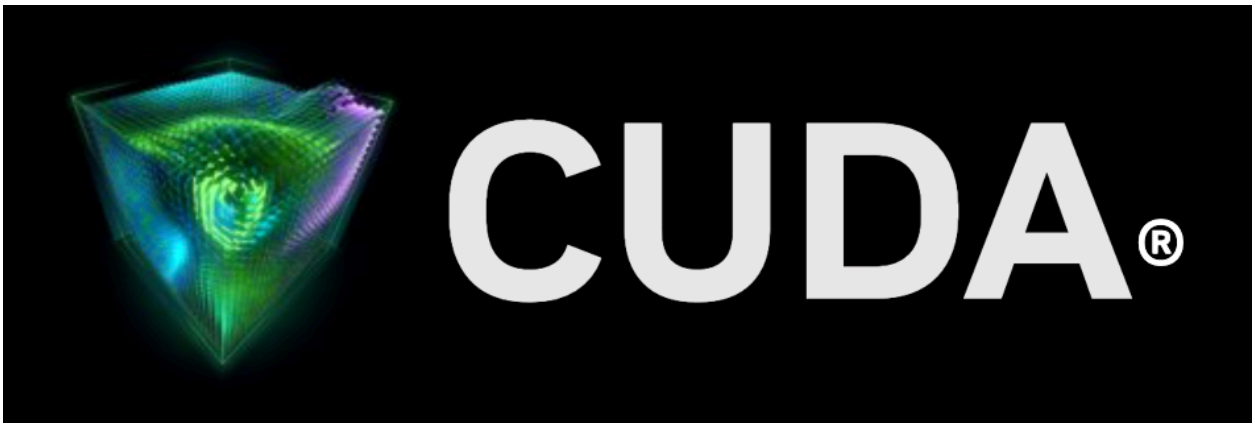


Installation Guide Windows

Release 13.5

NVIDIA Corporation

Nov 02, 2026

Contents

1 Overview **3**

2 Introduction **5**
2.1 System Requirements . . . 5
2.2 About This Document . . . 6

3 Installing CUDA Development Tools **7**
3.1 Verify You Have a CUDA-capable GPU . . . 7
3.2 Download the NVIDIA CUDA Toolkit . . . 7
3.3 Install the CUDA Software . . . 8
3.3.1 Uninstalling the CUDA Software . . . 11
3.4 Using Conda to Install the CUDA Software . . . 11
3.4.1 Conda Overview . . . 11
3.4.2 Installation . . . 11
3.4.3 Uninstallation . . . 11
3.4.4 Installing Previous CUDA Releases . . . 12
3.5 Use a Suitable Driver Model . . . 12
3.6 Verify the Installation . . . 12
3.6.1 Running the Compiled Examples . . . 13

4 Pip Wheels **15**
4.1 Prerequisites . . . 15
4.2 Procedure . . . 15
4.3 Packages . . . 16

5 Compiling CUDA Programs **17**
5.1 Compiling Sample Projects . . . 17
5.2 Sample Projects . . . 17
5.3 Build Customizations for New Projects . . . 18
5.4 Build Customizations for Existing Projects . . . 18

6 Additional Considerations **21**

7 Notices **23**
7.1 Notice . . . 23
7.2 OpenCL . . . 24
7.3 Trademarks . . . 24

CUDA Installation Guide for Microsoft Windows

Chapter 1. Overview

The CUDA Installation Guide for Microsoft Windows provides step-by-step instructions to help developers set up NVIDIA's CUDA Toolkit on Windows systems. It begins by introducing CUDA as NVIDIA's powerful parallel-computing platform—designed to accelerate compute-intensive applications by leveraging GPU capabilities. The guide details essential *system requirements*, including a CUDA-capable GPU and a compatible version of Windows, as well as the supported Visual Studio compilers. From there, it walks users through downloading the toolkit, installing both the CUDA driver and development tools, and verifying installation by compiling and running sample projects. With clear explanations geared toward both beginners and experienced developers, this guide ensures a smooth setup experience so you can start building GPU-accelerated applications on Windows right away.

Chapter 2. Introduction

CUDA® is a parallel computing platform and programming model invented by NVIDIA. It enables dramatic increases in computing performance by harnessing the power of the graphics processing unit (GPU).

CUDA was developed with several design goals in mind:

- Provide a small set of extensions to standard programming languages, like C, that enable a straightforward implementation of parallel algorithms. With CUDA C/C++, programmers can focus on the task of parallelization of the algorithms rather than spending time on their implementation.
- Support heterogeneous computation where applications use both the CPU and GPU. Serial portions of applications are run on the CPU, and parallel portions are offloaded to the GPU. As such, CUDA can be incrementally applied to existing applications. The CPU and GPU are treated as separate devices that have their own memory spaces. This configuration also allows simultaneous computation on the CPU and GPU without contention for memory resources.

CUDA-capable GPUs have hundreds of cores that can collectively run thousands of computing threads. These cores have shared resources including a register file and a shared memory. The on-chip shared memory allows parallel tasks running on these cores to share data without sending it over the system memory bus.

This guide will show you how to install and check the correct operation of the CUDA development tools.

2.1. System Requirements

To use CUDA on your system, you will need the following installed:

- A CUDA-capable GPU
- A supported version of Microsoft Windows
- A supported version of Microsoft Visual Studio
- NVIDIA CUDA Toolkit (available at https://developer.nvidia.com/cuda-downloads)

Supported Microsoft Windows® operating systems:

- Microsoft Windows 11 25H2
- Microsoft Windows 11 24H2
- Microsoft Windows 11 22H2-SV2
- Microsoft Windows 11 23H2

- Microsoft Windows 10 22H2
- Microsoft Windows Server 2022
- Microsoft Windows Server 2025

Table 1: Windows Compiler Support in CUDA 13.5

Compiler*	IDE	Native x86_64	Cross-compilation (32-bit on 64-bit)	C++ Dialect
MSVC Version 195x	Visual Studio 2026 18.x	YES	Not supported	C++14(default), C++17, C++20
MSVC Version 193x	Visual Studio 2022 17.x	YES	Not supported	C++14 (default), C++17, C++20
MSVC Version 192x	Visual Studio 2019 16.x	YES	Not supported	C++14 (default), C++17

* Support for Visual Studio 2015 is deprecated in release 11.1; support for Visual Studio 2017 is deprecated in release 12.5 and dropped in release 12.9.

32-bit compilation native and cross-compilation is removed from CUDA 12.0 and later Toolkit. Use the CUDA Toolkit from earlier releases for 32-bit compilation. CUDA Driver will continue to support running 32-bit application binaries on GeForce GPUs until Ada. Ada will be the last architecture with driver support for 32-bit applications. Hopper does not support 32-bit applications.

Support for running x86 32-bit applications on x86_64 Windows is limited to use with:

- CUDA Driver
- CUDA Runtime (cudart)
- CUDA Math Library (math.h)

2.2. About This Document

This document is intended for readers familiar with Microsoft Windows operating systems and the Microsoft Visual Studio environment. You do not need previous experience with CUDA or experience with parallel computation.

Chapter 3. Installing CUDA Development Tools

Basic instructions can be found in the Quick Start Guide. Read on for more detailed instructions.

The setup of CUDA development tools on a system running the appropriate version of Windows consists of a few simple steps:

- Verify the system has a CUDA-capable GPU.
- Download the NVIDIA CUDA Toolkit.
- Install the NVIDIA CUDA Toolkit.
- Test that the installed software runs correctly and communicates with the hardware.

3.1. Verify You Have a CUDA-capable GPU

You can verify that you have a CUDA-capable GPU through the **Display Adapters** section in the **Windows Device Manager**. Here you will find the vendor name and model of your graphics card(s). If you have an NVIDIA card that is listed in https://developer.nvidia.com/cuda-gpus, that GPU is CUDA-capable. The Release Notes for the CUDA Toolkit also contain a list of supported products.

The **Windows Device Manager** can be opened via the following steps:

1. Open a run window from the Start Menu
2. Run:

```
control /name Microsoft.DeviceManager
```

3.2. Download the NVIDIA CUDA Toolkit

The NVIDIA CUDA Toolkit is available at https://developer.nvidia.com/cuda-downloads. Choose the platform you are using and one of the following installer formats:

1. Network Installer: A minimal installer which later downloads packages required for installation. Only the packages selected during the selection phase of the installer are downloaded. This installer is useful for users who want to minimize download time.

2. Full Installer: An installer which contains all the components of the CUDA Toolkit and does not require any further download. This installer is useful for systems which lack network access and for enterprise deployment.

The CUDA Toolkit installs the tools needed to create, build and run a CUDA application as well as libraries, header files, and other resources. Starting with CUDA 13.1, the NVIDIA GPU driver on Windows is no longer bundled with the CUDA Toolkit and must be downloaded and installed separately from https://www.nvidia.com/drivers.

The download can be verified by comparing the MD5 checksum posted at https://developer.download.nvidia.com/compute/cuda/13.5.0/docs/sidebar/md5sum.txt with that of the downloaded file. If either of the checksums differ, the downloaded file is corrupt and needs to be downloaded again.

3.3. Install the CUDA Software

Before installing the toolkit, you should read the Release Notes, as they provide details on installation and software functionality.

Note: The driver and toolkit must be installed for CUDA to function. Starting with CUDA 13.1, the NVIDIA GPU driver on Windows is no longer bundled with the CUDA Toolkit and must be downloaded and installed separately from https://www.nvidia.com/drivers before installing the toolkit.

Note: The installation may fail if Windows Update starts after the installation has begun. Wait until Windows Update is complete and then try the installation again.

Graphical Installation

Install the CUDA Software by executing the CUDA installer and following the on-screen prompts.

Silent Installation

The installer can be executed in silent mode by executing the package with the `-s` flag. Additional parameters can be passed which will install specific subpackages instead of all packages. See the table below for a list of all the subpackage names.

Table 2: Possible Subpackage Names

Subpackage Name	Subpackage Description
Toolkit Subpackages (defaults to C:\Program Files\NVIDIA GPU Computing Toolkit\CUDA\v13.5)	
cccl_13.5	CUDA Core Compute Libraries (CCCL), including Thrust, CUB, and libcu++
crt_13.5	CUDA Runtime Library
ctadvisor_13.5	CUDA Threading/Compute Advisor
cublas_13.5 cublas_dev_13.5	cuBLAS Runtime Library
cuda_profiler_api_13.5	CUDA Profiler API

continues on next page

Table 2 – continued from previous page

Subpackage Name	Subpackage Description
cudart_13.5	CUDA Runtime Library
cudla_13.5	cuDLA Library (ARM64 installer only)
cufft_13.5	cuFFT Runtime Library
cufft_dev_13.5	
cuobjdump_13.5	Extracts information from cubin files
cupti_13.5	The CUDA Profiling Tools Interface for creating profiling and tracing tools that target CUDA applications
curand_13.5	cuRAND Runtime Library
curand_dev_13.5	
cusolver_13.5	cuSOLVER Runtime Library
cusolver_dev_13.5	
cusparse_13.5	cuSPARSE Runtime Library
cusparse_dev_13.5	
cuxxfilt_13.5	The CUDA cu++ filt Demangler Tool
nsight_compute_13.5	Nsight Compute
nsight_systems_13.5	Nsight Systems
nsight_vse_13.5	Installs the Nsight Visual Studio Edition plugin in all VS
npp_13.5	NPP Runtime Library
npp_dev_13.5	
nvcc_13.5	CUDA Compiler
nvdisasm_13.5	Extracts information from standalone cubin files
nvfatbin_13.5	Library for creating Fatbinaries at Runtime
nvjitlink_13.5	nvJitLink Library
nvjpeg_13.5	nvJPEG Library
nvjpeg_dev_13.5	
nvml_dev_13.5	NVML development libraries and headers
nvprune_13.5	Prunes host object files and libraries to only contain device code for the specified targets
nvptxcompiler_13.5	nvPTX Compiler Library
nvrtc_13.5	NVRTC Runtime Library
nvrtc_dev_13.5	NVRTC development library and headers
nvtx_13.5	NVTX on Windows

continues on next page

Table 2 – continued from previous page

Subpackage Name	Subpackage Description
nvvm_13.5	NVVM Compiler Library
opencl_13.5	OpenCL Library
sanitizer_13.5	Compute Sanitizer API
tileiras_13.5	CUDA Tile IR
visual_studio_integration_13.5	Installs CUDA project wizard and builds customization files in VS
Cross-Compilation Subpackages	
cublas_cross_13.5	cuBLAS libraries for cross-compilation
cudart_cross_13.5	CUDA Runtime libraries for cross-compilation
cufft_cross_13.5	cuFFT libraries for cross-compilation
cupti_cross_13.5	CUPTI libraries for cross-compilation
curand_cross_13.5	cuRAND libraries for cross-compilation
cusolver_cross_13.5	cuSOLVER libraries for cross-compilation
cusparse_cross_13.5	cuSPARSE libraries for cross-compilation
cuxxfilt_cross_13.5	cu++filt demangler for cross-compilation
npp_cross_13.5	NPP libraries for cross-compilation
nvcc_cross_13.5	CUDA Compiler components for cross-compilation
nvfatbin_cross_13.5	nvFatbin library for cross-compilation
nvjitlink_cross_13.5	nvJitLink library for cross-compilation
nvjpeg_cross_13.5	nvJPEG libraries for cross-compilation
nvml_dev_cross_13.5	NVML development libraries and headers for cross-compilation
nvptxcompiler_cross_13.5	nvPTX Compiler library for cross-compilation
nvrtc_cross_13.5	NVRTC libraries for cross-compilation
nvvm_cross_13.5	NVVM Compiler library for cross-compilation
opencl_cross_13.5	OpenCL library for cross-compilation

Subpackage names use the format `<component>_<major>.<minor>`, for example `nvcc_13.5`.

Subpackages whose names end in `_cross_13.5` install CUDA components for cross-compilation. All cross-compilation subpackages are available in both the x86_64 and ARM64 installers.

For example, to silently install only the CUDA compiler, the CUDA runtime, and CCCL:

```
<PackageName>.exe -s nvcc_13.5 cudart_13.5 cccl_13.5
```

Extracting and Inspecting the Files Manually

Sometimes it may be desirable to extract or inspect the installable files directly, such as in enterprise

deployment, or to browse the files before installation. The full installation package can be extracted using a decompression tool which supports the LZMA compression method, such as 7-zip or WinZip.

Once extracted, the CUDA Toolkit files will be in the `CUDAToolkit` folder, and similarly for CUDA Visual Studio Integration. Within each directory is a .dll and .nvi file that can be ignored as they are not part of the installable files.

Note: Accessing the files in this manner does not set up any environment settings, such as variables or Visual Studio integration. This is intended for enterprise-level deployment.

3.3.1. Uninstalling the CUDA Software

All subpackages can be uninstalled through the Windows Control Panel by using the Programs and Features widget.

3.4. Using Conda to Install the CUDA Software

This section describes the installation and configuration of CUDA when using the Conda installer. The Conda packages are available at https://anaconda.org/nvidia.

3.4.1. Conda Overview

The Conda installation installs the CUDA Toolkit. The installation steps are listed below.

3.4.2. Installation

To perform a basic install of all CUDA Toolkit components using Conda, run the following command:

```
conda install cuda -c nvidia
```

3.4.3. Uninstallation

To uninstall the CUDA Toolkit using Conda, run the following command:

```
conda remove cuda
```

3.4.4. Installing Previous CUDA Releases

All Conda packages released under a specific CUDA version are labeled with that release version. To install a previous version, include that label in the `install` command such as:

```
conda install cuda -c nvidia/label/cuda-11.3.0
```

Note: Some CUDA releases do not move to new versions of all installable components. When this is the case these components will be moved to the new label, and you may need to modify the install command to include both labels such as:

```
conda install cuda -c nvidia/label/cuda-11.3.0 -c nvidia/label/cuda-11.3.1
```

This example will install all packages released as part of CUDA 11.3.1.

3.5. Use a Suitable Driver Model

On Windows 10 and later, the operating system provides two driver models under which the NVIDIA Driver may operate:

- The WDDM driver model is used for display devices.
- The Tesla Compute Cluster (TCC) mode of the NVIDIA Driver is available for non-display devices such as NVIDIA Tesla GPUs and the GeForce GTX Titan GPUs; it uses the Windows WDM driver model.

TCC is enabled by default on most recent NVIDIA Tesla GPUs. To check which driver mode is in use and/or to switch driver modes, use the `nvidia-smi` tool that is included with the NVIDIA Driver installation (see `nvidia-smi -h` for details).

Note: Keep in mind that when TCC mode is enabled for a particular GPU, that GPU *cannot* be used as a display device.

Note: NVIDIA GeForce GPUs (excluding GeForce GTX Titan GPUs) do not support TCC mode.

3.6. Verify the Installation

Before continuing, it is important to verify that the CUDA toolkit can find and communicate correctly with the CUDA-capable hardware. To do this, you need to compile and run some of the included sample programs.

3.6.1. Running the Compiled Examples

The version of the CUDA Toolkit can be checked by running `nvcc -V` in a Command Prompt window. You can display a Command Prompt window by going to:

Start > All Programs > Accessories > Command Prompt

CUDA Samples are located in https://github.com/nvidia/cuda-samples. To use the samples, clone the project, build the samples, and run them using the instructions on the Github page.

To verify a correct configuration of the hardware and software, it is highly recommended that you build and run the `deviceQuery` sample program. The sample is located in https://github.com/NVIDIA/cuda-samples/tree/master/cpp/1_Utilities/deviceQuery and is built with CMake as described in the Windows build instructions in the CUDA Samples README.

This assumes that you used the default installation directory structure. If CUDA is installed and configured correctly, the output should look similar to *Figure 1*.

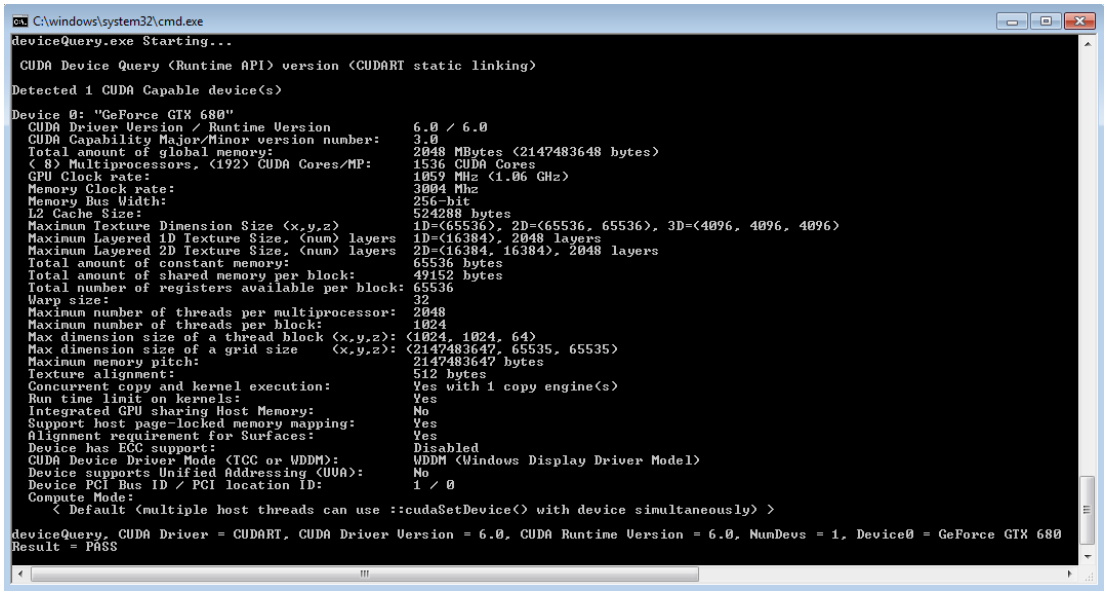


Figure 1: Valid Results from deviceQuery CUDA Sample

The exact appearance and the output lines might be different on your system. The important outcomes are that a device was found, that the device(s) match what is installed in your system, and that the test passed.

If a CUDA-capable device and the CUDA Driver are installed but `deviceQuery` reports that no CUDA-capable devices are present, ensure the device and driver are properly installed.

Running the `bandwidthTest` program, located in the same directory as `deviceQuery` above, ensures that the system and the CUDA-capable device are able to communicate correctly. The output should resemble *Figure 2*.

The device name (second line) and the bandwidth numbers vary from system to system. The important items are the second line, which confirms a CUDA device was found, and the second-to-last line, which confirms that all necessary tests passed.

If the tests do not pass, make sure you do have a CUDA-capable NVIDIA GPU on your system and make sure it is properly installed.

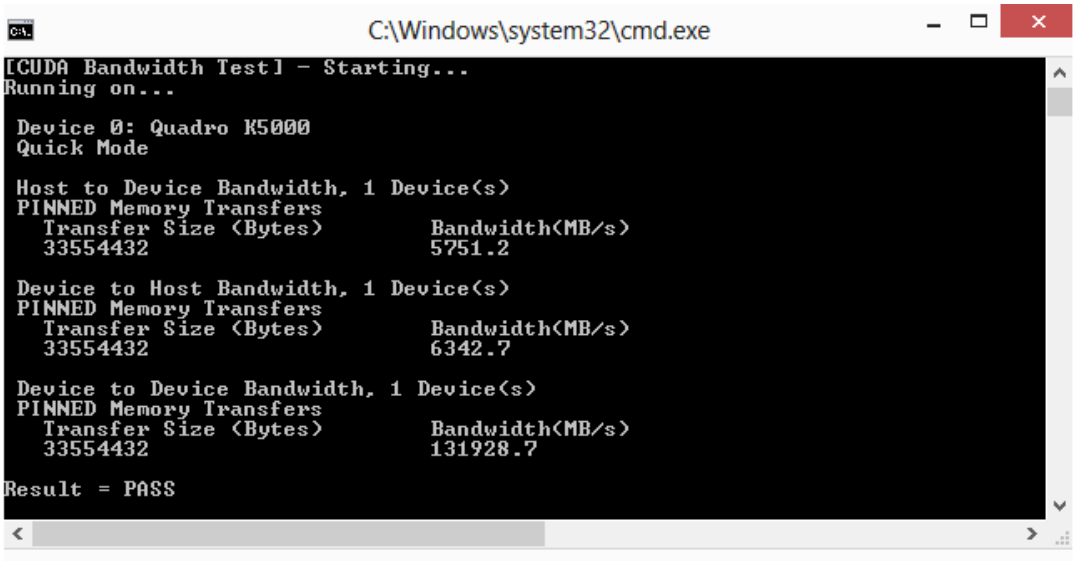


Figure 2: Valid Results from bandwidthTest CUDA Sample

To see a graphical representation of what CUDA can do, run the `particles` sample at

```
https://github.com/NVIDIA/cuda-samples/tree/master/cpp/2_Concepts_and_Techniques/
↪particles
```

Chapter 4. Pip Wheels

NVIDIA provides Python Wheels for installing CUDA through pip, primarily for using CUDA with Python. These packages are intended for runtime use and do not currently include developer tools (these can be installed separately).

Please note that with this installation method, CUDA installation environment is managed via pip and additional care must be taken to set up your host environment to use CUDA outside the pip environment.

4.1. Prerequisites

CUDA Python Wheels are published to the public Python Package Index (PyPI), so no additional package index needs to be configured. If your pip and setuptools Python modules are not up-to-date, then use the following command to upgrade these Python modules. If these Python modules are out-of-date then the commands which follow later in this section may fail.

```
py -m pip install --upgrade setuptools pip wheel
```

4.2. Procedure

Install the CUDA runtime package:

```
py -m pip install nvidia-cuda-runtime
```

Optionally, install additional packages as listed below using the following command:

```
py -m pip install nvidia-<library>
```

4.3. Packages

Starting with CUDA 13, wheel package names no longer include a CUDA major-version suffix (for example, `nvidia-cuda-runtime` replaces `nvidia-cuda-runtime-cu12`). The suffixed `-cu13` package names on PyPI are deprecated placeholders and should not be used.

The following packages will install the latest CUDA 13 version of the named component on Windows:

- `nvidia-cublas`
- `nvidia-cuda-cccl`
- `nvidia-cuda-cupti`
- `nvidia-cuda-nvcc`
- `nvidia-cuda-nvrtc`
- `nvidia-cuda-opencl`
- `nvidia-cuda-runtime`
- `nvidia-cuda-sanitizer-api`
- `nvidia-cufft`
- `nvidia-curand`
- `nvidia-cusolver`
- `nvidia-cusparse`
- `nvidia-npp`
- `nvidia-nvfatbin`
- `nvidia-nvjitlink`
- `nvidia-nvjpeg`
- `nvidia-nvml-dev`
- `nvidia-nvtx`

Chapter 5. Compiling CUDA Programs

The CUDA Samples in https://github.com/nvidia/cuda-samples are built with CMake. To build the Windows projects (for release or debug mode), use CMake to generate a Visual Studio solution file, `CUDA_Samples.sln`, as described in the Windows build instructions in the CUDA Samples README. Then open `CUDA_Samples.sln` in Visual Studio and build the desired configuration.

5.1. Compiling Sample Projects

The `deviceQuery` project is a good sample project to build and run. It is located in https://github.com/NVIDIA/cuda-samples/tree/master/cpp/1_Utilities/deviceQuery.

Build the program in Visual Studio and run the executable from the output directory specified in Visual Studio. If all works correctly, the output should be similar to *Figure 1*.

5.2. Sample Projects

The sample projects come in two configurations: debug and release (where release contains no debugging information) and different Visual Studio projects.

A few of the example projects require some additional setup.

These sample projects also make use of the `$CUDA_PATH` environment variable to locate where the CUDA Toolkit and the associated `.props` files are.

The environment variable is set automatically using the Build Customization `CUDA 13.5.props` file, and is installed automatically as part of the CUDA Toolkit installation process.

Table 3: CUDA Visual Studio .props locations

Visual Studio	**CUDA 13.5 .props file Install Directory**
Visual Studio 2019	C:\Program Files (x86)\Microsoft Visual Studio\2019\Professional\MSBuild\Microsoft\VC\v160\Build
Visual Studio 2022	C:\Program Files\Microsoft Visual Studio\2022\Professional\MSBuild\Microsoft\VC\v170\BuildCust

You can reference this `CUDA 13.5.props` file when building your own CUDA applications.

5.3. Build Customizations for New Projects

When creating a new CUDA application, the Visual Studio project file must be configured to include CUDA build customizations. To accomplish this, click File-> New | Project... NVIDIA-> CUDA->, then select a template for your CUDA Toolkit version. For example, selecting the "CUDA 13.5 Runtime" template will configure your project for use with the CUDA 13.5 Toolkit. The new project is technically a C++ project (.vcxproj) that is preconfigured to use NVIDIA's Build Customizations. All standard capabilities of Visual Studio C++ projects will be available.

To specify a custom CUDA Toolkit location, under **CUDA C/C++**, select **Common**, and set the **CUDA Toolkit Custom Dir** field as desired. Note that the selected toolkit must match the version of the Build Customizations.

Note: A supported version of MSVC must be installed to use this feature.

5.4. Build Customizations for Existing Projects

When adding CUDA acceleration to existing applications, the relevant Visual Studio project files must be updated to include CUDA build customizations. This can be done using one of the following two methods:

1. Open the Visual Studio project, right click on the project name, and select **Build Dependencies > Build Customizations...**, then select the CUDA Toolkit version you would like to target.
2. Alternatively, you can configure your project always to build with the most recently installed version of the CUDA Toolkit. First add a CUDA build customization to your project as above. Then, right click on the project name and select **Properties**. Under **CUDA C/C++**, select **Common**, and set the **CUDA Toolkit Custom Dir** field to `$(CUDA_PATH)` . Note that the `$(CUDA_PATH)` environment variable is set by the installer.

While Option 2 will allow your project to automatically use any new CUDA Toolkit version you may install in the future, selecting the toolkit version explicitly as in Option 1 is often better in practice, because if there are new CUDA configuration options added to the build customization rules accompanying the newer toolkit, you would not see those new options using Option 2.

If you use the `$(CUDA_PATH)` environment variable to target a version of the CUDA Toolkit for building, and you perform an installation or uninstallation of any version of the CUDA Toolkit, you should validate that the `$(CUDA_PATH)` environment variable points to the correct installation directory of the CUDA Toolkit for your purposes. You can access the value of the `$(CUDA_PATH)` environment variable via the following steps:

1. Open a run window from the Start Menu.
2. Run:

```
control sysdm.cpl
```

3. Select the **Advanced** tab at the top of the window.
4. Click **Environment Variables** at the bottom of the window.

Files which contain CUDA code must be marked as a CUDA C/C++ file. This can be done when adding the file by right clicking the project you wish to add the file to, selecting **Add New Item**, selecting **NVIDIA CUDA 13.5\CodeCUDA C/C++ File**, and then selecting the file you wish to add.

For advanced users, if you wish to try building your project against a newer CUDA Toolkit without making changes to any of your project files, go to the Visual Studio command prompt, change the current directory to the location of your project, and execute a command such as the following:

```
msbuild <projectname.extension> /t:Rebuild /p:CudaToolkitDir="drive:/path/to/new/
↪toolkit/"
```

Chapter 6. Additional Considerations

Now that you have CUDA-capable hardware and the NVIDIA CUDA Toolkit installed, you can examine and enjoy the numerous included programs. To begin using CUDA to accelerate the performance of your own applications, consult the CUDA C Programming Guide, located in the CUDA Toolkit documentation directory.

A number of helpful development tools are included in the CUDA Toolkit or are available for download from the NVIDIA Developer Zone to assist you as you develop your CUDA programs, such as NVIDIA® Nsight™ Visual Studio Edition, Nsight Systems, Nsight Compute and Compute Sanitizer.

For technical support on programming questions, consult and participate in the developer forums by clicking here.

Chapter 7. Notices

7.1. Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or

services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

7.2. OpenCL

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

7.3. Trademarks

NVIDIA and the NVIDIA logo are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

©2009-2026, NVIDIA Corporation & affiliates. All rights reserved