



Release 375 Quadro, NVS, Tesla, & Notebook Drivers - Version 377.83

RN-07954-377.83_v01 | January 18, 2018
Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

Release Notes



TABLE OF CONTENTS

1	Introduction to Release Notes	1
	Structure of the Document	1
	Changes in this Edition	1
2	Release 375 Driver Changes	2
	Version 377.83 Highlights	3
	Existing Support	3
	What's New in Version 377.83	3
	What's New in Release 375	4
	Discontinued and Unsupported Features in this Release	5
	Advanced Instructions for this Release	6
	Turning Off V-Sync to Boost Performance	6
	NVIDIA Application Configuration Engine (ACE)	7
	SLI Multi-OS - GPU Assignment in System Virtualization	7
	Using the WDDM Driver Model with Tesla GPU GOMs	10
	Changes in Version 377.83	12
	Security Updates in this Release	12
	Changes in Version 377.61	13
	Changes in Version 377.55	13
	Changes in Version 377.48	14
	Changes in Version 377.35	15
	Changes in Version 376.33	16
	Changes in Version 375.86	17
	Changes in Version 375.63	18
	Open Issues in Version 377.83	19
	Windows Server 2008 R2 Issues	19
	Windows Server 2012/2012 R2 Issues	19
	Issues Not Caused by NVIDIA Drivers	20
	Windows Server 2008 R2 and Windows Server 2012 Considerations	20
	OpenGL Application Issues For Windows Server 2008/Windows Server 2012	20
	Known Product Limitations	21
	Some APIs do not Report Total Available Graphics Memory Correctly	22
	Using HDMI/DisplayPort Audio with Displays that have a High Native	

TABLE OF CONTENTS

Resolution	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
Applying Workstation Application Profiles	25
1280x1024 @ 60 Hz not Available on BenQ FP241W Monitors	25
OpenGL Viewport Scaling Problem in Horizontal Span Mode.....	25
Image Sharpening Control not Available	25
3 The Release 375 Driver	26
Hardware and Software Support.....	26
Supported Operating Systems.....	26
Supported NVIDIA Products.....	27
Supported Languages.....	32
Driver Installation	33
Minimum Hard Disk Space	33
Before You Begin.....	33
Installation Instructions.....	35
Appendix A: Mode Support for Windows	36
General Mode Support Information.....	37
Default Modes Supported by GPU.....	38
Understanding the Mode Format	38
Tesla, Quadro, and NVS Family of GPUs.....	39
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	43

LIST OF TABLES

Table 1	GOM and Driver Model Compatibility	11
Table 2.1	Security Updates for NVIDIA Software Vulnerabilities	12
Table 3.1	Supported NVIDIA Workstation Products	27
Table 3.2	Supported NVIDIA Quadro Sync II Products	28
Table 3.3	Supported NVIDIA Quadro Sync Products	28
Table 3.6	Supported NVIDIA Quadro Plex Products	29
Table 3.7	Supported NVIDIA Quadro Notebook Products	29
Table 3.4	Supported NVIDIA Quadro G-Sync II Products	29
Table 3.5	Supported NVIDIA Quadro SDI Products	29
Table 3.8	Supported NVIDIA Quadro Blade/Embedded Graphics Board Series	30
Table 3.9	Supported NVIDIA Tesla Products	31
Table A.1	Modes Supported for High Resolution Displays	37
Table A.2	Non-standard Modes Supported	37
Table B.1	Default TCC Mode	43

01 INTRODUCTION TO RELEASE NOTES

This edition of *Release Notes* describes the Release 375 family of Quadro, NVS, Tesla, and GRID Drivers (versions 375.xx to 377.xx) for Microsoft® Windows® Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2. 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 375 Driver Changes” on page 2](#) gives a summary of changes, and fixed and open issues in this version.
- ▶ [“The Release 375 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 36](#) 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 Server includes information about NVIDIA graphics driver version 377.83, and lists changes made to the driver since version 369.07. These changes are discussed beginning with the chapter [“Release 375 Driver Changes” on page 2](#).

02 RELEASE 375 DRIVER CHANGES

This chapter describes open issues for version 377.83, and resolved issues and driver enhancements for versions of the Release 375 driver up to version 377.83. The chapter contains these sections:

- ▶ “Version 377.83 Highlights” on page 3
- ▶ “Advanced Instructions for this Release” on page 6
- ▶ “Changes in Version 377.83” on page 12
- ▶ “Changes in Version 377.61” on page 13
- ▶ “Changes in Version 377.55” on page 13
- ▶ “Changes in Version 377.48” on page 14
- ▶ “Changes in Version 377.35” on page 15
- ▶ “Changes in Version 375.63” on page 18
- ▶ “Open Issues in Version 377.83” on page 19
- ▶ “Issues Not Caused by NVIDIA Drivers” on page 20
- ▶ “Known Product Limitations” on page 21

Version 377.83 Highlights

This section provides highlights of version 377.83 of the NVIDIA Release 375 Driver for Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2

- ▶ Existing Support
- ▶ What's New in Version 377.83
- ▶ What's New in Release 375
- ▶ Discontinued and Unsupported Features in this Release

Existing Support

- ▶ This release supports the following APIs:
 - Open Computing Language (OpenCL™ software) 1.2 for NVIDIA® Kepler™, Maxwell™, and Pascal™ GPUs
 - OpenGL® 4.5
 - Vulkan® 1.0
 - DirectX 11
 - CUDA 8.0
- ▶ This driver installs **nView Desktop Manager** version 148.47.

What's New in Version 377.83

- ▶ This driver adds security updates for driver components. See “Security Updates in this Release” on page 12 for details.

What's New in Release 375

The section summarizes the following driver changes in Release 375 (since Release 367):

Quadro ODE Driver

Release 375 is an 'Optimal Drivers for Enterprise' (ODE) branch release. ODE branches are dedicated to relatively long term stability for ISV certification, OEMs, and Enterprise customers.

New Product Support

Added support for the following products:

- ▶ NVIDIA Quadro GP100
- ▶ NVIDIA Quadro P6000, P5000, P4000, P2000, P1000, P600, and P400
- ▶ NVIDIA Quadro M5000 SE
- ▶ NVIDIA Quadro M3000 SE

New Features Supported

- ▶ Added support for Quadro Sync II.

Quadro Sync II provides Quadro Sync functionality for the Pascal family of Quadro products.

Advanced features for Quadro Sync II will be available in a later ODE branch release.

- ▶ Video SDK Support

Hardware video encoding and decoding via Video SDK APIs are officially supported by FFMPEG and libAV as of Video SDK 7.0.

- ▶ OpenGL

NVIDIA has published the following OpenGL extensions that expose the Pascal GPU architecture's new hardware features. For more information, see [here](#).

- Single Pass Stereo: `GL_NV_stereo_view_rendering`
- Lens Matched Shading: `GL_NV_clip_space_w_scaling`
- Improved Conservative Rasterization:
`GL_NV_conservative_raster_pre_snap_triangles`
- Double Precision Atomic Operations: `GL_NV_shader_atomic_float64`

- ▶ The NVCapture SDK can now capture a Mosaic display.

- ▶ VR SDK

- OpenGL VR SDK 1.4: Implemented `GL_NV_gpu_multicast` extension

- Vulkan VR: Implemented cross-API interoperability between Vulkan and DirectX 11 for VR HMDs
 - VRDirect: Added support for Microsoft DirectX Graphics Infrastructure (DXGI) 2 Preview.
- ▶ Implemented FastSync: New VSync mode that allows DirectX applications to render faster than the display refresh rate, yet still display frames without tearing.

Discontinued and Unsupported Features in this Release

Discontinued Features

- ▶ Support for Quadro SDI products is discontinued for Windows Server 2012 and later operating systems.
- ▶ Discontinued support for NVIDIA GRID products.

Beginning with Release 375, the NVIDIA drivers no longer support the NVIDIA GRID K2, K1, K520, and K340 products.

Advanced Instructions for this Release

This section clarifies instructions for successfully accomplishing the following tasks:

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

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.

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
Pascal	<u>Quadro</u> : P2000, P4000, P5000, P6000, GP100
Maxwell	<u>Quadro</u> : K2200, M2000, M4000, M5000, M6000, M6000 24GB <u>Tesla</u> : M60, M6
Kepler	<u>Quadro</u> : K2000, K4000, K4200, K5000, K200, K6000 <u>Tesla</u> : K10, K20, K20x, K20Xm, K20c, K20s, K40m, K40c, K40s, K40st, K40t, K80
Fermi	<u>Quadro</u> : 2000, 4000, 5000, 6000 <u>Quadro-MXM</u> : 1000M, 3000M <u>Tesla</u> : C2050, C2075, M2050, M2070, M2070Q

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

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.

Changes in Version 377.83

The following sections list the important changes and the most common issues resolved since driver version 377.61.

Security Updates in this Release

For more information about NVIDIA's vulnerability management, visit the [NVIDIA Product Security](#) page.

Table 2.1 Security Updates for NVIDIA Software Vulnerabilities

CVE ID	NVIDIA Issue Number	More Information
CVE-2017-5753	2034846	Security Bulletin: NVIDIA Driver Security Updates for CPU Speculative Side Channel Vulnerabilities

Changes in Version 377.61

The following sections list the important changes and the most common issues resolved since driver version 377.55.

- ▶ This driver added security updates for driver components.

Changes in Version 377.55

The following sections list the important changes and the most common issues resolved since driver version 377.48.

- ▶ This driver added security updates for driver components.

Changes in Version 377.48

The following sections list the important changes and the most common issues resolved since driver version 377.35.

- ▶ Users may experience issues loading the Windows display driver on systems with greater than 1.75TB of system memory.
- ▶ [OpenGL][Cinema 4D]: Random driver crashes occur when running the OpenGL application.

Changes in Version 377.35

The following sections list the important changes and the most common issues resolved since driver version 376.33.

- ▶ This driver adds security updates for driver components.
- ▶ Fixed a bug that can cause the GPU InfoROM to be corrupted.

Changes in Version 376.33

The following sections list the important changes and the most common issues resolved since driver version 375.86.

- ▶ This driver adds security updates for driver components.

Changes in Version 375.86

The following sections list the important changes and the most common issues resolved since driver version 375.63.

- ▶ ITU conformance stream is not being decoded correctly.

Changes in Version 375.63

The following sections list the important changes and the most common issues resolved since driver version 369.07.

- ▶ Added security updates for driver components.
- ▶ [NVIDIA Control Panel][CATIA V5/V6]: Updated application profile settings.
- ▶ [Quandro M6000][Quantum 3D Mantis]: Fixed intermittent GPU stalling when running the application.
- ▶ [Quandro M5000][Quantum 3D Mantis][Quadro Sync]: Fixed swapgroup and swapbarrier failures when running the application with Quadro Sync enabled.
- ▶ [IC.IDO][VR]: Fixed “Error code 8” crash that occurred when running the application on VR hardware.

Open Issues in Version 377.83

As with every released driver, version 377.83 of the Release 375 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 Server 2008 R2 Issues

- ▶ [Quadro M6000][Quadro Sync]: Unexpected flashing may appear on the display connected to the second GPU.
Please consult your NVIDIA support team for details on resolving this behavior.
- ▶ [Quadro Sync][Quadro K5000]: Performance drop occurs when using Quadro Sync.
- ▶ [Quadro Sync]: Frame overruns occur on low resolution channels when using Quadro Sync swap synchronization.
- ▶ [Quadro 4000/K5000]: Rendering fails with `GL_INVALID_OPERATION` error in the draw call.
- ▶ [Quadro][Eyeon Fusion]: With Aero/desktop compositing and Quad-buffer stereo enabled, the application child windows flicker.

Windows Server 2012/2012 R2 Issues

- ▶ [Quadro M6000][Quadro Sync]: Unexpected flashing may appear on the display connected to the second GPU.
Please consult your NVIDIA support team for details on resolving this behavior.
- ▶ [Quadro Sync][Quadro K5000]: Performance drop occurs when using Quadro Sync.
- ▶ [Quadro 4000/K5000]: Rendering fails with `GL_INVALID_OPERATION` error in the draw call.

Issues Not Caused by NVIDIA Drivers

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 Server 2008 R2 and Windows Server 2012 Considerations” on page 20
- ▶ “OpenGL Application Issues For Windows Server 2008/Windows Server 2012” on page 20

Windows Server 2008 R2 and Windows Server 2012 Considerations

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>.

OpenGL Application Issues For Windows Server 2008/Windows Server 2012

The following are known compatibility issues for OpenGL applications developed under Windows Server 2003:

- ▶ 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.

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:

- ▶ “Some APIs do not Report Total Available Graphics Memory Correctly” on page 22
- ▶ “Using HDMI/DisplayPort Audio with Displays that have a High Native Resolution” 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
- ▶ “Applying Workstation Application Profiles” on page 25
- ▶ “1280x1024 @ 60 Hz not Available on BenQ FP241W Monitors” on page 25
- ▶ “OpenGL Viewport Scaling Problem in Horizontal Span Mode” on page 25
- ▶ “Image Sharpening Control not Available” on page 25

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

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).

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.

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 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.

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.

OpenGL Viewport Scaling Problem in Horizontal Span Mode

With nView Horizontal Span mode enabled, when opening an OpenGL model in a viewport, the model image is scaled too large to fit in the viewport. The problem occurs with such applications as Maya 5.0 and 3D Studio MAX 4.26.

This is not an NVIDIA bug, but a limitation in the application's ability to properly maintain the aspect ratio in Horizontal Span mode.

Image Sharpening Control not Available

The **Image sharpening** slider on the NVIDIA Control Panel-> Display->Adjust Desktop Color Settings page is grayed out.

This control is intentionally disabled because image sharpening is not supported on current GPUs.

03 THE RELEASE 375 DRIVER

This chapter covers the following main topics:

- ▶ “Hardware and Software Support” on page 26
- ▶ “Driver Installation” on page 33

Hardware and Software Support

Supported Operating Systems

The Release 375 driver, version 377.83, has been tested with the following Microsoft operating systems:

- ▶ Windows Server 2008R2 (64-bit)
- ▶ Windows Server 2012
- ▶ Windows Server 2012 R2 (64-bit)

Supported NVIDIA Products

The following tables list the NVIDIA products supported by the Release 375 driver, version 377.83.

- ▶ “Supported NVIDIA Workstation Products” on page 27
- ▶ “Supported NVIDIA Quadro Sync II Products” on page 28
- ▶ “Supported NVIDIA Quadro Sync Products” on page 28
- ▶ “Supported NVIDIA Quadro G-Sync II Products” on page 29
- ▶ “Supported NVIDIA Quadro SDI Products” on page 29
- ▶ “Supported NVIDIA Quadro Plex Products” on page 29
- ▶ “Supported NVIDIA Quadro Blade/Embedded Graphics Board Series” on page 30
- ▶ “Supported NVIDIA Tesla Products” on page 31

Table 3.1 Supported NVIDIA Workstation Products

Product	GPU Architecture
NVIDIA Quadro GP100	Pascal
NVIDIA Quadro P6000	Pascal
NVIDIA Quadro P5000	Pascal
NVIDIA Quadro P4000	Pascal
NVIDIA Quadro P2000	Pascal
NVIDIA Quadro P1000	Pascal
NVIDIA Quadro P600	Pascal
NVIDIA Quadro P400	Pascal
NVIDIA Quadro M6000 24 GB	Maxwell
NVIDIA Quadro M6000 24 GB	Maxwell
NVIDIA Quadro M6000	Maxwell
NVIDIA Quadro M5000	Maxwell
NVIDIA Quadro M4000	Maxwell
NVIDIA Quadro M2000	Maxwell
NVIDIA Quadro K6000	Kepler
NVIDIA Quadro K5200	Kepler
NVIDIA Quadro K5000	Kepler
NVIDIA Quadro K4200	Kepler
NVIDIA Quadro K4000	Kepler
NVIDIA Quadro K2200	Maxwell
NVIDIA Quadro K2000	Kepler

Table 3.1 Supported NVIDIA Workstation Products (continued)

Product	GPU Architecture
NVIDIA Quadro K1200	Maxwell
NVIDIA Quadro K620	Maxwell
NVIDIA Quadro K600	Kepler
NVIDIA Quadro K420	Kepler
NVIDIA Quadro 6000	Fermi
NVIDIA Quadro 5000	Fermi
NVIDIA Quadro 4000	Fermi
NVIDIA Quadro 2000D	Fermi
NVIDIA Quadro 2000	Fermi
NVIDIA Quadro 600	Fermi
NVIDIA Quadro 410	Fermi
NVIDIA NVS 810	Maxwell
NVIDIA NVS 510	Kepler
NVIDIA NVS 315	Fermi
NVIDIA NVS 310	Fermi

Table 3.2 Supported NVIDIA Quadro Sync II Products

Product	GPU Architecture
NVIDIA Quadro GP100	Pascal
NVIDIA Quadro P6000	Pascal
NVIDIA Quadro P5000	Pascal
NVIDIA Quadro P4000	Pascal

Table 3.3 Supported NVIDIA Quadro Sync Products

Product	GPU Architecture
NVIDIA Quadro M6000 24 GB	Maxwell
NVIDIA Quadro M6000	Maxwell
NVIDIA Quadro M5000	Maxwell
NVIDIA Quadro M4000	Maxwell
NVIDIA Quadro K6000	Kepler
NVIDIA Quadro K5200	Kepler
NVIDIA Quadro K5000	Kepler
NVIDIA Quadro K4200	Kepler

Table 3.4 Supported NVIDIA Quadro G-Sync II Products

Product	GPU Architecture
NVIDIA Quadro 6000	Fermi
NVIDIA Quadro 5000	Fermi

Table 3.5 Supported NVIDIA Quadro SDI Products

Product	GPU Architecture
NVIDIA Quadro K6000	Kepler
NVIDIA Quadro K5200	Kepler
NVIDIA Quadro K5000	Kepler
NVIDIA Quadro K4200	Kepler
NVIDIA Quadro K4000	Kepler
NVIDIA Quadro 6000	Fermi
NVIDIA Quadro 5000	Fermi
NVIDIA Quadro 4000	Fermi



Note: Quadro SDI products are not supported under Windows Server 2012.

Table 3.6 Supported NVIDIA Quadro Plex Products

Product	GPU Architecture
NVIDIA Quadro Plex 7000	Fermi

Table 3.7 Supported NVIDIA Quadro Notebook Products

Product	GPU Architecture
NVIDIA Quadro M5500	Maxwell
NVIDIA Quadro M5000M	Maxwell
NVIDIA Quadro M4000M	Maxwell
NVIDIA Quadro M3000M	Maxwell
NVIDIA Quadro M2000M	Maxwell
NVIDIA Quadro M1000M	Maxwell

Table 3.7 Supported NVIDIA Quadro Notebook Products (continued)

Product	GPU Architecture
NVIDIA Quadro M600M	Maxwell
NVIDIA Quadro M500M	Maxwell
NVIDIA Quadro K5100M	Kepler
NVIDIA Quadro K5000M	Kepler
NVIDIA Quadro K4100M	Kepler
NVIDIA Quadro K4000M	Kepler
NVIDIA Quadro K3100M	Kepler
NVIDIA Quadro K3000M	Kepler
NVIDIA Quadro K2200M	Maxwell
NVIDIA Quadro K2100M	Kepler
NVIDIA Quadro K2000M	Kepler
NVIDIA Quadro K1100M	Kepler
NVIDIA Quadro K1000M	Kepler
NVIDIA Quadro K620M	Maxwell
NVIDIA Quadro K610M	Kepler
NVIDIA Quadro K510M	Kepler
NVIDIA Quadro K500M	Kepler
NVIDIA Quadro 5010M	Fermi
NVIDIA Quadro 4000M	Fermi
NVIDIA Quadro 3000M	Fermi
NVIDIA Quadro 2000M	Fermi
NVIDIA Quadro 1000M	Fermi
NVIDIA Quadro 500M	Fermi

Table 3.8 Supported NVIDIA Quadro Blade/Embedded Graphics Board Series

Product	GPU Architecture
NVIDIA Quadro M5000 SE	Maxwell
NVIDIA Quadro M3000 SE	Maxwell
NVIDIA Quadro K3100M	Kepler
NVIDIA Quadro 4000M	Fermi
NVIDIA Quadro 3000M	Fermi
NVIDIA Quadro 1000M	Fermi

Table 3.8 Supported NVIDIA Quadro Blade/Embedded Graphics Board Series (continued)

Product	GPU Architecture
NVIDIA Quadro 500M	Fermi

Table 3.9 Supported NVIDIA Tesla Products

Product	GPU Architecture	Notes
Tesla P-Series Products		
NVIDIA Tesla P100	Pascal	
NVIDIA Tesla P40	Pascal	
NVIDIA Tesla P4	Pascal	
Tesla M-Series Products		
NVIDIA Tesla M60	Maxwell	
NVIDIA Tesla M40	Maxwell	
NVIDIA Tesla M6	Maxwell	
NVIDIA Tesla M4	Maxwell	
Tesla K-Series Products		
NVIDIA Tesla K80	Kepler	64-bit OS only
NVIDIA Tesla K40(m/c/s/st/t)	Kepler	64-bit OS only
NVIDIA Tesla K20(x/c/m/Xm/s)	Kepler	64-bit OS only
NVIDIA Tesla K10	Kepler	
Tesla S-Class Products		
NVIDIA Tesla S2050	Fermi	
Tesla M-Class Products		
NVIDIA Tesla M2090	Fermi	
NVIDIA Tesla M2075	Fermi	
NVIDIA Tesla M2070	Fermi	
NVIDIA Tesla M2070Q	Fermi	
NVIDIA Tesla M2050	Fermi	
Tesla C-Class Products		
NVIDIA Tesla C2075	Fermi	
NVIDIA Tesla C2070	Fermi	
NVIDIA Tesla C2050	Fermi	

Supported Languages

The Release 375 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 450 MB for English-only, and 900 MB for International.

The hard disk space requirement for 64-bit is minimum 570 MB for English-only, and 1158 MB for International.

Before You Begin

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.

Preserving Settings Before Upgrading Your Software

Before uninstalling or installing software, you can preserve your nView Desktop Manager and/or NVIDIA Display settings by using the nView Desktop Manager Profiles features.



Note: Follow the steps below and/or refer to the NVIDIA nView Desktop Manager User's Guide for details. Under Windows XP/2000 you must have, at least, Power User access privileges in order to create or save a profile. (Refer to Windows Help if you need an explanation of Power User access rights.)

Follow the steps below and/or refer to the *NVIDIA nView Desktop Manager User's Guide* for details.

- 1 Open the nView Desktop Manager Profiles page (Figure 4.1).
- 2 To preserve your current settings, you can use either the **Save** or the **New** option from the nView Desktop Manager Profiles page:
 - If you want to overwrite the currently loaded profile with your changed settings, use the **Save** option. Notice that a warning message indicates that you are about to overwrite the selected profile.
 - If you want to retain the currently loaded profile and want to save your changed settings to a new file, click the **New** option. Enter a name and description of the profile in the New Profile dialog box. For example, you can name this profile **My Settings**.
- 3 If you are an "advanced" user and want to customize certain settings in the saved profile, click **Advanced** << to expand the dialog box (Figure 4.2).

- 4 To customize the settings, you can select or clear any of the settings check boxes.
- 5 Click **Save** to return to the main Profiles page.

If you created a new profile, you will see the name of the newly created profile in the profiles list.

If you overwrote a current profile, the same profile name is retained in the list.



Note: nView Desktop Manager profile (.tvp) files are saved in the Windows\nView directory. Depending on the version of the NVIDIA driver previously installed, profiles may also be saved in the Documents and Settings\All Users\Application Data\nView_Profiles directory.

- 6 Now you can uninstall your current driver for a driver upgrade.
- 7 After you restart your computer following an NVIDIA new driver install, you can easily load the saved profile from the Profiles page of nView Desktop Manager.

About Using Saved Profiles in Another Computer

You can easily use any saved profile (.tvp file in the **Windows\nView** directory) from one computer and use it in another computer, if you want. You'll need to copy it to the **Windows\nView** directory of a computer that has the NVIDIA ForceWare graphics display driver, etc. installed properly. Then this profile can be loaded from another computer from the nView Desktop Manager Profiles page just as it can from your original computer.

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.
- 3 The license agreement dialog box appears.
- 4 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.
- 5 Extract the zip files to a temporary folder on your PC.
- 6 Open the NVIDIA driver installation .EXE file to launch the NVIDIA InstallShield Wizard.
- 7 Follow the instructions in the NVIDIA InstallShield Wizard to complete the installation.

APPENDIX A MODE SUPPORT FOR WINDOWS

This chapter details the Windows modes supported by the Release 375 driver for NVIDIA products. It contains these sections:

- ▶ “General Mode Support Information” on page 37
- ▶ “Default Modes Supported by GPU” on page 38

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 38.

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, and NVS Family of GPUs” on page 39

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, and NVS Family of GPUs

This section 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

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” 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 or supported Quadro 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/Quadro 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/Quadro 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/Quadro 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/Quadro 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/Quadro 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

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

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

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

Table B.1 Default TCC Mode

Product	Default Mode
K20Xm/c	TCC
K20m	TCC Off
C2075	TCC
C2050	TCC Off
C2070	TCC Off
S2050	TCC
M2050	TCC
M2070	TCC
M2075	TCC
M2070-Q	TCC Off
M2090	TCC
Quadro (Kepler/Maxwell)	TCC Off

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, VODOO, VODOO 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

© 2016-2018 NVIDIA Corporation. All rights reserved.