Release 331
- ▶ What's New in Version 331.40
- ▶ Discontinued and Unsupported Features in this Release
- ▶ Limitations in This Release

Existing Support

This release supports the following APIs:

- ▶ Open Computing Language (OpenCL) 1.1 in Quadro FX Series x700 and newer as well as the FX4600 and FX5600.
- ▶ OpenGL 4.4
- ▶ DirectX 11
- ▶ CUDA 5.5

What's New in Release 331

The section summarizes the following driver changes in Release 331 since Release 325:

Mosaic

- ▶ Improved Mosaic setup time
- ▶ Added default topology and display numbering in the NVIDIA Control Panel->Mosaic setup process.

NvWMI

- ▶ Updated to version 2.15

OpenGL

- ▶ Added support for OpenGL 4.4

What's New in Version 331.40

- ▶ This driver offers performance improvements over previous driver versions, including workstation compatibility fixes.
- ▶ This driver installs **nView Desktop Manager** version 140.75.

Discontinued and Unsupported Features in this Release

Discontinued Features

- ▶ The NVIDIA® AutoCAD Performance driver is no longer integrated in the graphics driver.

Standalone versions or version updates can still be downloaded from the NVIDIA driver download page.

- ▶ The following features are removed from the NVIDIA Control Panel
 - The Views option

You no longer need to select between Standard and Advanced views for many NVIDIA Control Panel controls.
 - The Profiles menu
- ▶ Support for Quadro SDI products is discontinued for Windows 8 and later operating systems.
- ▶ Legacy Support for Curie generation of Workstation products

Beginning with Release 310, the NVIDIA professional drivers no longer support the Curie generation of Workstation products.

NVIDIA Quadro FX 5500
NVIDIA Quadro FX 4500 X2
NVIDIA Quadro FX 4500
NVIDIA Quadro FX 4400
NVIDIA Quadro FX 4400G
NVIDIA Quadro FX 4000
NVIDIA Quadro FX 470
NVIDIA Quadro FX 350
NVIDIA Quadro VX 200
NVIDIA Quadro NVS 440
NVIDIA Quadro NVS 285

Release 304 drivers continue to support Curie generation Workstation products, and NVIDIA will continue to address driver issues for these products in driver branches up to and including Release 304. However, future driver enhancements and

optimizations in driver releases after Release 304 will not support Curie generation products.

Limitations in This Release

The following are features that are not currently supported or have limited support in this driver release:

► Video Memory Support

For Windows 7 and Windows Vista 64-bit, this driver recognizes up to the total available video memory on Quadro cards for Direct3D and OpenGL applications.

For Windows 7 and Windows Vista 32-bit, this driver recognizes only up to 4 GB of video memory on Quadro cards for DirectX, OpenGL, and CUDA applications.

► NVIDIA Control Panel Display Category

- The Graph tab on the Adjust Desktop Color Settings page is not available.

Advanced Instructions for this Release

This section clarifies instructions for successfully accomplishing the following tasks:

- ▶ Using Mosaic Technology
- ▶ Turning Off V-Sync to Boost Performance
- ▶ NVIDIA Application Configuration Engine (ACE)
- ▶ Using the WDDM Driver Model with Tesla GPU GOMs
- ▶ SLI Multi-OS – GPU Assignment in System Virtualization

Using Mosaic Technology

This release contains NVIDIA's Mosaic Technology, which lets you span 2-4 displays across a single desktop for select non-SLI Quadro products and configurations. The following table lists the products and configurations that support NVIDIA Mosaic in this release:

Table 2.1 Quadro Products Offering NVIDIA Mosaic

Products	Notes
Quadro 6000, 5000	SLI GPUs in a Quadro SLI-certified workstation will utilize Premium Mosaic automatically.
Quadro 4000, 2000, 600	
Quadro FX 5800	SLI GPUs in a Quadro SLI-certified workstation will utilize Premium Mosaic automatically.
Quadro FX 4800, 3800, 1800, 580, 380LP, 380	
Quadro NVS 450, 420, 295, 290	
NVS 300	

- ▶ NVIDIA Mosaic is available only under Windows 7 and later Windows operating systems.

Turning Off V-Sync to Boost Performance

To get the best benchmark and application performance measurements, turn V-Sync off as follows:

- 1 Open the NVIDIA Control Panel and make sure that *Advanced Settings* is selected from the control panel tool bar.
- 2 From the *Select a Task* pane, under 3D Settings, click **Manage 3D Settings**, then click the Global Settings tab.
- 3 From the Global presets pull-down menu, select **Base profile**.
- 4 From the Settings list box, select **Vertical sync** and change its value to **Force off**, then click **Apply**.
- 5 From the Global presets pull-down menu, select **3D App - Default Global Settings** (the driver's default profile) or use the application profile that matches the application you are testing, then click **Apply**.

Be sure to close the NVIDIA Control Panel completely —leaving it open will affect benchmark and application performance.

NVIDIA Application Configuration Engine (ACE)

This driver includes the NVIDIA Application Configuration Engine (ACE), which automatically detects the workstation application and configures the appropriate profile settings in the NVIDIA Control Panel.

See the *NVIDIA Quadro Professional Drivers: NVIDIA Control Panel Quick Start Guide* for more information about this feature.

Using the WDDM Driver Model with Tesla GPU GOMs

Tesla GPU Operation Modes

Tesla K20X and K20 GPU Accelerators offer the ability to specify the GPU operation mode using NV-SMI/NVML. (refer to: <https://developer.nvidia.com/nvidia-management-library-nvml>).

By setting the GPU operation mode, developers can selectively turn off certain features in the GPU to get the best performance per watt for certain workloads.

The following are the supported GOMs:

- ▶ **Compute-Only:** For running compute tasks only.
By default the Tesla K20X and K20 accelerator boards ship in this mode.
- ▶ **Low-Double Precision:** For graphics applications that don't require high bandwidth double precision.
This is recommended for workloads that are not sensitive to double precision but at the same time need graphics capabilities.
- ▶ **All On:** This is recommended only when the workload needs full double precision as well as graphics capabilities.

WDDM and TCC Driver Models

Along with the GPU operation mode, the developer needs to select the compatible driver model for the Tesla K20X and K20.

- ▶ **Tesla Compute Cluster (TCC):** Optimized for running compute workloads.
- ▶ **Windows Device Driver Model (WDDM):** Designed for graphics application and not recommended for compute workloads.

Compatibility Between GOM and Driver Models

Table 1 shows which GPU operation modes are compatible with which driver models.

Table 1 GOM and Driver Model Compatibility

GOM	TCC Driver Model	WDDM Driver Model	Use Case Support
All On	YES	YES	All use cases are supported.
Compute-Only	YES	NO	The following are unsupported : X11 and those that require X11 (GLInterop, OCL conformance and VIPER) 32-bit Windows OS
Low Double Precision	YES	YES	All use cases supported.

The compute-only GOM is supported only on the TCC driver model, while the WDDM driver model supports only GOM modes that enable graphics.

The compute-only GOM and WDDM are incompatible and should not be used simultaneously.

The Tesla K20 Active Accelerators for workstations ship in “compute-only” mode and cannot be modified. Therefore, use only the TCC driver model with these products.

Using the WDDM Driver Model

To use the WDDM driver model with Tesla K20x and K20 GPU Accelerators for servers, first switch the GOM mode from compute-only to All On, then switch from TCC to WDDM.

Do not attempt to specify the driver model by editing the registry. Doing so can result in compute-only GOM and WDDM being configured simultaneously, which might require a clean installation of the driver to fix.

Always use NVIDIA-provided tools to specify a processing mode or to switch between driver models. Such tools include nvidia-smi or the NVIDIA Control Panel->Manage Maximus Settings page. These tools provide warnings in the case of a conflict.

SLI Multi-OS - GPU Assignment in System Virtualization

On systems with two or more graphics cards installed, this driver supports a hypervisor's ability to directly assign GPUs to guest virtual machines (VMs). This direct assignment allows each guest VM to run on their own operating system with their own GPU and driver. The assignment allows full GPU performance and functionality in the guest VM.

Hardware Platform Requirements

To make use of GPU passthrough with virtual machines running Windows and Linux, the hardware platform must support the following features:

- ▶ A CPU with hardware-assisted instruction set virtualization: Intel VT-x or AMD-V.
- ▶ Platform support for I/O DMA remapping.

On Intel platforms the DMA remapper technology is called Intel VT-d.

On AMD platforms it is called AMD IOMMU.

Support for these feature varies by processor family, product, and system, and should be verified at the manufacturer's website.

Supported Hypervisors

The following hypervisors are supported:

Hypervisor	Notes
Citrix XenServer	Version 6.0 and later.
VMware vSphere (ESX / ESXi)	Version 5.1 and later.
Parallels Workstation Extreme	Version 4 and later

Supported Graphics Cards

The following GPUs are supported for device passthrough:

GPU Family	Boards supported
Kepler	<u>GRID</u> : K1, K2, K520, K340 <u>Quadro</u> : K2000, K4000, K5000, K6000 <u>Tesla</u> : K10, K20
Fermi	<u>Quadro</u> : 2000, 4000, 5000, 6000 <u>Quadro-MXM</u> : 1000M, 3000M <u>Tesla</u> : C2050, C2075, M2050, M2070, M2070Q

GPU Family	Boards supported
Tesla	<u>Quadro</u> FX1800, 3800, 4800, 5800 <u>Quadro-MXM</u> : FX880M, FX2800M <u>Tesla</u> : M1060, C1060

Notes and Known Issues

VMware

- PCI I/O hole may need to be changed for Windows 64-bit VMs.
Windows 64-bit VMs may require that you edit the VM configuration file to configure a larger PCI I/O hole for the GPU.
- Access Control Services is required on some switches.
Starting with ESX 5.0 Update 1, Access Control Services (ACS) is required on any switches in the PCIe hierarchy above a PCIe device that is to be used for passthrough. If ACS is not present, ESX will not allow the device to be assigned directly to a VM.
To allow assignment of devices below switches that do not support ACS, you can disable the ACS check as follows:

```
esxcfg-advcfg -k true disableACSCheck
```
- Configuring passthrough to a Windows Server 2008R2 VM.
The VMware WDDM driver must be manually installed on Windows Server 2008 R2.
Follow the steps provided at <http://communities.vmware.com/message/1423263#1425288>.
- MSI translation must be disabled.
Blue-screen crashes may occur on VMs with assigned GPUs if MSI is initially enabled for passthrough devices.
To avoid this issue, disable MSI translation by setting `pciPassthru0.msiEnabled = "FALSE"` in the VM's VMX file.

Citrix Xenserver

- XenServer 6.1 workaround for 64-bit MMIO failures on PCI passthrough.
Dom0 kernel may relocate GPU BARs to an invalid location immediately beyond the end of physical RAM. There are two workaround options to address this condition:
Workaround option #1: Add 'pci=use_crs' to the kernel parameters by running the following command in Dom0:

```
/opt/xensource/libexec/xen-cmdline --set-dom0 pci=use_crs
```

Workaround option #2: Install the hotfix provided by Citrix at <http://support.citrix.com/article/CTX137645>):

Multi-Monitor Support on GRID

Multi-monitor support on GRID boards K1 and K2 requires the following VBIOS versions:

- GRID K1: **80.07.AF.00.00** or later
- GRID K2: **80:04:BA:00.00** or later

Open Issues in Version 331.40

As with every released driver, version 331.40 of the Release 331 driver has open issues and enhancement requests associated with it. This section includes lists of issues that are either not fixed or not implemented in this version. Some problems listed may not have been thoroughly investigated and, in fact, may not be NVIDIA issues. Others may have workaround solutions.

- ▶ [“Windows Vista/Windows 7 32-bit Issues”](#) on page 13
- ▶ [“Windows Vista/Windows 7 64-bit Issues”](#) on page 13

Windows Vista/Windows 7 32-bit Issues

- ▶ Cinema 4D—performance improvements are requested.

Windows Vista/Windows 7 64-bit Issues

- ▶ [Cinema 4D]: Performance improvements are requested.
- ▶ [Eyeon Fusion]: With Aero/desktop compositing enabled, the application Display Views flickers when dragged from partially off-screen back to on-screen.
- ▶ [Eyeon Fusion]: With Aero/desktop compositing and Quad-buffer stereo enabled, the application child windows flicker.
- ▶ [Kepler-class GPUs][Mudbox 2014]: Shadows are not visible.
- ▶ [Kepler-class GPUs]: 10-bit support doesn't work if a 10-bit capable monitor is connected to DVI when the computer boots, and then later transferred to the DisplayPort connector.
- ▶ [Quadro Plex 7000][G Sync]: With a projector cluster setup where each projector is driven by one GPU, p-state transitions cause one of the channels to go blank.
- ▶ [Quadro K5000][Eyeon Fusion][Quad-buffered stereo]: Blue-screen crash occurred pointing to `nvlddmkm.sys`.
- ▶ [Quadro 6000]: `clCreateFromGLTexture2D` fails with Mipmaps.
- ▶ [Quadro 6000][RTTScale]: `GLCache` directory causes the application to crash upon launch.
- ▶ Quadro 6000: Performance drop in octree program occurs with atomic counters.
- ▶ Quadro 6000/5000: `clCreateFromGLBuffer` returns `CL_OUT_OF_HOST_MEMORY`.
- ▶ [Quadro 4000/K5000]: Rendering fails with `GL_INVALID_OPERATION` error in the draw call.
- ▶ [Quadro FX 1800M]: The `clCreateContext` operation fails.

Windows 8 64-bit Issues

- ▶ [Kepler-class GPUs][Mudbox 2014]: Shadows are not visible.

Not NVIDIA Issues

This section lists issues that are not due to the NVIDIA driver as well as features that are not meant to be supported by the NVIDIA driver for Windows Vista/Windows 7.

- ▶ “Windows Vista Considerations” on page 15
- ▶ “Windows 7 Considerations” on page 15
- ▶ “Unsupported Features” on page 16
- ▶ “OpenGL Application Issues” on page 18
- ▶ “OpenGL Application Issues” on page 18
- ▶ “Application Issues” on page 18
- ▶ “Other Issues” on page 19

Windows Vista Considerations

These are behaviors that may be different from Windows XP and are related directly to the Windows Vista operating system.

- ▶ **Gamma ramps are inconsistent between single and two-headed systems.**
- ▶ **NVIDIA TurboCache**

Windows Vista now controls the allocation of system memory to the GPU for TurboCache functions. The Windows Vista Display Properties pages show the shared system memory (SSM), or how much memory is allocated for NVIDIA GPUs to use for TurboCache.

For more information on graphics memory reporting under Windows Vista, visit <http://www.microsoft.com/whdc/device/display/graphicsmemory.mspx>.

Windows 7 Considerations

Windows DWM Disabled for SLI Mosaic Mode

Due to compatibility issues, when using SLI Mosaic mode the driver turns off the Windows 7 Desktop Window Manager (DWM). As a result, DWM-managed desktop features such as Windows Aero or Windows Flip 3D will not be available.

Hotplug Action

Unlike the hotplug activity under Windows Vista, the default settings are not applied when a new display is hotplugged, and there is no message balloon alert stating that a new display was detected. Under Windows 7, all display connection and detection events are handled through the Windows 7 Connecting and Configuring Displays (CCD) mechanism.

NVIDIA Control Panel Rotate Display Page

The rotation radio button labels are changed slightly under Windows 7 to be consistent with the Microsoft panel

Table 2.1 NVIDIA Control Panel Rotation Page Radio Buttons

Clockwise Rotation	Windows 7 Label	Windows Vista Label
0 degrees	Landscape	No rotation (Landscape)
90 degrees	Portrait	90 degrees to the right (Inverted Portrait)
180 degrees	Landscape (flipped)	180 degree rotation (Inverted landscape)
270 degrees	Portrait (flipped)	90 degrees to the left (Portrait)

Limitation

- ▶ When switching the refresh rate from 59 Hz to 60Hz, the refresh rate remains at 59 Hz.
See the Microsoft KB article KB2006076 at <http://support.microsoft.com/kb/2006076>.

Unsupported Features

The following are features and functionality that were available in driver releases supporting Windows XP, but are not—and will not be—available in driver releases for Windows Vista:

- ▶ **High resolution scaling desktop (HRSD)**
- ▶ **MultiView Display Mode** (for NVIDIA Quadro NVS graphics cards)
- ▶ **NVKeystone**
- ▶ **Unified back buffer (UBB) controls**
- ▶ **OpenGL Video Overlays**
This is an operating system limitation.
- ▶ **Overclocking**
GPU overclocking is no longer supported in the default GPU driver control panel. This feature is available in the NVIDIA System Tools software, which you can download from NVIDIA.com.
- ▶ **GPU Temperature Monitoring**
Temperature monitoring is no longer supported in the default GPU driver control panel. This feature is available in the NVIDIA System Tools software, which you can download from NVIDIA.com.
- ▶ **AGP Settings Adjustment**

- ▶ **Video Zoom**
- ▶ **Pan & Scan** - the process of panning across the desktop in order to display a desktop on a monitor with lower resolution
- ▶ **Per-display Desktop Color Setting Adjustments**

For Clone mode, the desktop color setting adjustments through the NVIDIA Control Panel can only be made across all displays in a system, and not on a per-display basis.
- ▶ **Per-display Video Color Setting Adjustments**

For Dualview mode, the video color setting adjustments through the NVIDIA Control Panel can only be made across all displays in a system, and not on a per-display basis.
- ▶ **Edge Blending**
- ▶ **Run display optimization wizard**
- ▶ **Run multiple display wizard**
- ▶ **Run television setup wizard**
- ▶ **nView Horizontal and Vertical Span Modes**

Due to architectural changes in the new Windows Vista Window Display Driver Model (WDDM), span mode is available only with NVIDIA Mosaic Technology.
- ▶ **Display/Connection Wizard** (such as was provided with Windows Media Center Edition)
- ▶ **DVD/MPEG Extensions** (such as was provided with Windows Media Center Edition)
- ▶ **Audio Extensions** (such as was provided with Windows Media Center Edition)

OpenGL Application Issues

The following are known compatibility issues for OpenGL applications developed under Windows XP:

- ▶ Mixed GDI and OpenGL rendering does not work.
 - A number of applications use GDI to render UI components and object highlighting. This is not supported in the Windows Vista driver model.
 - NVIDIA recommends converting GDI rendering to OpenGL.
 - The following are some applications that are known to have this issue:
 - Maya 7.01
- ▶ Applications, Tools, and Benchmarks not Supported Under Windows Vista
 - GLperf
 - 3ds max 8 (later releases may be supported)
 - CATIA V5R15 (V5R16 is supported)
 - PTC's CDRS 2001
- ▶ Front buffered rendering may be slow, especially when DWM is enabled.
 - Flushing the rendering queue while rendering to the front buffer may cause the window manager to recomposite. Applications should therefore minimize the frequency with which they flush the rendering queue.

Application Issues

- ▶ Softimage–The application crashes when thumbing the CgFX scene model while in wireframe display mode.
- ▶ Solidworks 2009–Application profile is not shown in the NVIDIA Control Panel when SolidWorks 2009 is installed.
 - This is an issue with the application shortcut.*
- ▶ ArchiCAD–the driver crashes when navigating 3D shadows.
- ▶ ArchiCAD12–OpenGL speed is half as fast on Windows Vista than on Windows XP.
- ▶ CATIA V5R20–not all drawing elements appear if the drawing is created using Approximate mode.

Other Issues

- ▶ Quadro 6000: OpenCL/OpenGL interoperability performance suffers with `clEnqueueReleaseGLObjects`.

To avoid this issues, use CUDA.

- ▶ The Windows Vista display mode switches from Aeroglass to Basic when a quad-buffer for stereo is created.

Quad-buffered windowed stereo is only supported with Aeroglass turned off.

- ▶ The NVIDIA Control Panel->Set Up Multiple Displays page does not provide the capability of setting the dual monitor order under Windows Vista as it does under Windows XP.

This capability is provided through the Windows Vista Display Properties Settings page.

Known Product Limitations

This section describes problems that will not be fixed. Usually, the source of the problem is beyond the control of NVIDIA. Following is the list of problems and where they are discussed in this document:

- ▶ “System Should not Enter Sleep/Hibernate Mode With SDI Capture in Progress” on page 20
- ▶ “Once Installed, the SDI Capture Driver Release Must Always Match the Graphics Driver” on page 21
- ▶ “Some APIs do not Report Total Available Graphics Memory Correctly” on page 21
- ▶ “Using HDMI/DisplayPort Audio with Displays that have a High Native Resolution” on page 23
- ▶ “Using HDMI/DisplayPort Displays that do not Support Audio” on page 23
- ▶ “Using HDMI/DisplayPort Audio in Dualview or Clone Mode Configurations” on page 24
- ▶ “GPU Runs at a High Performance Level (full clock speeds) in Multi-display Modes” on page 24
- ▶ “Aero Must be Enabled for Windowed SLI AFR Mode Under Vista” on page 24
- ▶ “SLI Connector Requirement on NVIDIA Quadro SLI Cards” on page 25
- ▶ “Applying Workstation Application Profiles” on page 25
- ▶ “1280x1024 @ 60 Hz not Available on BenQ FP241W Monitors” on page 25
- ▶ “Gigabyte GA-6BX Motherboard” on page 25

System Should not Enter Sleep/Hibernate Mode With SDI Capture in Progress

SDI capture programs will become non-responsive upon resume from sleep or hibernate modes.

This is not a typical use case, so power management tests should not be performed when testing SDI.

Once Installed, the SDI Capture Driver Release Must Always Match the Graphics Driver

To use the Quadro SDI Capture card, you need to install the NVIDIA WDM driver for the Quadro SDI Capture card after installing the Quadro professional driver. The WDM driver version must be from the same branch as the installed NVIDIA graphics driver.

If at a later time you install a different Quadro professional driver version—for example, when upgrading from a Release 260 driver to a Release 265 driver—you must reinstall the WDM driver from the same driver branch as well—in this example, Release 265.

You must install the updated WDM driver even if you remove the SDI Capture card and uninstall the WDM driver. Portions of the old WDM driver remain on the system, and the resulting mismatch with the updated graphics driver results in a faulty driver installation. Installing a matching WDM driver corrects this issue.

Some APIs do not Report Total Available Graphics Memory Correctly

Background-TAG Memory

In the Windows Display Driver Model (WDDM), Total Available Graphics (TAG) memory is reported as the sum of

- Dedicated Video Memory (video memory dedicated for graphics use)
- Dedicated System Memory (system memory dedicated for graphics use), and
- Shared System Memory (system memory shared between the graphics subsystem and the CPU).

The values for each of these components are computed according to WDDM guidelines when the NVIDIA Display Driver is loaded.

Issue

NVIDIA has found that some TAG-reporting APIs represent video memory using 32-bits instead of 64-bits, and consequently do not properly report available graphics memory when the TAG would otherwise exceed 4 gigabytes (GB). This results in under reporting of available memory and potentially undesirable behavior of applications that rely on these APIs to report available memory.

The reported memory can be severely reduced. For example, 6 GB might be reported as 454 MB, and 8 GB might be reported as 1259 MB.

NVIDIA Action for Some GeForce-based Systems

For GeForce GPUs with 2.75 GB or less of video memory, the NVIDIA display driver constrains TAG memory to just below 4 GB¹. In this scenario, the Shared System Memory component of TAG is limited first, before limiting Dedicated Video Memory.

This is a policy decision within the driver, and results in reliable reporting of sub-4 GB TAG memory.

When TAG Reporting Would Not Be Limited

For GeForce-based GPUs with more than 2.75 GB of video memory, as well as all Quadro and Tesla GPUs, the NVIDIA display driver does not constrain TAG memory reporting.

The disadvantage of constraining TAG on systems with larger amounts of video and system memory is that memory which otherwise would be available for graphics use is no longer available. Since shared system memory is limited first, driver components and algorithms utilizing shared system memory may suffer performance degradation when TAG is constrained.

Since these and similar scenarios are prevalent in many Workstation applications, the NVIDIA driver avoids constraining TAG on all Quadro and Tesla-based systems. Likewise, the driver does not constrain TAG for GeForce-based systems with more than 2.75 GB of video memory.

1. The WDDM guidelines dictate minimum and maximum values for the components, but the display driver may further constrain the values that are reported (within the allowed minimum and maximum).

Using HDMI/DisplayPort Audio with Displays that have a High Native Resolution

To use HDMI/DisplayPort audio with some displays that have a native resolution higher than 1920x1080, you must set the display to a lower HD resolution.

Some HDMI TV's have a native resolution that exceeds the maximum supported HD mode. For example, TVs with a native resolution of 1920x1200 exceed the maximum supported HD mode of 1920x1080.

Applying this native mode results in display overscan which cannot be resized using the NVIDIA Control Panel since the mode is not an HD mode.

To avoid this situation and provide a better user experience, the driver treats certain TVs— such as the Viewsonic VX2835wm and the Westinghouse LVM- 37w3—as a DVI monitor when applying the native mode. Because the driver does not treat the TV as an HDMI in this case, the HDMI audio is not used.

Using HDMI/DisplayPort Displays that do not Support Audio

Some HDMI/DisplayPort displays do not support audio, or have issues with Quadro FX family and earlier NVIDIA graphics cards.

The NVIDIA driver attempts to identify such displays and automatically disables the audio. For example, the NVIDIA driver disables HDMI audio for all Philips HDMI TVs, as these have been identified as having issues with Quadro FX family and earlier NVIDIA graphics cards.

There may be cases where either the driver disables audio even though there is no problem, or does not disable the audio when in fact the audio does not work. The following sections describe these situations and provides guidance for handling them.

Corrupted video and no audio

The driver has not disabled audio and the display's audio signal is incompatible with the graphics card, causing video corruption.

With a different display connected in order to establish video, disable audio for the HDMI display using the NVIDIA Control Panel-> Change Resolution page. From the connector list, select **HDMI-HDTV (Audio Disabled)**.

Video but no audio

- ▶ Check the connector list on the NVIDIA Control Panel->Change Resolution page.
- ▶ If **HDMI-HDTV (Audio Disabled)** is selected and you want to test whether your HDMI audio does, in fact, work, then select **HDMI-HDTV (Audio Enabled)** and the driver will prompt you with instructions for testing HDMI audio with the display.
- ▶ If **HDMI-HDTV (Audio Enabled)** is selected, then the driver has not successfully detected that an incompatible display is connected.
Future drive versions will properly identify such displays and disable audio.
- ▶ If there is no HDMI connector option in the NVIDIA Control Panel->Change Resolution page, the display does not support audio and has properly reported this to the NVIDIA driver.

Using HDMI/DisplayPort Audio in Dualview or Clone Mode Configurations

Two Audio-enabled Ports

In a multi-display configuration where both HDMI/DisplayPort audio ports are enabled, only the primary display will provide the audio.

One Audio-enabled Port

In a multi-display configuration where only one audio port is enabled, such as when one display is a DVI display, then the HDMI/DisplayPort display can provide the audio whether is it the primary or secondary display.

GPU Runs at a High Performance Level (full clock speeds) in Multi-display Modes

This is a hardware limitation and not a software bug. Even when no 3D programs are running, the driver will operate the GPU at a high performance level in order to efficiently drive multiple displays. In the case of SLI or multi-GPU PCs, the second GPU will always operate with full clock speeds; again, in order to efficiently drive multiple displays. Today, all hardware from all GPU vendors have this limitation.

Aero Must be Enabled for Windowed SLI AFR Mode Under Vista

Windows 7 Aero must be enabled in order to achieve SLI acceleration using windowed AFR mode.

SLI Connector Requirement on NVIDIA Quadro SLI Cards

The SLI connector that links two SLI cards is needed for proper SLI operation. However, the connector can be removed if you do not intend to enable SLI mode. If you remove the connector, then you must make sure that SLI mode is disabled from the NVIDIA control panel. Enabling SLI mode without the SLI connector installed will result in video corruption.

Applying Workstation Application Profiles

► Background

The workstation application profiles are software settings used by the NVIDIA Display Drivers to provide optimum performance when using a selected application. The profile also works around known application issues and bugs.

If there is an available setting for an application, it should be used, otherwise incorrect behavior or reduced performance is likely to occur.

► Issues

Configuration changes require that you restart the application.

Once an application is running, it does not receive notification of configuration changes. Therefore, if you change the configuration while the application is running, you must exit and restart the application for the configuration changes to take effect.

1280x1024 @ 60 Hz not Available on BenQ FP241W Monitors

Even though the monitor EDID lists 1280x1024 @ 60 Hz, the screen turns blank when using an HDMI connection. This is an issue with the monitor and not the NVIDIA driver.

Because of this issue with the monitor, the NVIDIA driver blocks the problem mode (1280x1024 @ 60 Hz) and makes it unavailable.

Gigabyte GA-6BX Motherboard

This motherboard uses a LinFinity regulator on the 3.3-V rail that is rated to only 5 A—less than the AGP specification, which requires 6 A. When diagnostics or applications are running, the temperature of the regulator rises, causing the voltage to the NVIDIA chip to drop as low as 2.2 V. Under these circumstances, the regulator cannot supply the current on the 3.3-V rail that the NVIDIA chip requires.

This problem does not occur when the graphics board has a switching regulator or when an external power supply is connected to the 3.3-V rail.

03 THE RELEASE 331 DRIVER

This chapter covers the following main topics:

- ▶ “Hardware and Software Support” on page 26
- ▶ “Driver Installation” on page 34

Hardware and Software Support

Supported Operating Systems

The Release 331 driver, version 331.40, has been tested with

- ▶ Microsoft Windows® 8.1, and supports both 32-bit and 64-bit versions.
- ▶ Microsoft Windows® 8, and supports both 32-bit and 64-bit versions.
- ▶ Microsoft Windows® 7, and supports both 32-bit and 64-bit versions.
- ▶ Microsoft Windows® Vista, and supports both 32-bit and 64-bit versions.

Supported NVIDIA Products

The following tables list the NVIDIA products supported by the Release 331 driver, version 331.40.

- ▶ Table 3.1, “Supported NVIDIA Quadro & NVS Products” on page 27
- ▶ Table 3.2, “Supported NVIDIA Quadro Sync Products” on page 29
- ▶ Table 3.3, “Supported NVIDIA Quadro G-Sync II Products” on page 29
- ▶ Table 3.4, “Supported NVIDIA Quadro Plex Products” on page 29
- ▶ Table 3.5, “Supported NVIDIA Quadro Blade/Embedded Graphics Board Series” on page 29
- ▶ Table 3.6, “Supported NVIDIA Tesla Products” on page 31
- ▶ Table 3.7, “Supported NVIDIA GRID Products” on page 32

NVIDIA Quadro & NVS Product

Table 3.1 Supported NVIDIA Quadro & NVS Products

Product	Notes
NVIDIA Quadro K6000	
NVIDIA Quadro K5000	
NVIDIA Quadro K4000	
NVIDIA Quadro K2000	
NVIDIA Quadro K600	
NVIDIA Quadro 6000	
NVIDIA Quadro 5000	
NVIDIA Quadro 4000	
NVIDIA Quadro 2000D	
NVIDIA Quadro 2000	
NVIDIA Quadro 600	
NVIDIA Quadro 410	
NVIDIA Quadro 400	
NVIDIA Quadro FX 5800	
NVIDIA Quadro FX 5600	
NVIDIA Quadro FX 4800	
NVIDIA Quadro FX 4700 X2	
NVIDIA Quadro FX 4600	
NVIDIA Quadro FX 3800	
NVIDIA Quadro FX 3700	
NVIDIA Quadro FX 3450	

Table 3.1 Supported NVIDIA Quadro & NVS Products (continued)

Product	Notes
NVIDIA Quadro FX 1800	
NVIDIA Quadro FX 580	
NVIDIA Quadro FX 570	
NVIDIA Quadro FX 380	
NVIDIA Quadro FX 380 LP (low profile)	
NVIDIA Quadro FX 370	
NVIDIA Quadro FX 370 LP (low profile)	
NVIDIA Quadro CX	
NVIDIA NVS 510	
NVIDIA NVS 310	
NVIDIA NVS 300	
NVIDIA Quadro NVS 450	
NVIDIA Quadro NVS 420	
NVIDIA Quadro NVS 295	
NVIDIA Quadro NVS 290	

NVIDIA Quadro Sync Products

Table 3.2 Supported NVIDIA Quadro Sync Products

Product	Notes
NVIDIA Quadro K5000	

NVIDIA Quadro G-Sync II Products

Table 3.3 Supported NVIDIA Quadro G-Sync II Products

Product	Notes
NVIDIA Quadro 6000	
NVIDIA Quadro 5000	
NVIDIA Quadro FX 5800	
NVIDIA Quadro FX 4800	
NVIDIA Quadro FX 5600	
NVIDIA Quadro FX 4600	

NVIDIA Quadro Plex Products

Table 3.4 Supported NVIDIA Quadro Plex Products

Product	Notes
NVIDIA Quadro Plex 7000	
NVIDIA Quadro Plex D Series	
NVIDIA Quadro Plex S Series	
NVIDIA Quadro Plex Model IV	

NVIDIA Quadro Blade/Embedded Graphics Board

Table 3.5 Supported NVIDIA Quadro Blade/Embedded Graphics Board Series

Product	Notes
NVIDIA Quadro K5000M	
NVIDIA Quadro K4000M	
NVIDIA Quadro K3000M	
NVIDIA Quadro 4000M	
NVIDIA Quadro 3000M	
NVIDIA Quadro 1000M	

Table 3.5 Supported NVIDIA Quadro Blade/Embedded Graphics Board Series (continued)

Product	Notes
NVIDIA Quadro 500M	
NVIDIA Quadro FX 3600M	
NVIDIA Quadro FX 2800M	
NVIDIA Quadro FX 880M	
NVIDIA Quadro FX 770M	
NVIDIA Quadro FX 370M	

NVIDIA Tesla Products

The Tesla driver package is designed for systems that have one or more Tesla products installed.

- ▶ Only one GHIC can be connected to the server in a Tesla system.
- ▶ This release of the Tesla driver supports CUDA C/C++ applications and libraries that rely on the CUDA C Runtime and/or CUDA Driver API.

Table 3.6 Supported NVIDIA Tesla Products

Product	Notes
Tesla K-Series Products	
NVIDIA Tesla K20Xm	64-bit OS only
NVIDIA Tesla K20m	64-bit OS only
NVIDIA Tesla K20c	64-bit OS only
NVIDIA Tesla K10	
Tesla S-Class Products	
NVIDIA Tesla S2050	
NVIDIA Tesla S1070	
Tesla M-Class Products	
NVIDIA Tesla M2090	
NVIDIA Tesla M2075	
NVIDIA Tesla M2070	
NVIDIA Tesla M2070Q	
NVIDIA Tesla M2050	
NVIDIA Tesla M1060	
Tesla C-Class Products	
NVIDIA Tesla C2075	
NVIDIA Tesla C2070	
NVIDIA Tesla C2050	
NVIDIA Tesla C1060	
NVIDIA Tesla C870	

NVIDIA GRID Products

Table 3.7 Supported NVIDIA GRID Products

Product	Notes
NVIDIA GRID K2	Windows 7 Windows 8
NVIDIA GRID K1	Windows 7 Windows 8

Supported Languages

The Release 331 Graphics Drivers supports the following languages in the main driver Control Panel:

English (USA)	German	Portuguese (Euro/ Iberian)
English (UK)	Greek	Russian
Arabic	Hebrew	Slovak
Chinese (Simplified)	Hungarian	Slovenian
Chinese (Traditional)	Italian	Spanish
Czech	Japanese	Spanish (Latin America)
Danish	Korean	Swedish
Dutch	Norwegian	Thai
Finnish	Polish	Turkish
French	Portuguese (Brazil)	

Driver Installation

Minimum Hard Disk Space

The hard disk space requirement for 32-bit is minimum 220 MB for English-only, and 300 MB for International.

The hard disk space requirement for 64-bit is minimum 320 MB for English-only, and 400 MB for International.

Before You Begin

nTune

If you have previously installed NVIDIA nTune, NVIDIA recommends that you uninstall nTune before installing this driver. After the driver install is complete, you can reinstall NVIDIA nTune.

SLI Mosaic Mode

You must make sure SLI Mosaic mode is disabled before installing a new driver over a previously installed driver. If SLI Mosaic mode is active on your displays when you install the new driver, the driver will not install properly.

Installation Instructions

- 1 Follow the instructions on the NVIDIA .com Web site driver download page to locate the appropriate driver to download, based on your hardware and operating system.
- 2 Click the driver download link.
The license agreement dialog box appears.
- 3 Click **Accept** if you accept the terms of the agreement, then either open the file or save the file to your PC and open it later.
- 4 Open the NVIDIA driver installation .EXE file to launch the NVIDIA InstallShield Wizard.
- 5 Follow the instructions in the NVIDIA InstallShield Wizard to complete the installation.



Note: If you are overinstalling the driver (installing over a previous driver without first removing the previous driver), then you must reboot your computer in order to complete the installation.

APPENDIX A MODE SUPPORT FOR WINDOWS

This chapter details the Windows modes supported by the Release 331 driver for NVIDIA products. It contains these sections:

- ▶ “General Mode Support Information” on page 36
- ▶ “Default Modes Supported by GPU” on page 37
- ▶ “Modes Supported by TV Encoders” on page 40

General Mode Support Information

The NVIDIA graphics driver includes a standard list of display modes that are supported by default. These modes are listed in the section “Default Modes Supported by GPU” on page 37.

The actual modes available depend on the capabilities of the display. In addition, the NVIDIA graphics driver has a “dynamic EDID detection” capability and will make available *additional* modes that are listed in the display EDID, provided the graphics hardware can support it.

The NVIDIA graphics driver also supports the high resolutions available with the displays listed in Table A.1 as well as the non-standard modes listed in Table A.2.

Table A.1 Modes Supported for High Resolution Displays

Display	Maximum Resolution
Apple 30" Cinema HD Display (Dual link DVI)	2560x1600 @ 60 Hz
Dell WFP 3007 (Dual Link DVI)	2560x1600 @ 60 Hz
HP LP3065 dual-link DVI flat panel	2560x1600 @ 60Hz.

Table A.2 Non-standard Modes Supported

Resolution		
1680 x 1050		
1366 x 768		

Default Modes Supported by GPU

This section lists the modes that are included by default in the driver INF for the following product families:

► “Tesla, Quadro, & NVS GPUs” on page 38

Understanding the Mode Format

Figure A.1 gives an example of how to read the mode information presented in this section.




Figure A.1 Mode Format

Note:

- Horizontal spanning modes of 3840x1080 and above, and vertical spanning modes of 1920x2160 and above generally require at least 32 MB of video memory at 32 bpp.
- An “i” next to the refresh rate indicates an interlaced refresh rate.

Tesla, Quadro, & NVS GPUs

This sections lists the supported display resolutions, color depths, and refresh rates for the products listed in [“Supported NVIDIA Products”](#) on page 27.

Standard Modes

640 x 480	8	60 70 72 75 85 100 120 140 144 150 170 200 240
720 x 480	8	60
720 x 576	8	50
800 x 600	8	60 70 72 75 85 100 120 140 144 150 170 200 240
1024 x 768	8	60 70 72 75 85 100 120 140 144 150 170 200 240
1152 x 864	8	60 70 72 75 85 100 120 140 144 150 170 200
1280 x 720	8	60
1280 x 768	8	60 70 72 75 85 100 120 140 144 150 170
1280 x 800	8	60 70 72 75 85 100 120 140 144 150 170
1280 x 960	8	60 70 72 75 85 100 120 140 144 150 170
1280 x 1024	8	60 70 72 75 85 100 120 140 144 150 170
1360 x 768	8	60 70 72 75 85 100 120 140 144 150 170
1600 x 900	8	60 70 72 75 85 100 120 140 144 150
1600 x 1024	8	60 70 72 75 85 100 120
1600 x 1200	8	60 70 72 75 85 100 120
1680 x 1050	8	60
1920 x 1080	8	60
1920 x 1200	8	60 70 72 75 85 100
1920 x 1440	8	60 70 72 75 85
2048 x 1536	8	60

640 x 480	16	60 70 72 75 85 100 120 140 144 150 170 200 240
720 x 480	16	60
720 x 576	16	50
800 x 600	16	60 70 72 75 85 100 120 140 144 150 170 200 240
1024 x 768	16	60 70 72 75 85 100 120 140 144 150 170 200 240
1152 x 864	16	60 70 72 75 85 100 120 140 144 150 170 200
1280 x 720	16	60
1280 x 768	16	60 70 72 75 85 100 120 140 144 150 170
1280 x 800	16	60 70 72 75 85 100 120 140 144 150 170
1280 x 960	16	60 70 72 75 85 100 120 140 144 150 170
1280 x 1024	16	60 70 72 75 85 100 120 140 144 150 170
1360 x 768	16	60 70 72 75 85 100 120 140 144 150 170
1600 x 900	16	60 70 72 75 85 100 120 140 144 150
1600 x 1024	16	60 70 72 75 85 100 120
1600 x 1200	16	60 70 72 75 85 100 120
1680 x 1050	16	60
1920 x 1080	16	60
1920 x 1200	16	60 70 72 75 85 100

```

1920 x 1440 16      60 70 72 75 85
2048 x 1536 16      60
-----
 640 x  480 32      60 70 72 75 85 100 120 140 144 150 170 200 240
 720 x  480 32      60
 720 x  576 32      50
 800 x  600 32      60 70 72 75 85 100 120 140 144 150 170 200 240
1024 x  768 32      60 70 72 75 85 100 120 140 144 150 170 200 240
1152 x  864 32      60 70 72 75 85 100 120 140 144 150 170 200
1280 x  720 32      60
1280 x  768 32      60 70 72 75 85 100 120 140 144 150 170
1280 x  800 32      60 70 72 75 85 100 120 140 144 150 170
1280 x  960 32      60 70 72 75 85 100 120 140 144 150 170
1280 x 1024 32      60 70 72 75 85 100 120 140 144 150 170
1360 x  768 32      60 70 72 75 85 100 120 140 144 150 170
1600 x  900 32      60 70 72 75 85 100 120 140 144 150
1600 x 1024 32      60 70 72 75 85 100 120
1600 x 1200 32      60 70 72 75 85 100 120
1680 x 1050 32      60
1920 x 1080 32      60
1920 x 1200 32      60 70 72 75 85 100
1920 x 1440 32      60 70 72 75 85
2048 x 1536 32      60
-----

```

Modes Supported by TV Encoders

Table A.3 and Table A.4 list the NTSC, PAL, and HDTV TV-Out modes supported by the NVIDIA driver.

Table A.3 Mode Support for S-Video and Composite Out

Resolution	Bit depth	Comments
320x200	8, 16, 32	DirectDraw mode; not selectable as a Windows desktop
320x240	8, 16, 32	DirectDraw mode; not selectable as a Windows desktop
640x400	8, 16, 32	DirectDraw mode; not selectable as a Windows desktop
640x480	8, 16, 32	
720x480	8, 16, 32	Overscans (for video)
720x576	8, 16, 32	Overscans (for video)
800x600	8, 16, 32	
1024x768	8, 16, 32	Conexant 25871 only

Table A.4 Mode Support for Component YPrPb Out and DVI Out

Resolution	Comments
480i (SDTV)	Supported on graphics boards with Conexant 875 or Philips 7108 TV encoders and compatible connectors, and compatible GeForce 8 Series and later GPUs.
480p (EDTV)	
720p (HDTV)	
1080i (HDTV)	
576i (PAL)	
576p (PAL)	

The driver supports manual overscan correction for component and DVI outputs. See the online NVIDIA Control Panel Help for instructions on how to use the overscan correction features.

APPENDIX B NVIDIA TESLA COMPUTE CLUSTER MODE

This chapter describes the Tesla Compute Cluster (TCC) mode.

- ▶ “About Tesla Compute Cluster Mode” on page 41
- ▶ “Operating on Systems with non-TCC NVIDIA GPUs” on page 43
- ▶ “Setting TCC Mode for Tesla Products” on page 43

About Tesla Compute Cluster Mode

TCC Overview

Tesla Compute Cluster (TCC) mode is designed for compute cluster nodes that have one or more Tesla products installed.

Benefits

- ▶ TCC drivers make it possible to use NVIDIA GPUs in nodes with non-NVIDIA integrated graphics.
- ▶ NVIDIA GPUs on systems running the TCC drivers will be available via Remote Desktop, both directly and via cluster management systems that rely on Remote Desktop.
- ▶ NVIDIA GPUs will be available to applications running as a Windows service (i.e. in Session 0) on systems running the Tesla driver in TCC mode.

TCC Does not Support Graphics Acceleration

- ▶ TCC mode does not provide CUDA–DirectX/OpenGL interoperability.

It is a “non-display” driver, and NVIDIA GPUs using this driver will not support DirectX or OpenGL hardware acceleration.

Running CUDA Applications

- ▶ This release of the Tesla driver supports CUDA C/C++ applications and libraries that rely on the CUDA C Runtime and/or CUDA Driver API.
- ▶ NVIDIA GPUs running the Tesla driver in TCC mode will be available for CUDA applications running via services or Remote Desktop.
- ▶ In this release, all GPUs will be in compute exclusive mode. As a result, only one CUDA context may exist on a particular device at a time.
- ▶ SDK applications that use graphics will not run properly under TCC mode. The following are examples of CUDA SDK applications that are not supported:

bicubicTexture	boxFilter	cudaDecodeD3D9	smokeParticles
cudaDecodeGL	fluidsD3D9	fluidsGL	SobelFilter
imageDenoising	Mandelbrot	marchingCubes	volumeRender
nbody	oceanFFT	particles	
postProcessGL	recursiveGaussian	simpleD3D10	
simpleD3D10Texture	simpleD3D11Texture	simpleD3D9	
simpleD3D9Texture	simpleGL	simpleTexture3D	

Operating on Systems with non-TCC NVIDIA GPUs

- ▶ NVIDIA GPUs running under TCC mode may coexist with other display devices.
- ▶ The Tesla driver is overinstalled over any NVIDIA display driver in the system—the NVIDIA Tesla driver then becomes the only driver for all NVIDIA GPUs in the system.

If the Tesla driver is uninstalled at a later time, the previous driver is not restored.

- ▶ NVIDIA GPUs that do not support TCC mode will appear as “VGA adapters” in the Windows Device Manager and can be used to drive displays.

Non-supported NVIDIA GPUs can still function as CUDA devices, but the GPU’s graphics functionality is not available to applications.

Setting TCC Mode for Tesla Products

The following table shows the default TCC/non-TCC mode for supported Tesla products.

Table B.1 Default TCC Mode for Tesla Products

Tesla Product	Default Mode	Notes
K20Xm	TCC	Windows 7 64-bit, Windows 8 64-bit Windows Server 2012 64-bit
K20c	TCC	Windows 7 64-bit, Windows 8 64-bit Windows Server 2012 64-bit
K20m	non-TCC	Windows 7 64-bit, Windows 8 64-bit Windows Server 2012 64-bit
C1060	TCC	Windows Server 2012 64-bit
C2075	TCC	Windows 7, Windows Server 2008 & Windows Server 2008 R2 Windows Server 2012 64-bit
	non-TCC	Windows Vista
C2050	TCC	Windows Server 2008 & Windows Server 2008 R2 Windows Server 2012 64-bit
	non-TCC	Windows Vista, Windows 7
C2070	TCC	Windows Server 2008 & Windows Server 2008 R2 Windows Server 2012 64-bit
	non-TCC	Windows Vista, Windows 7
S2050	TCC	
S1070	TCC	
M1060	TCC	
M2050	TCC	
M2070	TCC	
M2075	TCC	

Table B.1 Default TCC Mode for Tesla Products

Tesla Product	Default Mode	Notes
M2070-Q	non-TCC	
M2090	TCC	

To change the TCC mode, use the NVIDIA **smi** utility as follows:

```
nvidia-smi -g (GPU ID) -dm (0 for WDDM, 1 for TCC)
```


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.

Macrovision Compliance Statement

NVIDIA Products that are Macrovision enabled can only be sold or distributed to buyers with a valid and existing authorization from Macrovision to purchase and incorporate the device into buyer's products.

Macrovision copy protection technology is protected by U.S. patent numbers 5,583,936; 6,516,132; 6,836,549; and 7,050,698 and other intellectual property rights. The use of Macrovision's copy protection technology in the device must be authorized by Macrovision and is intended for home and other limited pay-per-view uses only, unless otherwise authorized in writing by Macrovision. Reverse engineering or disassembly is prohibited.

OpenCL Notice

Portions of the NVIDIA system software contain components licensed from third parties under the following terms:

Clang & LLVM:

Copyright (c) 2003-2008 University of Illinois at Urbana-Champaign.

All rights reserved.

Portions of LLVM's System library:

Copyright (C) 2004 eXtensible Systems, Inc.

Developed by:

LLVM Team

University of Illinois at Urbana-Champaign

<http://llvm.org>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal with the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimers.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimers in the documentation and/or other materials provided with the distribution.

Neither the names of the LLVM Team, University of Illinois at Urbana-Champaign, nor the names of its contributors may be used to endorse or promote products derived from this Software without specific prior written permission.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE CONTRIBUTORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS WITH THE SOFTWARE.

Trademarks

NVIDIA, the NVIDIA logo, NVIDIA nForce, GeForce, NVIDIA Quadro, NVDDVD, NVIDIA Personal Cinema, NVIDIA Soundstorm, Vanta, TNT2, TNT, RIVA, RIVA TNT, VOODOO, VOODOO GRAPHICS, WAVEBAY, Accuvision, Antialiasing, Detonator, Digital Vibrance Control, ForceWare, NVRotate, NVSensor, NVSync, PowerMizer, Quincunx Antialiasing, Sceneshare, See What You've Been Missing, StreamThru, SuperStability, T-BUFFER, The Way It's Meant to be Played Logo, TwinBank, TwinView and the Video & Nth Superscript Design Logo are registered trademarks or trademarks of NVIDIA Corporation in the United States and/or other countries. Other company and product names may be trademarks or registered trademarks of the respective owners with which they are associated.

Copyright

© 2009, 2013 NVIDIA Corporation. All rights reserved.