



# **NVIDIA RTX Enterprise Drivers**

## **Release 470, Version 474.82**

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

Release Notes

# Table of Contents

Introduction to Release Notes.....	1
Structure of the Document .....	2
Changes in This Edition .....	2
Release 470 Driver Changes.....	3
About Release 470.....	3
Existing Support.....	3
New Features .....	3
Discontinued and Unsupported Features.....	6
Limitations in This Release.....	6
External Graphics .....	7
Deep Color (10 bpc) on HDR-Capable Monitors Under Windows 10 RS2 and RS3..	7
OpenCL 3.0 Known Issues.....	7
Cross-Adapter Clone Mode .....	7
Optimus Systems .....	8
Version 474.82 Highlights .....	8
New Features .....	8
Fixed Issues.....	8
Open Issues.....	8
Windows Issues.....	8
Application Issues (Not Caused by NVIDIA Driver).....	8
Changes in Earlier Versions.....	9
Version 474.44 .....	9
Version 474.30.....	9
Version 474.14 .....	9
Version 474.04.....	9
Version 473.81 .....	9
New Features .....	9
Fixed Issues.....	9
Open Issues.....	10
Windows Issues.....	10
Application Issues (Not Caused by NVIDIA Driver).....	10
Version 473.62 .....	10
Version 472.96 .....	10
Version 472.84.....	10
Fixed Issues.....	11
Version 472.39.....	11

Fixed Issues.....	11
Version 472.12.....	11
Fixed Issues.....	11
Version 471.68.....	12
Fixed Issues.....	12
Version 471.41.....	12
Fixed Issues.....	12
Version 471.11.....	12
Fixed Issues.....	12
The Release 470 Driver.....	13
Driver Security.....	13
Restricting and Enabling Access to GPU Performance Counters.....	13
Enabling Access to GPU Performance Counters Using NVIDIA Control Panel.....	13
Restricting and Enabling Access to GPU Performance Counters Across an Enterprise Using Scripts.....	14
Advanced Driver Information.....	14
Turning Off V-Sync to Boost Performance.....	14
NVIDIA Application Configuration Engine (ACE).....	15
Using Windows Display Driver Model with Tesla GPU GOMs.....	15
Tesla GPU Operation Modes.....	15
WDDM and TCC Driver Models.....	16
Compatibility Between GOM and Driver Models.....	16
Using the WDDM Driver Model.....	16
Multi-OS – GPU Assignment in System Virtualization.....	17
Hardware Platform Requirements.....	17
Supported Hypervisors.....	17
Supported Graphics Cards.....	17
Notes and Known Issues.....	18
VMware.....	18
Citrix XenServer.....	18
Known Product Limitations.....	19
Issues Installing NVIDIA Control Panel from the Windows Store.....	19
Some APIs do not Report Total Available Graphics Memory Correctly.....	19
NVIDIA Action for Some GeForce-based Systems.....	20
When TAG Reporting Would Not Be Limited.....	20
Using HDMI/DisplayPort Audio with Displays That Have High Native Resolution.....	20
Using HDMI/DisplayPort Audio in Dualview or Clone Mode Configurations.....	21
GPU Runs at a High Performance Level in Multi-Display Modes.....	21
Applying Workstation Application Profiles.....	21
Hardware and Software Support.....	21

Supported Operating Systems.....	21
Supported NVIDIA Workstation GPUs .....	22
NVIDIA Quadro, NVS, and RTX Products.....	22
NVIDIA Quadro Sync II.....	23
NVIDIA Quadro Sync Products.....	24
NVIDIA Quadro Blade/Embedded Graphics Board .....	24
NVIDIA Data Center Products .....	25
Supported NVIDIA Notebook GPUs.....	26
Supported Languages.....	27
Driver Installation .....	27
Minimum Hard Disk Space.....	27
Before You Begin.....	27
nTune.....	27
Notebooks .....	27
SLI Mosaic Mode.....	28
Installation Instructions.....	28
Appendix A. NVIDIA Tesla Compute Cluster Mode.....	29
About Tesla Compute Cluster Mode .....	29
TCC Overview .....	29
Benefits .....	29
TCC Does Not Support Graphics Acceleration .....	29
Running CUDA Applications .....	29
Operating on Systems with non-TCC NVIDIA GPUs.....	30
Setting TCC Mode.....	30

# Table of Tables

- Table 1. GOM and Drive Model Compatibility ..... 16
- Table 2. Supported Hypervisors ..... 17
- Table 3. GPUs Supported for Device Passthrough..... 17
- Table 4. Supported NVIDIA Quadro and NVS Products ..... 22
- Table 5. Supported NVIDIA Quadro Sync II Products ..... 23
- Table 6. Supported NVIDIA Quadro Sync Products..... 24
- Table 7. Supported NVIDIA Quadro Blade/Embedded Graphics Board Series ..... 24
- Table 8. Supported NVIDIA Data Center Products ..... 25
- Table 9. Supported Notebook GPUs ..... 26
- Table 10. Default TCC Mode ..... 30

---

# Introduction to Release Notes

This edition of Release Notes describes version 474.82 of the release 470 family of Microsoft Windows 10 and Windows 11 drivers for:

- > NVIDIA RTX®
- > NVIDIA® Quadro®
- > NVIDIA® NVS™
- > Data Center
- > Quadro Notebook

Release 470 is a **Production Branch** (PB) release. PB drivers are designed and tested to provide long-term stability and availability, making these drivers ideal for enterprise customers and other users who require application and hardware certification from ISVs and OEMs respectively.

PB drivers are a superset of the NVIDIA Studio Drivers and provide all the benefits of the Studio Driver of the same version in addition to NVIDIA RTX-specific enhancements and testing.

NVIDIA provides these notes to describe performance improvements and bug fixes in each documented version of the drivers.



**Note:** For convenience, these release notes use brief forms to refer to certain software components.

- > **Windows** refers to Microsoft Windows 10 and Microsoft Windows 11, except where stated otherwise. Version 474.82 of these drivers supports Windows 10 and Windows 11.
- > **The driver** means “the NVIDIA GPU drivers in version 474.82 of release 470,” or “the driver in version 474.82 that supports your hardware and software configuration.”

# Structure of the Document

This document is organized in the following sections:

- > [Chapter 2. Release 470 Driver Changes](#) gives a summary of changes and of fixed and open issues in this version.
- > [Chapter 3. The Release 470 Driver](#) describes the NVIDIA products and languages supported by this driver, the driver's system requirements, and how to install the driver.
- > [Appendix A. NVIDIA Tesla Compute Cluster Mode](#).

## Changes in This Edition

This edition of the *NVIDIA RTX Enterprise Drivers Release Notes* for Windows includes information about NVIDIA graphics driver version 474.82, and lists changes made to the driver since version 474.44.

These changes are discussed beginning with the chapter [Release 470 Driver Changes](#) on page 2.

---

# Release 470 Driver Changes

This chapter describes open issues for version 474.82, and resolved issues and driver enhancements for versions of the Release 470 driver up to version 474.82. It describes [Version 474.82 Highlights: New Features, Fixed Issues, and Open Issues](#).

## About Release 470

This section contains information that applies to all versions of release 470.

## 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
- > CUDA 11.4

This driver installs NVIDIA RTX Desktop Manager version 203.05.

## New Features

This section summarizes driver changes in release 470 (since release 465).

Release 470 is the latest Production Branch release of the NVIDIA RTX® Enterprise Driver. This new Production Branch provides improvements over the previous branch in the areas of application performance, API interoperability (e.g., OpenCL™/Vulkan®), and application power management.

The specific new features in release 470 are:

- > Added support for Microsoft Windows 11
- > Added support for NVIDIA CUDA® 11.4
- > Added NVIDIA® GPUDirect® support for Windows in NVIDIA Rivermax SDK



- > Added NVAPI support to configure desired 3-pin DIN stereo output and target display head for synchronization
- > Added support for WDDM 3.0
- > Added 10-bit color support for custom resolutions
- > Security updates:
  - Version 474.64: See *Security Bulletin: NVIDIA GPU Display Driver – October 2023*, listed on the [NVIDIA Product Security page](#).
  - Version 474.44: See *Security Bulletin: NVIDIA GPU Display Driver – June 2023*, listed on the [NVIDIA Product Security page](#).
  - Version 474.30: See *Security Bulletin: NVIDIA GPU Display Driver – March 2023*, listed on the [NVIDIA Product Security page](#).
  - Version 474.14: See *Security Bulletin: NVIDIA GPU Display Driver – November 2022*, listed on the [NVIDIA Product Security page](#).
  - Version 473.81: See *Security Bulletin: NVIDIA GPU Display Driver – August 2022*, posted shortly after the release date of this version of the driver, and then listed on the [NVIDIA Product Security page](#).
  - Version 472.39 added security updates for driver components.
    - a) Refer to the NVIDIA Security Bulletin 5230 for details.
  - Version 471.41 added security updates for driver components.
    - b) See the NVIDIA Security Bulletin 5211 for details of updates in 471.41.
- > Added support for the following NVIDIA GPUs:
  - NVIDIA RTX A4500
  - NVIDIA RTX A2000 12GB
  - NVIDIA T1000 8GB
  - NVIDIA T400 4GB
  - NVIDIA RTX A2000 GPU
- > Added support for the new product NVIDIA® Omniverse™
- > Added features to further assist 3D artists and designers in collaborative creative workflows:
  - Omniverse VR for real-time ray traced VR
  - Omniverse Remote for connecting artists to view and interact with fully ray traced assets through iOS and Android devices
  - Omniverse Showroom for quick and easy demonstrations of Omniverse technology
- > Application profile updates:
  - **Adobe Substance 3D Stager**: Fixed the Optimus profile to run on the discrete GPU by default
  - **IntelliCAD**: Updates to address stability issues

- > Studio Application updates:
  - **Blender 3.0:** Support for new features, including native USD and Omniverse integration, utilizing NVIDIA RTX GPU Tensor Cores for NVIDIA OptiX™ AI denoising and RT Cores for NVIDIA OptiX ray tracing.
  - **Isotropix Clarisse:** Improved performance with Angie renderer, enabling interaction with large scenes using exclusive ray traced acceleration for NVIDIA RTX GPUs
  - **Powerdirector:** Integration for NVIDIA Broadcast and AI tools for video and audio denoising in post-production
  - **NVIDIA Texture Tools:** Support for new API allowing users to create compressed, high-quality texture files with substantially reduced memory usage.
  - **Twinmotion 2022.1:** Support for accelerated RTX ray racing by RT cores on NVIDIA RTX GPUs
  - **OctaneRender 2021.1:** Motion blur scenes are accelerated on NVIDIA RTX GPUs, allowing designers to achieve realistic reflections with motion blur.
  - **V-Ray 5 Update 2:** Support for new V-Ray Decal feature that adds decals onto surfaces at any angle without extra UVW work or disturbing any underlying materials, and a new Sharpen/Blur layer for more artistic control within applications.
  - **Houdini 19:** Support for NVIDIA OptiX in NVIDIA RTX GPUs in new Karma XPU render engine.
  - **Adobe Lightroom and Lightroom Classic:** support for Select Subject and Select Sky, which are AI-powered masking tools for isolating people and capturing skylines respectively.
  - **Adobe Substance 3D:** NVIDIA RTX-accelerated ray tracing support speeds up interactive image rendering and improves performance over CPU alone.
  - **Adobe Premiere Elements:** Support for Auto Reframe feature to automatically frame the most important areas in shots.
  - **DXO Photolab:** Tensor core AI integration for DeepPRIME technology to reduce digital noise, especially in photos taken in low light conditions.
  - **Gamecaster Live:** NVIDIA Broadcast support for different levels of noise removal by audio source to further fine tune output
  - **NVIDIA Canvas:** Added support for custom styles.
  - **Jianying video editing software:** Added support.
  - **Maxine AR SDK:** Updated to improve quality and stability of AI-driven body tracking in applications such as Notch.
  - **Topaz Sharpen AI v3.2:** Refinements added to AI support by RTX GPUs and Tensor Cores to accelerate motion blur and further reduce image artifacts.
  - **Reallusion Character Creator v3.43:** Now allows artists to export characters from Character Creator to Omniverse as USD files with Audio2Face-compliant meshes, allowing facial and lip animations to be completely AI-driven solely from voice input.

- **Capture One 21 v14.3.0:** Increased performance for new Magic Brush tool to create complex masks for layer editing.
- **Unity 2021.2 Beta:** Native DLSS (Deep Learning Super Sampling) SDK 2.2.1 support enables optimal image quality in "Auto" mode.
- **Unreal Engine 4.27:** Eye-tracked foveated rendering brings improved performance; GPU Lightmass baking built on RTX ray tracing provides improved lighting and detail control.
- **NVIDIA Omniverse:** Platform for 3D content creation and collaboration, built from the ground up to be easily extensible and customizable with a modular development framework.
- **Enscape 3.1:** Support for new VulkanRT renderer and NVIDIA Real-Time Denoiser to provide smoother viewport visibility, as well as DLSS support to render at lower resolution for higher framerates.
- **Marmoset Toolbag 4.03:** RTX-accelerated AI denoising added to allow artists to quickly visualize materials with photorealistic lighting and shadows.
- **Topaz Video Enhance AI:** Added Slow Motion, a new RTX GPU Tensor Core-powered AI feature that generates high quality, smooth slow-motion capture with minimal artifacts.
- **Pixar Animation Studios RenderMan 24:** Added RenderMan XPU, a look-development focused GPU-accelerated ray tracer; together with AI denoising in the viewport, RenderMan XPU enables artists to interactively create their art and view an image that is predictive of the final frame render.
- **Discord:** Support for NVDEC for accelerated video decoding, reducing video resources for better gaming/app performance and stability
- **NVIDIA Canvas**, enabling NVIDIA RTX and NVIDIA® Quadro RTX™ users to turn doodles into stunning landscapes; beta download available here: <https://www.nvidia.com/en-us/studio/canvas/>.
- **Adobe Substance 3D Stager and 3D Modeler:** Support for new applications, with RTX acceleration for ray-traced rendering and light baking, plus AI-enhanced features.

## Discontinued and Unsupported Features

Release 470 will be the last driver branch to support desktop GPUs from the NVIDIA Kepler™ architecture generation.

## 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. Disconnecting the external GPU while applications are running is not advised.
- > **Mixed GeForce/Quadro products:** Mixed GeForce/Quadro products are supported (NVIDIA® GeForce® GPU with NVIDIA® Quadro® eGPU, or Quadro GPU with GeForce eGPU), but require 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 connected, Windows 10 may initially set the color depth to 8 bpc (bits per channel). You can then enable deep color (10 bpc) through 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 reverts to the 8 bpc color space.

## OpenCL 3.0 Known Issues

- > **Device-Side-Enqueue related queries:** Such queries may return 0 values, although the kernel can safely use corresponding built-ins.  
This is in accordance with conformance requirements described on the Khronos Group web site page [The OpenCL Specification](#).
- > **Certain denormalized results not flushed to zero:** The denormalized results for some math functions that were flushed to zero when `cl-fastrelaxed-math` is passed as a compiler option on OpenCL 1.2 drivers are not flushed to zero with OpenCL 3.0 drivers.
- > **Shared virtual memory:** The current implementation of shared virtual memory is limited to 64-bit platforms.

## Cross-Adapter Clone Mode

Under Windows, 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 NVIDIA Control Panel under Windows 10.

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

## Optimus Systems

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

## Version 474.82 Highlights

This section provides highlights of version 474.82 of the driver.

### New Features

- > **Security updates:** See *Security Bulletin: NVIDIA GPU Display Driver – February 2024*, posted shortly after the release date of this version of the driver, and then listed on the [NVIDIA Product Security](#) page.

### Fixed Issues

This version of the driver has no critical issues to highlight.

### Open Issues

As with every released driver, version 474.82 of this driver has open issues and enhancement requests associated with it. This section includes lists of issues that are not fixed and features that are not implemented in this version. Some of the issues may not have been thoroughly investigated and, in fact, may not be NVIDIA issues. Others may have workarounds.

### Windows Issues

- > [Alchemist]: The application crashes when either changing the display resolution or switching to another display with a different resolution.  
As a workaround, avoid changing the resolution or disconnecting from or connecting to an external display with the application opened.

### Application Issues (Not Caused by NVIDIA Driver)

This section lists issues that affect the NVIDIA driver but are caused by applications.

- > [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 in 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.

## Changes in Earlier Versions

This section describes issues fixed and new features added in earlier versions of release 470.

### Version 474.44

Only security issues are fixed in version 474.44 of the NVIDIA release 470 driver for Windows.

### Version 474.30

Only security issues are fixed in version 474.40 of the NVIDIA release 470 driver for Windows.

### Version 474.14

Only security issues are fixed in version 474.14 of the NVIDIA release 470 driver for Windows.

### Version 474.04

Only security issues are fixed in version 474.04 of the NVIDIA release 470 driver for Windows.

### Version 473.81

## New Features

- > **Security updates:** See *Security Bulletin: NVIDIA GPU Display Driver – August 2022*, posted shortly after the release date of this version of the driver, and then listed on the [NVIDIA Product Security](#) page.

## Fixed Issues

This version of the driver has no fixed issues.

## Open Issues

As with every released driver, version 474.82 of this driver has open issues and enhancement requests associated with it. This section includes lists of issues that are not fixed and features that are not implemented in this version. Some of the issues may not have been thoroughly investigated and, in fact, may not be NVIDIA issues. Others may have workarounds.

### Windows Issues

- > [Alchemist]: The application crashes when either changing the display resolution or switching to another display with a different resolution.  
As a workaround, avoid changing the resolution or disconnecting from or connecting to an external display with the application opened.

### Application Issues (Not Caused by NVIDIA Driver)

This section lists issues that affect the NVIDIA driver but are caused by applications.

- > [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 in 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.

## Version 473.62

Only security issues are fixed in version 473.62 of the NVIDIA release 470 driver for Windows.

## Version 472.96

Only security issues are fixed in version 472.96 of the NVIDIA release 470 driver for Windows.

## Version 472.84

The following sections list the important changes and the most common issues resolved in this driver version.

## Fixed Issues

- > [REDCINE-X Pro]: Poor performance when using NVIDIA P620 and T400 GPUs.
- > [Redshift]: Denoiser alpha channel shows artifacts when using kernel prediction mode with "denoise alpha" enabled.
- > [Redshift3d][Ampere GPUs][NVIDIA OptiX Ray Tracing]: NVIDIA OptiX may crash.
- > [Dassault 3DEXPERIENCE][OpenGL/Vulkan Interoperability]: The driver may crash when using dedicated memory allocation.
- > [NVIDIA RTX Experience][DaVinci Resolve]: Names of games in the application tiles are incorrect.
- > [AutoCAD][Ampere GPUs]: Recreating an `ID3D12CommandQueue` causes subsequent compute shader writes to fail.
- > [Maxon C4D]: Dark flickering appears in the viewport.

## Version 472.39

The following sections list the important changes and the most common issues resolved in this driver version.

### Fixed Issues

- > [Autodesk VRED]: Flickering occurs in the application when rendering in 16-bit float.
- > [AutoCAD][Notebook]: The application might crash when switching between drawings.
- > [Altair]: Visualization issues occur when disabling multisampling.
- > [Redshift3D]: GPU clocks remain at maximum rate even when the application is idle.

## Version 472.12

The following sections list the important changes and the most common issues resolved in this driver version.

### Fixed Issues

- > [Rheinmetall][NVIDIA Turing™ GPUs]: The GPU driver may crash when using MVR with VRWorks.
- > [Dassault 3DEXPERIENCE][Quadro M2000]: OpenGL/Vulkan interoperability-related errors may occur.
- > [Redshift3D]: GPU clocks "stuck" at maximum rate.



## Version 471.68

The following sections list the important changes and the most common issues resolved in this driver version.

### Fixed Issues

- > [Luxmark][RTX A6000]: Application performance decreases on the RTX A6000.
- > [Topaz Video Enhance AI]: The application fails to process video with driver TDR error.
- > [D5 Render]: Black corruption overlay occurs with NVIDIA DLSS ON.

## Version 471.41

The following sections list the important changes and the most common issues resolved in this driver version.

### Fixed Issues

- > [Chaos] Artifacts depending on scratch buffer content appear in kernel prediction modes.
- > [Catia V6 3DEXPERIENCE 2017x]: Corruption appears in the viewport.
- > [Direct2IR] [NVIDIA Maxwell® and NVIDIA Pascal™ GPUs]: NVIDIA OptiX crashes (pointing to an invalid instruction) when D2IR is used with V-Ray.

## Version 471.11

The following sections list the important changes and the most common issues resolved in this driver version.

### Fixed Issues

- > [Blender]: There is image corruption with CUDA/OpenGL interoperability on multiple GPUs when HWS is active.
- > V-Ray performance drop occurs.
- > ANSYS is unable to use NVML compiled with CUDA 10.2.

---

# The Release 470 Driver

This chapter discusses matters that apply to all versions of release 470, such as driver security and known product limitations.

## 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 you run as administrator, do not elevate UAC privileges for activities or programs that don't need them.

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

## Restricting and 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. Developers need GPU performance counters 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-administrative users.

For more information about CVE-2018-6260, see NVIDIA Security Bulletin 4772, [NVIDIA GPU Display Driver - February 2019](#).

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

Access to GPU performance counters can be enabled for non-administrative users who need to use NVIDIA developer tools. You can enable access to GPU performance counters through the **NVIDIA Control Panel → Developer → Manage GPU Performance**

**Counters** page. See the **Developer → Manage GPU Performance Counters** section of NVIDIA Control Panel Help for instructions.



**Note:** Keep access to GPU performance counters disabled for non-administrative users who do not need to use NVIDIA developer tools.

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

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



**Caution:** This operation should be performed only by enterprise administrators. Changes to the registry must be made with care. Changing the registry incorrectly can cause system instability.

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

The data value 1 restricts access to administrative users; 0 allows access to all users.

You must reboot for the changes to take effect.

## Advanced Driver Information

This section clarifies instructions for successfully accomplishing the several tasks for users with special requirements and advanced skills.

### Turning Off V-Sync to Boost Performance

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

1. Open 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 preset** dropdown menu, select **Base profile**.
4. From the **Settings** list box, select **Vertical sync** and change its value to **Force off**; then click **Apply**.

- From the **Global presets** dropdown 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**.



**Note:** Be sure to close NVIDIA Control Panel completely, as leaving it open affects 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 NVIDIA Control Panel.

See the NVIDIA Knowledge Base article [How do I select application profiles for my workstation applications?](#) for more information about this feature.

## Using Windows Display Driver Model with Tesla GPU GOMs

**Windows Display Driver Model** is a display driver architecture defined by Microsoft to ensure compatibility between Windows and third-party display drivers. WDDM drivers are certified by Microsoft to comply with WDDM.

### Tesla GPU Operation Modes

Tesla K20X and K20 GPU accelerators offer the ability to specify the GPU operation mode using `nvidia-smi` or NVML. See [NVIDIA Management Library \(NVML\)](#) on the NVIDIA developer web site for more information.

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 GOMs are supported:

- > **Compute only:** For running compute tasks only.  
Tesla K20X and K20 accelerator boards ship in this mode.
- > **Low double precision:** For graphics applications that don't require high-bandwidth double precision.  
NVIDIA recommends this GOM for workloads that are not sensitive to double precision but need graphics capabilities.
- > **All On:** NVIDIA recommends this GOM 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 must 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 applications; 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 Drive 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: <ul style="list-style-type: none"> <li>&gt; X11 and use cases that require X11 (GLInterop, OCL conformance and VIPER)</li> <li>&gt; 32-bit Windows OS</li> </ul>
Low double precision	YES	YES	All use cases supported.

The “compute only” GOM is supported only on the TCC driver model. The WDDM driver model supports only GOM modes that enable graphics.

The “compute only” GOM and WDDM are incompatible and must not be used simultaneously.

Tesla K20 Active Accelerators for workstations ship in “compute only” mode, and their mode cannot be changed. Consequently, you must 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.



**Caution:** Do not attempt to specify the driver model by editing the registry. Doing so could result in the “compute only” GOM and WDDM being configured simultaneously, which could 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` and NVIDIA Control Panel’s **Manage Maximus Settings** page. These tools provide warnings in the case of a conflict.

## Multi-OS – GPU Assignment in System Virtualization

On systems with two or more graphics cards, this driver supports a hypervisor's ability to assign GPUs directly to guest virtual machines (VMs). Such direct assignment allows each guest VM to run its own operating system with its own GPU and driver. Direct 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 have 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 features vary by processor family, product, and system, and should be verified at the manufacturer's website.

### Supported Hypervisors

This table shows the supported hypervisors.

**Table 2. Supported Hypervisors**

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

This table shows GPUs that are supported for device passthrough.

**Table 3. GPUs Supported for Device Passthrough**

GPU Family	Supported Boards
NVIDIA Ampere	NVIDIA Data Center: A100, A40, A30, A16, A10, A2 NVIDIA RTX®: A6000, A5500, A5000, A4500, A4000
NVIDIA Turing™	NVIDIA® Quadro®: RTX 8000, RTX 6000, RTX 5000, RTX 4000 NVIDIA Data Center: T4
NVIDIA Volta™	Quadro: GV100 NVIDIA Data Center: V100
NVIDIA Pascal™	Quadro: P2000, P4000, P5000, P6000, GP100 Tesla: P100, P40, P4

GPU Family	Supported Boards
NVIDIA Maxwell®	Quadro: K2200, M2000, M4000, M5000, M6000, M6000 24GB Tesla: M60, M10, M6
NVIDIA Kepler™	Quadro: K2000, K4000, K4200, K5000, K5200, K6000 Tesla: 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 switch in the PCIe hierarchy above a PCIe device that is to be used for passthrough. If ACS is not present, ESX does 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 with this command:
 

```
C...> esxcfg-advcfg -k true disableACScheck
```
- > 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.  
You need not disable MSI translation for ESXi 6.0 and later.

### Citrix XenServer

- > PCI passthrough may can 64-bit MMIO failures on XenServer 6.1:  
The dom0 kernel may relocate GPU BARs to an invalid location immediately beyond the end of physical RAM.  
There are two alternative workarounds for this condition:
  - Add `pci=use_crs` to the kernel parameters by running this command in dom0:
 

```
C:...> /opt/xensource/libexec/xen-cmdline --set-dom0 pci=use_crs
```
  - Install the hotfix in Citrix Support Knowledge Center article CTX213256, [XenApi Crashes with error: oom\\_kill\\_process, Memory leak for Dom0](#).

# Known Product Limitations

This section describes problems that will not be fixed. Typically, the source of such a problem is beyond the control of NVIDIA.

## Issues Installing NVIDIA Control Panel from the Windows Store

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

- > The download process from the Windows Store freezes at the “Starting download ...” stage.
- > NVIDIA Control Panel fails to download after you initiate the download from the notification popup that appears when you install the driver.

For assistance with installing 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, 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

**TAG memory, background:** 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)
- > Shared System Memory (system memory shared between the graphics subsystem and the CPU)

The value for each of these components is 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-bit addresses instead of 64-bit, and consequently do not report available graphics memory properly when the TAG would otherwise exceed 4 GB. This results in underreporting of available memory and potentially undesirable behavior for 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 NVIDIA® 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 Dedicated Video Memory.

This is a policy decision for 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, and for 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.

Because these and similar scenarios are common in many workstation applications, the NVIDIA driver avoids constraining TAG on all Quadro and Tesla-based systems, and on GeForce-based systems with more than 2.75 GB of video memory.

## Using HDMI/DisplayPort Audio with Displays That Have 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 TVs 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 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 DVI monitors 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.

---

<sup>1</sup> The WDDM guidelines dictate minimum and maximum values for the components, but the display driver may further constrain the values that are reported.

## Using HDMI/DisplayPort Audio in Dualview or Clone Mode Configurations

**With two audio-enabled ports:** In a multi-display configuration where both HDMI/DisplayPort audio ports are enabled, only the primary display provides audio.

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

## GPU Runs at a High Performance Level in Multi-Display Modes

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

## Applying Workstation Application Profiles

**Background:** A workstation application profile is a group software setting used to make NVIDIA display drivers deliver optimum performance with a specific application. Workstation application profiles also can implement workarounds for known application issues and bugs.

If a profile is available for a given application, NVIDIA recommends that you use it. Otherwise, the application is likely to exhibit incorrect behavior or reduced performance.

**Issue:** Configuration changes require you to 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 section discusses topics such as supported operating systems, supported GPUs, and supported languages.

### Supported Operating Systems

The release 470 driver, version 474.82, has been tested with:

- > Microsoft Windows 11

- > Microsoft Windows 10, 64-bit (versions 1909 and later)

## Supported NVIDIA Workstation GPUs

The following tables list the NVIDIA products supported by the version 474.82 of the driver.

- > [Table 4. Supported NVIDIA Quadro and NVS Products](#), on page 22
- > [Table 5. Supported NVIDIA Quadro Sync II Products](#), on page 23
- > [Table 6. Supported NVIDIA Quadro Sync Products](#), on page 24
- > [Table 7. Supported NVIDIA Quadro Blade/Embedded Graphics Board Series](#), on page 24
- > [Table 8. Supported NVIDIA Data Center Products](#), on page 25
- > [Table 9. Supported Notebook GPUs](#) on page 26

### NVIDIA Quadro, NVS, and RTX Products

**Table 4. Supported NVIDIA Quadro and NVS Products**

Product	GPU Architecture
NVIDIA RTX A6000	NVIDIA Ampere
NVIDIA RTX A5000	
NVIDIA RTX A4500	
NVIDIA RTX A4000	
NVIDIA RTX A2000 12GB	
NVIDIA RTX A2000	
NVIDIA Quadro RTX 8000	Turing
NVIDIA Quadro RTX 6000	
NVIDIA Quadro RTX 5000	
NVIDIA Quadro RTX 4000	
NVIDIA Quadro RTX 3000	
NVIDIA T1000 8GB	
NVIDIA T1000	
NVIDIA T600	
NVIDIA T400 4GB	
NVIDIA T400	
NVIDIA Quadro GV100	Volta
NVIDIA Quadro GP100	Pascal
NVIDIA Quadro P6000	
NVIDIA Quadro P5000	
NVIDIA Quadro P4000	

Product	GPU Architecture
NVIDIA Quadro P2200	
NVIDIA Quadro P2000	
NVIDIA Quadro P1000	
NVIDIA Quadro P600	
NVIDIA Quadro P400	
NVIDIA Quadro M6000 24GB	Maxwell
NVIDIA Quadro M6000	
NVIDIA Quadro M5000	
NVIDIA Quadro M4000	
NVIDIA Quadro M2000	
NVIDIA Quadro K6000	Kepler
NVIDIA Quadro K5200	
NVIDIA Quadro K5000	Kepler
NVIDIA Quadro K4200	
NVIDIA Quadro K4000	
NVIDIA Quadro K2200	Maxwell
NVIDIA Quadro K2000	Kepler
NVIDIA Quadro K1200	Maxwell
NVIDIA Quadro K620	
NVIDIA Quadro K600	Kepler
NVIDIA Quadro K420	
NVIDIA Quadro 410	
NVIDIA® NVS™ 810	Maxwell
NVIDIA NVS 510	Kepler

## NVIDIA Quadro Sync II

**Table 5. Supported NVIDIA Quadro Sync II Products**

Product	GPU Architecture
NVIDIA RTX A6000	NVIDIA Ampere
NVIDIA RTX A5000	
NVIDIA RTX A4500	
NVIDIA RTX A4000	
NVIDIA Quadro RTX 8000	Turing
NVIDIA Quadro RTX 6000	
NVIDIA Quadro RTX 5000	
NVIDIA Quadro RTX 4000	

Product	GPU Architecture
NVIDIA Quadro RTX 3000	
NVIDIA T1000 8GB	
NVIDIA T1000	
NVIDIA T600	
NVIDIA T400 4GB	
NVIDIA T400	
NVIDIA Quadro GV100	Volta
NVIDIA Quadro GP100	Pascal
NVIDIA Quadro P6000	
NVIDIA Quadro P5000	
NVIDIA Quadro P4000	

## NVIDIA Quadro Sync Products

**Table 6. Supported NVIDIA Quadro Sync Products**

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

## NVIDIA Quadro Blade/Embedded Graphics Board

**Table 7. Supported NVIDIA Quadro Blade/Embedded Graphics Board Series**

Product	GPU Architecture
NVIDIA RTX A4500 Embedded GPU	Ampere
NVIDIA RTX A2000 Embedded GPU	
NVIDIA Quadro RTX 5000	Turing
NVIDIA Quadro RTX 3000	
NVIDIA Quadro T 1000	
NVIDIA Quadro P5000	Pascal
NVIDIA Quadro P3000	

Product	GPU Architecture
NVIDIA Quadro P2000	Maxwell
NVIDIA Quadro P1000	
NVIDIA Quadro M5000 SE	
NVIDIA Quadro M3000 SE	

## 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 NVIDIA CUDA® C/C++ applications and libraries that rely on the CUDA C Runtime and/or CUDA Driver API.

**Table 8. Supported NVIDIA Data Center Products**

Product	GPU Architecture
<b>NVIDIA A-Series Products</b>	
NVIDIA A100	NVIDIA Ampere
NVIDIA A40	
NVIDIA A30	
NVIDIA A16	
NVIDIA A10	
NVIDIA A2	
<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	
NVIDIA Tesla P4	
<b>Tesla M-Series Products</b>	
NVIDIA Tesla M60	Maxwell
NVIDIA Tesla M6	
<b>Tesla K-Series Products</b>	
NVIDIA Tesla K80	Kepler
NVIDIA Tesla K40(m/c/s/st/t)	
NVIDIA Tesla K20(x/c/m/Xm/s)	
NVIDIA Tesla K10	

## 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 table lists the NVIDIA notebook GPUs supported by version 474.82 of the driver.

**Table 9. Supported Notebook GPUs**

<b>Notebook Product</b>	<b>GPU Architecture</b>
NVIDIA RTX A5500 Laptop GPU NVIDIA RTX A5000 Laptop GPU NVIDIA RTX A4500 Laptop GPU NVIDIA RTX A4000 Laptop GPU NVIDIA RTX A3000 Laptop GPU NVIDIA RTX A2000 8GB Laptop GPU NVIDIA RTX A2000 Laptop GPU	NVIDIA Ampere
NVIDIA T1200 Laptop GPU NVIDIA T600 Laptop GPU NVIDIA T550 Laptop GPU	Turing
Quadro RTX 6000 Quadro RTX 5000 Quadro RTX 4000 Quadro RTX 3000	Turing
Quadro T2000 Quadro T1000	Turing
Quadro P5200 Quadro P5000 Quadro P4200 Quadro P4000 Quadro P3200 Quadro P3000	Pascal
Quadro P620 Quadro P520 Quadro P600 Quadro P500	Pascal
Quadro M5500 Quadro M5000M Quadro M4000M Quadro M3000M Quadro M2000M Quadro M1000M	Maxwell

Notebook Product	GPU Architecture
Quadro M600M Quadro M500M	
Quadro K2200M Quadro K620M	Maxwell

## Supported Languages

The release 470 drivers support the following languages in NVIDIA Control Panel.

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

## Driver Installation

### Minimum Hard Disk Space

The hard disk space requirement is approximately 1.5× the size of the installation download to accommodate extracted and temporary files.

### Before You Begin

Certain applications and systems require preparation before you install the driver.

#### nTune

If you have previously installed NVIDIA nTune, NVIDIA recommends that you uninstall nTune before installing this driver. After the driver installation 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 26.

NVIDIA recommends that you back up your current system configuration before installation.



If you own a Dell Inspiron 1420, Dell XPS M1330, or Dell XPS M1530, or Dell Latitude D630 or D630c, NVIDIA strongly recommends that you apply all relevant updates to your system software before you install the driver.

## 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 cannot install properly.

## Installation Instructions

1. If a previous version of the driver is installed, uninstall it
2. Follow the instructions on the NVIDIA web site's [driver download page](#) to locate the appropriate driver to download, based on your hardware and operating system.
3. Click the driver download link. 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. Open the NVIDIA driver installation .exe file to launch the NVIDIA driver installer.
6. Follow the instructions in the NVIDIA driver installer to complete the installation.



**Note:** If you are installing the driver over a previous driver without first removing the previous driver, you must reboot your computer to complete the installation.

---

# Appendix A. NVIDIA Tesla Compute Cluster Mode

## 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 NVIDIA® 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 are available via Remote Desktop, both directly and via cluster management systems that rely on Remote Desktop.
- > NVIDIA GPUs are available to applications running as Windows services (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 NVIDIA CUDA®-DirectX/OpenGL® interoperability. A driver operating in TCC mode is a “non-display” driver, and NVIDIA GPUs that use this driver do 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 are available for CUDA applications running via services or Remote Desktop.

In this release, all GPUs are operated 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 do not run properly under TCC mode. Following are examples of CUDA SDK applications that are not supported:

bicubicTexture	Mandelbrot	simpleD3D10	smokeParticles
boxFilter	marchingCubes	simpleD3D10Texture	SobelFilter
cudaDecodeD3D9	Nbody	simpleD3D11Texture	volumeRender
cudaDecodeGL	oceanFFT	simpleD3D9	
fluidsD3D9	Particles	simpleD3D9Texture	
fluidsGL	postProcessGL	simpleGL	
imageDenoising	recursiveGaussian	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 installed over any NVIDIA display driver in the system. It then becomes the exclusive driver for all NVIDIA GPUs in the system.  
If the Tesla/Quadro driver is later uninstalled, the previous driver is not restored.
- > NVIDIA GPUs that do not support TCC mode appear as “VGA adapters” in 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:

```
C...> 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 10. Default TCC Mode**

Product	Default Mode
K20Xm/c inc	TCC
K20m	TCC Off
C2075	TCC
C2050	TCC Off
C2070	TCC Off

<b>Product</b>	<b>Default Mode</b>
S2050	TCC
M2050	TCC
M2070	TCC
M2075	TCC
M2070	Q TCC Off
M2090	TCC
Quadro (NVIDIA Kepler™, NVIDIA Maxwell®, and NVIDIA 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 or 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:

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.

## Copyright

© 2020–2024 NVIDIA CORPORATION & AFFILIATES. All rights reserved.