



# NVIDIA RTX Enterprise Release 465 Drivers, Version 465.89

For NVIDIA RTX, Quadro, Quadro Notebook, NVS, and Data Center GPUs  
Windows 10

Release Notes

# Table of Contents

<b>1</b>	<b>Introduction to Release Notes</b> .....	<b>1</b>
	Structure of the Document.....	1
	Changes in this Edition.....	1
<b>2</b>	<b>Release 465 Driver Changes</b> .....	<b>2</b>
	Version 465.89 Highlights.....	3
	Existing Support.....	3
	What's New in Version 465.89.....	3
	What's New in Release 465.....	5
	Discontinued and Unsupported Features.....	6
	Limitations in This Release.....	6
	Open Issues in Version 465.89.....	8
	Windows 10 Issues.....	8
	Issues Not Caused by NVIDIA Drivers.....	9
	Application Issues.....	9
<b>3</b>	<b>The Release 465 Driver</b> .....	<b>10</b>
	Driver Security.....	11
	Restricting/Enabling Access to GPU Performance Counters.....	11
	Advanced Driver Information.....	13
	Turning Off V-Sync to Boost Performance.....	13
	NVIDIA Application Configuration Engine (ACE).....	13
	Using the WDDM Driver Model with Tesla GPU GOMs.....	14
	SLI Multi-OS – GPU Assignment in System Virtualization.....	16
	Known Product Limitations.....	19
	Issues Installing the NVIDIA Control Panel from the Windows Store.....	20
	Some APIs do not Report Total Available Graphics Memory Correctly.....	20
	Using HDMI/DisplayPort Audio with Displays that have a High Native Resolution.....	22
	Using HDMI/DisplayPort Audio in Dualview or Clone Mode Configurations.....	22
22	GPU Runs at a High Performance Level (full clock speeds) in Multi-display Modes.....	22
	SLI Connector Requirement on NVIDIA Quadro SLI Cards.....	23
	Applying Workstation Application Profiles.....	23
	Hardware and Software Support.....	24
	Supported Operating Systems.....	24
	Supported NVIDIA Workstation GPUs.....	24
	Supported NVIDIA Notebook GPUs.....	28
	Supported Languages.....	29

# Table of Contents

Driver Installation.....	30
Minimum Hard Disk Space .....	30
Before You Begin .....	30
Installation Instructions.....	30
<b>NVIDIA Tesla Compute Cluster Mode 32</b>	
About Tesla Compute Cluster Mode.....	32
TCC Overview.....	32
Running CUDA Applications .....	33
Operating on Systems with non-TCC NVIDIA GPUs .....	34
Setting TCC Mode.....	34

# List of Tables

Table 1	GOM and Driver Model Compatibility .....	15
Table 5.1	Supported NVIDIA Quadro & NVS Products.....	25
Table 5.2	Supported NVIDIA Quadro Sync II Products .....	26
Table 5.3	Supported NVIDIA Quadro Sync Products .....	26
Table 5.4	Supported NVIDIA Quadro Blade/Embedded Graphics Board Series .....	27
Table 5.5	Supported NVIDIA Data Center Products.....	27
Table 5.6	NVIDIA Quadro Notebook GPU Support .....	28
Table A.1	Default TCC Mode .....	34

---

# Chapter 1. Introduction to Release Notes

This edition of *Release Notes* describes the Release 465 family of NVIDIA<sup>®</sup> RTX, Quadro<sup>®</sup>, NVS<sup>™</sup>, Data Center, and Quadro Notebook Drivers for Microsoft<sup>®</sup> Windows<sup>®</sup> 10. 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 465 Driver Changes](#)” on [page 2](#) gives a summary of changes, and fixed and open issues in this version.
- ▶ “[The Release 465 Driver](#)” on [page 10](#) describes the NVIDIA products and languages supported by this driver, the system requirements, and how to install the driver.
- ▶ “[NVIDIA Tesla Compute Cluster Mode](#)” on [page 32](#) describes the Tesla Compute Cluster mode.

## Changes in this Edition

This edition of the *Release Notes* for Windows 10 includes information about NVIDIA graphics driver version 465.89, and lists changes made to the driver since version 461.92.

These changes are discussed beginning with the chapter “[Release 465 Driver Changes](#)” on [page 2](#).

---

## Chapter 2. Release 465 Driver Changes

This chapter describes open issues for version 465.89, and resolved issues and driver enhancements for versions of the Release 465 driver up to version 465.89.

The chapter contains these sections:

- ▶ “Version 465.89 Highlights” on page 3
- ▶ “Open Issues in Version 465.89” on page 8
- ▶ “Issues Not Caused by NVIDIA Drivers” on page 9

## Version 465.89 Highlights

This section provides highlights of version 465.89 of the NVIDIA Release 465 Driver for Windows 10.

- ▶ Existing Support
- ▶ What's New in Version 465.89
- ▶ What's New in Release 465
- ▶ Discontinued and Unsupported Features
- ▶ Limitations in This Release

### Existing Support

- ▶ This release supports the following APIs:
  - Open Computing Language (OpenCL™ software) 3.0 for NVIDIA® Maxwell™ and later GPUs
  - OpenGL® 4.6
  - Vulkan® 1.2
  - DirectX 11
  - DirectX 12
  - NVIDIA® CUDA® 11.2
- ▶ This driver installs nView Desktop Manager version 201.66.

## What's New in Version 465.89

### New Features

- ▶ Displays configured through Mosaic can now be independently mirrored; useful for matching bezel sizes of displays (requires NVIDIA RTX GPU)
- ▶ Stereo scanouts now offered in normal, vertical/horizontal mirror, and 180 degree rotation modes; useful for compatibility with beam splitter stereo deployments (requires NVIDIA RTX GPU)

### OpenCL 3.0

Added support for OpenCL 3.0<sup>1</sup>, the latest major version of OpenCL maintaining backward compatibility with OpenCL 1.2. NVIDIA OpenCL 3.0 continues to support existing OpenCL 1.2

---

1. Khronos has recently released OpenCL 3.0 spec ([https://www.khronos.org/registry/OpenCL/specs/3.0-unified/pdf/OpenCL\\_API.pdf](https://www.khronos.org/registry/OpenCL/specs/3.0-unified/pdf/OpenCL_API.pdf))

functionality as well as Khronos and vendor extensions that are already supported with NVIDIA OpenCL 1.2 drivers.

- ▶ The following new features beyond existing NVIDIA OpenCL 1.2 features are supported by NVIDIA OpenCL 3.0
  - RGBA vector component naming in OpenCL C kernels
  - `pragma_unroll` hint
  - `opencl_3d_image_writes`
  - `clCreate*WithProperties` APIs which can be used as replacement for existing `clCreateBuffer/Image` APIs.
  - `clSetContextDestructorCallback`
  - `clCloneKernel` from OpenCL 2.1
  - `clEnqueueSVMMigrateMem` from OpenCL 2.1
- ▶ Incorporates the following experimental 2.0 features:
  - Device side enqueue
    - > The current implementation is limited to 64-bit platforms only.
    - > Allows kernels to be enqueued with `global_work_size` larger than the compute capability of the NVIDIA GPU. The current implementation supports only combinations of `global_work_size` and `local_work_size` that are within the compute capability of the NVIDIA GPU.
 

The maximum supported CUDA grid and block size of NVIDIA GPUs is available at <http://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html#compute-capabilities>. For a given grid dimension, the `global_work_size` can be determined by CUDA grid size x CUDA block size.
    - > For executing kernels (whether from the host or the device), OpenCL 3.0 supports non-uniform ND-ranges where `global_work_size` does not need to be divisible by the `local_work_size`. This capability is not yet supported in the NVIDIA driver, and therefore not supported for device side kernel enqueues.



**Note:** Other OpenCL 2.X entry-points which are now optional and are not supported in NVIDIA OpenCL 3.0 will behave as described at [https://www.khronos.org/registry/OpenCL/specs/3.0-unified/html/OpenCL\\_API.html#opencl-3.0-backwards-compatibility](https://www.khronos.org/registry/OpenCL/specs/3.0-unified/html/OpenCL_API.html#opencl-3.0-backwards-compatibility)



## What's New in Release 465

The section summarizes the following driver changes in Release 465 (since Release 460).

### Driver Name Change

With the introduction of NVIDIA RTX products powering the next generation of professional visual, compute, and AI platform solutions for the enterprise, the Quadro driver will be transitioning to the NVIDIA RTX Enterprise brand to better align with new and future products. NVIDIA RTX Enterprise drivers will continue to provide users with the same level of enterprise-class quality, reliability, performance, and security as previous Quadro drivers. These new drivers will not only deliver the latest features and improvements for NVIDIA RTX GPUs, but also for current and previous generation Quadro GPUs.

### NVIDIA RTX New Feature Branch Driver

Release 465 is a New Feature Branch (NFB) driver. NFB drivers provide users the flexibility to try new features made available outside the launch cycle of longer-lived Production Branch drivers. NFB drivers are tested for workstation environments and are recommended only for those users who need specific features mentioned in the release highlights or release notes.

For the most stable and fully supported enterprise driver please use the Production Branch graphics drivers downloadable from the main NVIDIA driver download page.

### New Features

For this driver, the new features are the same as those reported in [“What's New in Version 465.89”](#) on page 3.

## Discontinued and Unsupported Features

### Discontinued Support in this Release

#### ► Implicit SLI

After Release 450, the Quadro driver will no longer support Implicit SLI. R450 will be the last driver branch to support Implicit SLI—future driver branches will support Explicit SLI only.

Implicit SLI means applications do not manage the SLI state, but it is the driver that makes assumptions with application profiles to achieve GPU scaling. Explicit SLI means the application knows the SLI state and uses extensions (such as DirectX 12 linked nodes, Vulkan device groups, or OpenGL multicast extensions) to issue commands to each device in the SLI group.

### Limitations in This Release

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

#### External Graphics

##### ► External GPU Surprise Removal

Not all applications have been designed to address surprise removal of the external GPU; disconnection of the external GPU while applications are running is not advised.

##### ► Mixed GeForce/Quadro Products

Mixed GeForce/Quadro products are supported (Geforce GPU + Quadro eGPU, or Quadro GPU + Geforce eGPU), but requires installation of the GeForce driver package. The Quadro package does not install GeForce drivers.

### Deep Color (10 bpc) on HDR-Capable Monitors Under Windows 10 RS2 and RS3

Depending on the monitors you have connected, Windows may initially set the color depth to 8-bpc (bits per channel). You can then enable deep color (10bpc) through the NVIDIA Control Panel. If you enable deep color on HDR-capable monitors, the screen may become blank due to the current SDR/HDR (standard dynamic range/high dynamic range) management policies in Windows 10 RS2 and RS3. In this event, the driver will revert back to the 8-bpc color space.

### OpenCL 3.0 Known Issues

#### ► Device-Side-Enqueue related queries may return 0 values, although corresponding built-ins can be safely used by kernel.

This is in accordance with conformance requirements described at [https://www.khronos.org/registry/OpenCL/specs/3.0-unified/html/OpenCL\\_API.html#opencl-3.0-backwards-compatibility](https://www.khronos.org/registry/OpenCL/specs/3.0-unified/html/OpenCL_API.html#opencl-3.0-backwards-compatibility)

- ▶ The denormalized results for some math functions that were flushed to zero when `cl-fast-relaxed-math` is passed as a compiler option on OpenCL 1.2 drivers won't be flushed to zero with OpenCL 3.0 drivers.
- ▶ Shared virtual memory - the current implementation of shared virtual memory is limited to 64-bit platforms only.

## Cross-adapter Clone Mode

Under Windows 10, clone mode across graphics processors is handled by the Windows OS, and not by the NVIDIA driver. Consequently, there are no controls for cross-adapter clone mode in the NVIDIA Control Panel under Windows 10.

Functionality of some display controls in the NVIDIA Control Panel may be affected if cross-adapter clone mode is enabled through the Windows Control Panel. In this case, use the Windows Control Panel to adjust display settings.

## Optimus Systems

Because cross-adapter clone mode is handled by the Windows 10 OS, there are no display controls in the NVIDIA Control Panel if no display is connected to the NVIDIA GPU.

## Open Issues in Version 465.89

As with every released driver, version 465.89 of the Release 465 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 10 Issues

- ▶ [Alchemist]: The application crashes when either changing the display resolution or switching to another display with a different resolution.

*To workaround, avoid changing the resolution, or disconnecting from or connecting to an external display with the application opened.*

# 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 7.

- ▶ “Application Issues” on page 9

## Application Issues

- ▶ [3ds Max 2017]: DirectX 9, Mode Texture of some models disappears after some time if no operation (such as zoom, pan, or rotate) is performed on the model.
- ▶ [Quadro K2200][Bunkspeed Shot 2014]: Render Argument exception occurs in Accurate Mode, and Internal Rendering Error occurs with Fast Mode.

*Quadro K2200/K60 are not supported by Bunkspeed software.*

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

*To avoid this issue, use CUDA.*

---

## Chapter 3. The Release 465 Driver

This chapter covers the following main topics:

- ▶ “Driver Security” on page 11
- ▶ “Advanced Driver Information” on page 13
- ▶ “Known Product Limitations” on page 19
- ▶ “Hardware and Software Support” on page 24

# Driver Security

Follow these safe computing practices:

- ▶ Only download or execute content and programs from trusted third parties.
- ▶ Run your system and programs with the least privilege necessary. Users should run without administrator rights whenever possible.
- ▶ When running as administrator, do not elevate UAC privileges for activities or programs that don't need them.

This section describes additional actions to take to mitigate specific known security issues.

## Restricting/Enabling Access to GPU Performance Counters

The NVIDIA graphics driver contains a vulnerability (CVE-2018-6260) that may allow access to application data processed on the GPU through a side channel exposed by the GPU performance counters. GPU performance counters are needed by developers in order to use NVIDIA developer tools such as CUPTI, Nsight Graphics, and Nsight Compute. In order to address CVE-2018-6260 the driver (starting with version 419.67) automatically disables access for non-admin users. For more information about CVE-2018-6260 visit the [NVIDIA Security Bulletin 4772](#).

## Enabling Access to GPU Performance Counters Using the NVIDIA Control Panel

Access to GPU performance counters can be enabled for non-admin users who need to use NVIDIA developer tools. Enabling access to GPU performance counters can be accomplished through the NVIDIA Control Panel->*Developer*->*Manage GPU Performance Counters* page. Refer to the *Developer*->*Manage GPU Performance Counters* section of the NVIDIA Control Panel Help for instructions.

**Note:** Access to GPU performance counters should be kept disabled for non-admin users who do not need to use NVIDIA developer tools.

## Restricting/Enabling Access to GPU Performance Counters Across an Enterprise Using Scripts

Enterprise administrators can use scripts to programmatically apply the settings. The scripts should incorporate the registry key information provided below to automate the deployment.

**CAUTION:** These instructions should be performed only by enterprise administrators. Changes to the registry must be made with care. System instability can result if performed incorrectly.

```
[HKLM\SYSTEM\CurrentControlSet\Services\nvlddmkm\Global\NVTweak]
Value: "RmProfilingAdminOnly"
Type: DWORD
Data: 00000001
```

The data value 1 restricts access to admin users, whereas data value 0 allows access to all users.

A system reboot is required for the changes to take effect.



# Advanced Driver Information

This section clarifies instructions for successfully accomplishing the following tasks:

- ▶ 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

## 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
NVIDIA Ampere	NVIDIA Data Center: A100, A40 NVIDIA RTX: A6000
Turing	Quadro: RTX 8000, RTX 6000, RTX 5000, RTX 4000 NVIDIA Data Center: T4

GPU Family	Boards supported
Volta	<u>Quadro</u> : GV100 NVIDIA Data Center: V100
Pascal	<u>Quadro</u> : P2000, P4000, P5000, P6000, GP100 <u>Tesla</u> : P100, P40, P4
Maxwell	<u>Quadro</u> : K2200, M2000, M4000, M5000, M6000, M6000 24GB <u>Tesla</u> : M60, M6
Kepler	<u>Quadro</u> : K2000, K4000, K4200, K5000, K5200, K6000 <u>Tesla</u> : K10, K20, K20x, K20Xm, K20c, K20s, K40m, K40c, K40s, K40st, K40t, K80

## 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 for ESXi 5.1 and 5.5.

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.

Disabling MSI translation is not needed for ESXi 6.0 and later.

### 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>):

# 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:

- ▶ “Issues Installing the NVIDIA Control Panel from the Windows Store” on page 20
- ▶ “Some APIs do not Report Total Available Graphics Memory Correctly” on page 20
- ▶ “Using HDMI/DisplayPort Audio with Displays that have a High Native Resolution” on page 22
- ▶ “Using HDMI/DisplayPort Audio in Dualview or Clone Mode Configurations” on page 22
- ▶ “GPU Runs at a High Performance Level (full clock speeds) in Multi-display Modes” on page 22
- ▶ “SLI Connector Requirement on NVIDIA Quadro SLI Cards” on page 23
- ▶ “Applying Workstation Application Profiles” on page 23

## Issues Installing the NVIDIA Control Panel from the Windows Store

You may encounter issues when attempting to install the NVIDIA Control Panel from the Windows Store under Windows 10, such as

- ▶ The download process from the Windows Store freezes at the “Starting download ...” stage.
- ▶ The NVIDIA Control Panel fails to download after initiating the download from the notification popup that appears upon installing the driver.

For assistance with installing the NVIDIA Control Panel from the Microsoft Windows Store, see the NVIDIA Knowledge Base Article, [NVIDIA Control Panel Windows Store App](#).

For information about the DCH vs Standard drivers for Windows 10, see the NVIDIA Knowledge Base Article, [NVIDIA DCH/Standard Display Drivers for Windows 10 FAQ](#).

## 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 GB<sup>1</sup>. 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 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.

## 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.

# Hardware and Software Support

This chapter covers the following main topics:

- ▶ “Supported Operating Systems” on page 24
- ▶ “Supported NVIDIA Workstation GPUs” on page 24
- ▶ “Supported NVIDIA Notebook GPUs” on page 28
- ▶ “Supported Languages” on page 29

## Supported Operating Systems

The Release 465 driver, version 465.89, has been tested with

- ▶ Microsoft Windows® 10, 64-bit (versions 1709 and later)

## Supported NVIDIA Workstation GPUs

The following tables list the NVIDIA products supported by the Release 465 driver, version 465.89.

- ▶ Table 5.1, “Supported NVIDIA Quadro & NVS Products” on page 25
- ▶ Table 5.2, “Supported NVIDIA Quadro Sync II Products” on page 26
- ▶ Table 5.3, “Supported NVIDIA Quadro Sync Products” on page 26
- ▶ Table 5.4, “Supported NVIDIA Quadro Blade/Embedded Graphics Board Series” on page 27
- ▶ Table 5.5, “Supported NVIDIA Data Center Products” on page 27

## NVIDIA Quadro, NVS, & RTX Product

Table 5.1 Supported NVIDIA Quadro & NVS Products

Product	GPU Architecture
NVIDIA RTX A6000	NVIDIA Ampere
NVIDIA Quadro RTX 8000	Turing
NVIDIA Quadro RTX 6000	Turing
NVIDIA Quadro RTX 5000	Turing
NVIDIA Quadro RTX 4000	Turing
NVIDIA Quadro RTX 3000	Turing
NVIDIA Quadro GV100	Volta
NVIDIA Quadro GP100	Pascal
NVIDIA Quadro P6000	Pascal
NVIDIA Quadro P5000	Pascal
NVIDIA Quadro P4000	Pascal
NVIDIA Quadro P2200	Pascal
NVIDIA Quadro P2000	Pascal
NVIDIA Quadro P1000	Pascal
NVIDIA Quadro P600	Pascal
NVIDIA Quadro P400	Pascal
NVIDIA Quadro M6000 24GB	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
NVIDIA Quadro K1200	Maxwell
NVIDIA Quadro K620	Maxwell
NVIDIA Quadro K600	Kepler
NVIDIA Quadro K420	Kepler
NVIDIA Quadro 410	Kepler

Table 5.1 Supported NVIDIA Quadro & NVS Products (continued)

Product	GPU Architecture
NVIDIA NVS 810	Maxwell
NVIDIA NVS 510	Kepler

## NVIDIA Quadro Sync II

Table 5.2 Supported NVIDIA Quadro Sync II Products

Product	GPU Architecture
NVIDIA RTX A6000	NVIDIA Ampere
NVIDIA Quadro RTX 8000	Turing
NVIDIA Quadro RTX 6000	Turing
NVIDIA Quadro RTX 5000	Turing
NVIDIA Quadro RTX 4000	Turing
NVIDIA Quadro RTX 3000	Turing
NVIDIA Quadro GV100	Volta
NVIDIA Quadro GP100	Pascal
NVIDIA Quadro P6000	Pascal
NVIDIA Quadro P5000	Pascal
NVIDIA Quadro P4000	Pascal

## NVIDIA Quadro Sync Products

Table 5.3 Supported NVIDIA Quadro Sync Products

Product	GPU Architecture
NVIDIA Quadro M6000 24GB	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

## NVIDIA Quadro Blade/Embedded Graphics Board

Table 5.4 Supported NVIDIA Quadro Blade/Embedded Graphics Board Series

Product	GPU Architecture
NVIDIA Quadro RTX 5000	Turing
NVIDIA Quadro RTX 3000	Turing
NVIDIA Quadro T1000	Turing
NVIDIA Quadro P5000	Pascal
NVIDIA Quadro P3000	Pascal
NVIDIA Quadro P2000	Pascal
NVIDIA Quadro P1000	Pascal
NVIDIA Quadro M5000 SE	Maxwell
NVIDIA Quadro M3000 SE	Maxwell

## NVIDIA Data Center Products

The driver package is designed for systems that have one or more NVIDIA Data Center products installed.

- ▶ Only one GHIC can be connected to the server in a 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 5.5 Supported NVIDIA Data Center Products

Product	GPU Architecture
<b>NVIDIA A-Series Products</b>	
NVIDIA A100	NVIDIA Ampere
NVIDIA A40	NVIDIA Ampere
<b>NVIDIA T-Series Products</b>	
NVIDIA T4	Turing
<b>NVIDIA V-Series Products</b>	
NVIDIA V100	Volta
<b>Tesla P-Series Products</b>	
NVIDIA Tesla P100	Pascal
NVIDIA Tesla P40	Pascal
NVIDIA Tesla P4	Pascal

Table 5.5 Supported NVIDIA Data Center Products (continued)

Product	GPU Architecture
<b>Tesla M-Series Products</b>	
NVIDIA Tesla M60	Maxwell
NVIDIA Tesla M6	Maxwell
<b>Tesla K-Series Products</b>	
NVIDIA Tesla K80	Kepler
NVIDIA Tesla K40(m/c/s/st/t)	Kepler
NVIDIA Tesla K20(x/c/m/Xm/s)	Kepler
NVIDIA Tesla K10	Kepler

## Supported NVIDIA Notebook GPUs

The notebook driver is part of the NVIDIA Verde Notebook Driver Program, and can be installed on supported NVIDIA notebook GPUs. However, please note that your notebook original equipment manufacturer (OEM) provides certified drivers for your specific notebook on their website. NVIDIA recommends that you check with your notebook OEM about recommended software updates for your notebook. OEMs may not provide technical support for issues that arise from the use of this driver.

The following tables list the NVIDIA notebook products supported by the Release 465 driver, version 465.89:

Table 5.6 NVIDIA Quadro Notebook GPU Support

Notebook Products	GPU Architecture
Quadro RTX 6000	Turing
Quadro RTX 5000	Turing
Quadro RTX 4000	Turing
Quadro RTX 3000	Turing
Quadro T2000	Turing
Quadro T1000	Turing
Quadro P5200	Pascal
Quadro P5000	Pascal
Quadro P4200	Pascal
Quadro P4000	Pascal
Quadro P3200	Pascal
Quadro P3000	Pascal



Table 5.6 NVIDIA Quadro Notebook GPU Support (continued)

Notebook Products	GPU Architecture
Quadro P620	Pascal
Quadro P520	Pascal
Quadro P600	Pascal
Quadro P500	Pascal
Quadro M5500	Maxwell
Quadro M5000M	Maxwell
Quadro M4000M	Maxwell
Quadro M3000M	Maxwell
Quadro M2000M	Maxwell
Quadro M1000M	Maxwell
Quadro M600M	Maxwell
Quadro M500M	Maxwell
Quadro K2200M	Maxwell
Quadro K620M	Maxwell

## Supported Languages

The Release 465 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 is approximately 1.5x the size of the installation download to accommodate extracted and temporary files.

## 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.

## Notebooks

- ▶ Check to make sure that your notebook has a supported GPU (see “Supported NVIDIA Notebook GPUs” on page 28).
- ▶ It is recommended that you back up your current system configuration.
- ▶ If you own a Dell Inspiron 1420, Dell XPS M1330, or Dell XPS M1530, or Dell LatitudeD630 or D630c, it is highly recommended that you first install this [Dell software update](#).

## 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



**Obtaining NVIDIA Drivers:** Microsoft now provides only SHA-2 signed drivers. If your Windows 7 system is NOT equipped to detect SHA-2, you need to install the SHA-2 update support patches. For details as well as access to SHA-2 update support patches, see the Microsoft KBA “[2019 SHA-2 Code Signing Support requirement for Windows and WSUS](#)”.

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 NVIDIA Tesla Compute Cluster Mode

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

- ▶ “About Tesla Compute Cluster Mode” on page 32
- ▶ “Operating on Systems with non-TCC NVIDIA GPUs” on page 34
- ▶ “Setting TCC Mode” on page 34

## 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 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 A.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/Pascal)	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 NON-INFRINGEMENT, 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.

## 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.

## Third Party Notice

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

Clang & LLVM:

Copyright (c) 2003-2015 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 NON-INFRINGEMENT. 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.

## Vulkan Notice

Copyright (c) 2015-2016 The Khronos Group Inc.

Copyright (c) 2015-2016 LunarG, Inc.

Copyright (c) 2015-2016 Valve Corporation

The Vulkan Runtime is comprised of 100% open source components (MIT, and Apache 2.0). The text of such licenses is included below along with the copyrights.

ALL INFORMATION HERE IS PROVIDED "AS IS." LUNARG MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, WITH REGARD TO THIS LIST OR ITS ACCURACY OR COMPLETENESS, OR WITH RESPECT TO ANY RESULTS TO BE OBTAINED FROM USE OR DISTRIBUTION OF THE LIST. BY USING OR DISTRIBUTING THIS LIST, YOU AGREE THAT IN NO EVENT SHALL LUNARG BE HELD LIABLE FOR ANY DAMAGES WHATSOEVER RESULTING FROM ANY USE OR DISTRIBUTION OF THIS LIST, INCLUDING, WITHOUT LIMITATION, ANY SPECIAL, CONSEQUENTIAL, INCIDENTAL OR OTHER DIRECT OR INDIRECT DAMAGES.

Softfloat Notice

Portions of the driver use the Softfloat floating point emulation library.

Softfloat Release 3e (<http://www.jhauser.us/arithmetric/SoftFloat.html>) is provided under the following terms:

Copyright 2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018 The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1.Redistributions of source code must retain the above copyright notice, this list of conditions, and the following disclaimer.
- 2.Redistributions in binary form must reproduce the above copyright notice, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3.Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS", AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

=====**Apache 2.0**=====

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

=====**MIT**=====

Copyright (c) 2009 Dave Gamble

Copyright (c) 2015-2016 The Khronos Group Inc.

Copyright (c) 2015-2016 Valve Corporation

Copyright (c) 2015-2016 LunarG, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in 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:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

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 AUTHORS 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 IN THE SOFTWARE.

=====





## Trademarks

NVIDIA, the NVIDIA logo, NVIDIA nForce, GeForce, NVIDIA Quadro, are registered trademarks or trademarks of NVIDIA Corporation in the United States and/or other countries.

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

OpenGL® and the oval logo are trademarks or registered trademarks of Silicon Graphics, Inc. in the United States and/or other countries worldwide. Additional license details are available on the SGI website.

OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

Vulkan and the Vulkan logo are trademarks of the Khronos Group Inc.

Other company and product names may be trademarks or registered trademarks of the respective owners with which they are associated.

## Copyright

© 2020 NVIDIA Corporation. All rights reserved.