




NVIDIA CUDA C INSTALLATION AND VERIFICATION ON MICROSOFT XP, MICROSOFT VISTA, AND WINDOWS 7 SYSTEMS

DU-80003-001_v01 | April 2010

Getting Started



DOCUMENT CHANGE HISTORY

DU-80003-001_v01

Version	Date	Authors	Description of Change
01	April 20, 2-010	CW, TS	Release

TABLE OF CONTENTS

Introduction	4
CUDA-Supercomputing on Desktop Systems	4
System Requirements	5
About This Document	5
Installing CUDA Development Tools	6
Verify You Have a CUDA-Enabled System	6
CUDA Software	7
Identify Your Driver	7
Install CUDA Software	8
Verify the Installation	9
Compiling the Examples	9
Compiling CUDA Programs	11
Compiling Sample Projects	11
Sample Projects	12
Additional Considerations	13

INTRODUCTION

CUDA-SUPERCOMPUTING ON DESKTOP SYSTEMS

NVIDIA® CUDA™ is a general purpose parallel computing architecture introduced by NVIDIA. It includes the CUDA Instruction Set Architecture (ISA) and the parallel compute engine in the GPU. To program to the CUDA architecture, developers can use C, one of the most widely used high-level programming languages, which can then be run at great performance on a CUDA-enabled processor.

The CUDA architecture and its associated software were 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 and C for CUDA, 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 both the CPU and GPU without contention for memory resources.

CUDA-enabled GPUs have hundreds of cores that can collectively run thousands of computing threads. Each core has shared resources, including registers and 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.

SYSTEM REQUIREMENTS

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

- ▶ CUDA-enabled GPU
- ▶ Device driver
- ▶ CUDA software (available at no cost from <http://www.nvidia.com/cuda>)
- ▶ Microsoft Visual Studio 2005 or 2008, or the corresponding versions of Microsoft Visual C++ Express

ABOUT THIS DOCUMENT

This document is intended for readers familiar with Microsoft Windows XP, Microsoft Windows Vista, or Microsoft Windows 7 operating systems and the Microsoft Visual Studio environment. You do not need previous experience with CUDA or experience with parallel computation.

INSTALLING CUDA DEVELOPMENT TOOLS

The installation of CUDA development tools on a system running the appropriate version of Linux consists of four simple steps:

- ▶ Verify the system has a CUDA-enabled GPU
- ▶ Download the CUDA software
- ▶ Install the driver for Windows XP, Windows Vista or Windows 7 (if necessary)
- ▶ Install the CUDA software

Test your installation by compiling and running one of the sample programs in the CUDA software to validate that the hardware and software are running correctly and communicating with each other.

VERIFY YOU HAVE A CUDA-ENABLED SYSTEM

Many NVIDIA products today contain CUDA-enabled GPUs. These include:

- ▶ NVIDIA GeForce® 8, 9, 200, and 400 series GPUs
- ▶ NVIDIA Tesla™ computing solutions
- ▶ Many of the NVIDIA Quadro® products

An up-to-date list of CUDA-enabled GPUs can be found on the NVIDIA CUDA Web site at http://www.nvidia.com/object/cuda_learn_products.html.

The Release Notes for the CUDA Toolkit also contain a list of supported products.

To verify which video adapter your Windows system uses, open the Control Panel (Start → Control Panel) and double click on System. In the System Properties window that opens, click the Hardware tab, then Device Manager. Expand the Display adapters entry. There you will find the vendor name and model of your graphics card.

CUDA SOFTWARE

The CUDA driver is integrated in the NVIDIA graphics driver and is available from the main CUDA download site at http://www.nvidia.com/object/cuda_get.html.

Choose the platform you are using and download the following:

- ▶ Driver
- ▶ The CUDA Toolkit
The CUDA Toolkit contains the tools needed to compile and build a CUDA application in conjunction with Microsoft Visual Studio. It includes tools, libraries, header files, and other resources.
- ▶ The CUDA SDK
The CUDA SDK (software development kit) includes sample projects that have all the necessary project configuration and build files to perform one-click builds using Microsoft Visual Studio.

Both software packages are available for 32-bit Windows (called x86 on the download site) and 64-bit Windows (called x86-64 on the download site). Download instructions appear in an earlier section of this chapter.

Before installing these packages, you should read the *Release Notes* bundled with each, as these notes provide details on installation and software functionality.

Identify Your Driver

The CUDA driver is integrated in the NVIDIA graphics driver. To use the CUDA Toolkit, you must have at least the version of the NVIDIA ForceWare® graphics driver specified in the *CUDA Toolkit Release Notes*. In most cases, if you are running a recent NVIDIA graphics adapter that has support for CUDA, you already have installed the CUDA driver.

To identify the version of your NVIDIA driver, open the NVIDIA Control Panel by right clicking on the desktop and selecting **NVIDIA Control Panel**. Click the **System Information** button in the lower left corner of the main panel to display a dialog box that specifies the version of the driver installed on your system.



Note: New versions of the CUDA Software can require updates of the driver, so always verify that you are running the right release of the driver for the version of CUDA you are using.

Install CUDA Software

Use the following procedure to download and install the CUDA software:

1. Download the NVIDIA CUDA software.

Go to http://www.nvidia.com/object/cuda_get.html and save the installer to your desktop.

2. Uninstall any previous versions of the NVIDIA CUDA Toolkit and SDK.

You can uninstall the NVIDIA CUDA Toolkit through the Start menu:

Start → All Programs → NVIDIA Corporation → CUDA Toolkit → Uninstall CUDA.

Uninstalling the CUDA SDK uses the same sequence.

3. Install CUDA Toolkit.

Install the CUDA Toolkit by executing the Toolkit installer package and following the on-screen prompts. The CUDA Toolkit installation default to **C:\CUDA**.

4. Define the environment variables.

- **CUDA_BIN_PATH** (defaults to either **C:\CUDA\bin** or **C:\CUDA\bin64**) contains the compiler executables and runtime libraries.
- **CUDA_INC_PATH** (defaults to **C:\CUDA\include**) contains the include files needed to compile CUDA programs.
- **CUDA_LIB_PATH** (defaults to either **C:\CUDA\lib** or **C:\CUDA\lib64**) contains the libraries needed for linking CUDA codes.

In addition to these directories, the CUDA Toolkit installation also includes a documentation directory (**C:\CUDA\doc**) containing the *CUDA Programming Guide*, *Compiler Guide*, and guides for the CUDA implementation of the **BLAS** and **FFT** libraries.

5. Install CUDA SDK.

Install the CUDA SDK by executing the installer package and following the on-screen prompts. The CUDA SDK is installed in **C:\Documents and Settings\All Users\Application Data\NVIDIA Corporation\NVIDIA GPU Computing SDK** and contains source code for many example problems and templates for Microsoft Visual Studio.

VERIFY THE INSTALLATION

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

Compiling the 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**

The CUDA SDK includes sample programs in both *source* and *compiled* form. To verify a correct configuration of the hardware and software, it is highly recommended that you run the **bandwidthTest** program located in **C:\Documents and Settings\All Users\Application Data\NVIDIA Corporation\NVIDIA GPU Computing SDK\C\bin\win32\Release**, presuming that you used the default installation directory structure. (On 64-bit versions of Windows, the directory name ends with **\win64\Release**.) If CUDA is installed and configured correctly, the output should look similar to Figure 1.

```

C:\Documents and Settings\All Users\Application ...
Running on.....
device 0:Quadro NUS 140M
Quick Mode
Host to Device Bandwidth for Pageable memory
.
Transfer Size (Bytes)  Bandwidth(MB/s)
33554432                1739.7

Quick Mode
Device to Host Bandwidth for Pageable memory
.
Transfer Size (Bytes)  Bandwidth(MB/s)
33554432                472.7

Quick Mode
Device to Device Bandwidth
.
Transfer Size (Bytes)  Bandwidth(MB/s)
33554432                4227.9

8888 Test PASSED

Press ENTER to exit...

```

Figure 1. Valid Results from Sample CUDA bandwidthTest Program

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-enabled NVIDIA GPU on your system and make sure it is properly installed.

To see a graphical representation of what CUDA can do, run the sample Particles executable in:

- ▶ For XP
`c:\Documents and Settings\All Users\Application Data\
NVIDIA GPU Computing SDK\C\bin\win32\Release` (or
...\`win64\Release` on 64-bit Windows)
- ▶ For Vista and Win7
`C:\ProgramData\NVIDIA Corporation\NVIDIA GPU Computing
SDK\C\bin\win32\Release` (or ...\`win64\Release` on 64-bit Windows)

COMPILING CUDA PROGRAMS

The project files in the CUDA SDK have been designed to provide simple, one-click builds of the programs that include all source code. To build the 32-bit or 64-bit Windows projects (for release, debug, or emulated release and debug-called *emurelease* and *emudebug*, respectively), use the provided ***.sln** solution files for Microsoft Visual Studio 2005 and ***_vc90.sln** for Microsoft Visual Studio 2008. (Likewise for the corresponding versions of Microsoft Visual C++ Express Edition.) You can use either the solution files located in each of the examples directories in **NVIDIA GPU Computing SDK\C\src** or the global solution files **Release.sln** or **Release_vc90.sln** located in **NVIDIA GPU Computing SDK\C\src**.

COMPILING SAMPLE PROJECTS

The **bandwidthTest** project is a good sample project to build and run. It is located in the **...NVIDIA Corporation\NVIDIA GPU Computing SDK\C\src\bandwidthTest** directory.

The output is placed in **NVIDIA GPU Computing SDK\C\bin\win32\Debug**. (As mentioned previously, the **\win32** segment of this address will be **\win64** on 64-bit versions of Windows.) This location presumes that you used the default installation directory structure.

Build the program using the appropriate solution file and run the executable. If all works correctly, the output should be similar to Figure 1.

SAMPLE PROJECTS

The sample projects come in four configurations: debug and release (where release contains no debugging information), and emulated versions of both. The emulated versions are for developing and running CUDA software in the absence of a CUDA GPU.

A few of the example projects require some additional setup. The **simpleD3D9** example requires the system to have a Direct3D SDK installed and the Visual C++ directory paths (located in **Tools**→**Options...**) properly configured. Consult the Direct3D documentation for additional details.

Most samples link to a utility library called **cutil** whose source code is in **NVIDIA GPU Computing SDK\C\common\src**. The release and emurelease versions of these samples link to **cutil32.lib** (or **cutil64.lib**) and dynamically load **cutil32.dll** (or **cutil64.dll**). The debug and emudebug versions of these samples link to **cutil32D.lib** and dynamically load **cutil32D** (or their 64-bit equivalents on 64-bit versions of Windows).

To build the Win32 release and/or debug configurations of the **cutil library**, use the solution files located in **NVIDIA GPU Computing SDK\C\common**. The output of the compilation process should be placed in **NVIDIA GPU Computing SDK\C\common\lib**:

cutil32.lib and **cutil32D.lib** (or **cutil64.lib** and **cutil64D.lib**) are the release and debug import libraries.

cutil32.dll and **cutil32D.dll** (or **cutil64.dll** and **cutil64D.dll**) are the release and debug dynamic-link libraries, which also are copied to **NVIDIA GPU Computing SDK\C\bin\win32**[**release**|**emurelease**] and **NVIDIA GPU Computing SDK\C\bin\win32**[**debug**|**emudebug**] respectively. (Substitute **\win64** for **\win32** on 64-bit Windows.)

ADDITIONAL CONSIDERATIONS

Now that you have CUDA-enabled hardware and the software installed, you can examine and enjoy the numerous included programs. To begin using CUDA to accelerate the performance of your own applications, consult the *NVIDIA CUDA Programming Guide*, located in **c:\CUDA\doc**.

For technical support on programming questions, consult and participate in the bulletin board and mailing list at <http://forums.nvidia.com/index.php?showforum=71>.

Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

HDMI

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

ROVI Compliance Statement

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

This device is protected by U.S. patent numbers 6,516,132; 5,583,936; 6,836,549; 7,050,698; and 7,492,896 and other intellectual property rights. The use of ROVI Corporation's copy protection technology in the device must be authorized by ROVI Corporation and is intended for home and other limited pay-per-view uses only, unless otherwise authorized in writing by ROVI Corporation. Reverse engineering or disassembly is prohibited.

OpenCL

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

Trademarks

NVIDIA, the NVIDIA logo, GeForce, Tesla, Quadro, and ForceWare 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

© 2010 NVIDIA Corporation. All rights reserved.