




NVIDIA CUDA C INSTALLATION AND VERIFICATION ON MAC OS

DU-80002-001_v01 | April 2010

Getting Started



DOCUMENT CHANGE HISTORY

DU-80002-001_v01

Version	Date	Authors	Description of Change
01	April 20, 2010	CW, TS	Release

TABLE OF CONTENTS

Introduction	5
CUDA-Supercomputing on Desktop Systems	5
System Requirements.....	6
About This Document	6
Installing CUDA Development Tools	7
Verify You Have a CUDA-Enabled System	7
Verify the Correct Version of Mac OS X	8
Verify That gcc Is Installed.....	8
Download the NVIDIA Driver and CUDA Software.....	9
Install the CUDA Driver and Software	10
Verify the Installation	11
Compiling the Examples	11
Running the Binaries	11
Additional Considerations	14

LIST OF FIGURES

Figure 1. About This Mac Dialog Box	8
Figure 2. Valid Results from Sample CUDA deviceQuery Program	12
Figure 3. Valid Results from Sample CUDA bandwidthTest Program	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
- ▶ Mac OS X v. 10.5.6 or later
- ▶ The gcc compiler and toolchain installed using Xcode
- ▶ CUDA software (available at no cost from <http://www.nvidia.com/cuda>)

ABOUT THIS DOCUMENT

This document is intended for readers familiar with the Mac OS X environment and the compilation of C programs from the command line. 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 Mac OS X consists of three simple steps:

- ▶ Verify the system has a CUDA-enabled GPU, a supported version of Mac OS X, and that gcc has been installed via Xcode.
- ▶ Download the CUDA driver and software.
- ▶ Install the CUDA driver and 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 Mac OS X system uses, under the **Apple** menu select **About This Mac**, click the **More Info ...** button, and then select **Graphics/Displays** under the **Hardware** list.

VERIFY THE CORRECT VERSION OF MAC OS X

The CUDA Development Tools require an Intel-based Mac running Mac OS X v. 10.5.6 or later. To check which version you have, go to the Apple menu on the desktop and select **About This Mac**. You should see a dialog box similar to Figure 1.



To use the CUDA Development Tools, you need to check for version 10.5.6 or later.

Figure 1. About This Mac Dialog Box

Verify That gcc Is Installed

The **gcc** compiler and toolchain are installed using the installation of Xcode. The Xcode development environment is found on the **Xcode Developer Tools DVD** that ships with new Mac systems and with Leopard, if you buy the operating-system upgrade. When installing Xcode, the package that contains **gcc** and the necessary tools is called *Developer Tools Essentials*. You can verify that **gcc** is installed entering the command `/usr/bin/gcc --help` from a Terminal window.

DOWNLOAD THE NVIDIA DRIVER AND CUDA SOFTWARE

Once you have verified that you have a supported NVIDIA processor and a supported version of the Mac OS, you need to download the NVIDIA driver and CUDA software.

The NVIDIA driver and CUDA software are available at no cost from the main CUDA download site at http://www.nvidia.com/object/cuda_get.html .

Choose **Mac OS** from the **Operating System** menu, and download the following packages for the latest version of the Development Tools:

- ▶ CUDA Driver
- ▶ CUDA Toolkit
- ▶ CUDA SDK

Note that there are two CUDA driver packages; one for NVIDIA GeForce GPUs and one for NVIDIA Quadro GPUs.

The CUDA Toolkit contains the tools needed to compile and build a CUDA application in conjunction with the **nvcc** compilation driver. It includes tools, libraries, header files, and other resources.

The CUDA SDK includes sample projects that provide source code and other resources for constructing CUDA programs.

INSTALL THE CUDA DRIVER AND SOFTWARE

Use the following procedure to successfully install the CUDA driver and software. For information not listed here, see the documentation under `/usr/local/cuda/doc` in the download location.

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

Then, follow these few steps for a successful installation.

1. **Uninstall any previous versions of the CUDA Toolkit and CUDA SDK.**
Do this by deleting the files from the default location at `/usr/local/cuda`. Adjust accordingly if you placed the files elsewhere. (If you wish to keep the files so you can compile for different versions of CUDA software, then rename the existing directories and modify your **makefile** accordingly.)
2. **Install the CUDA Driver.**
Installing the CUDA driver package based on the GPU in the system (either NVIDIA *GeForce* GPU or NVIDIA *Quadro* GPU) by executing the installer and following the on-screen prompts.
3. **Install the CUDA Toolkit.**
Install the CUDA Toolkit by executing the Toolkit installer package and following the on-screen prompts. The CUDA Toolkit installation defaults to `/usr/local/cuda`.
4. **Define the environment variables.**
 - The **PATH** variable needs to include `/usr/local/cuda/bin`.
 - **DYLD_LIBRARY_PATH** needs to contain `/usr/local/cuda/lib`.

The typical way to place these values in your environment is with the following commands:

```
export PATH=/usr/local/cuda/bin:$PATH
export DYLD_LIBRARY_PATH=/usr/local/cuda/lib:$DYLD_LIBRARY_PATH
```

To make these settings permanent, place them in `~/ .bash_profile`.

5. **Install the CUDA SDK.**
The installation process places the files in `/Developer/GPU Computing`. Note that this differs from the default location of the previous version of the SDK, which was placed in `/Developer/CUDA`.

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 terminal window. The `nvcc` command runs the compiler driver that compiles CUDA programs. It calls the `gcc` compiler for C code and the NVIDIA PTX compiler for the CUDA code.

NVIDIA includes sample programs in source form in the CUDA SDK. You should compile them all by changing to `/Developer/GPU Computing/C` and type `make`. The resulting binaries will be installed under the home directory in `/Developer/GPU Computing/C/bin/darwin/release`.

Running the Binaries

The sample projects use libraries pointed to by `DYLIB_LIBRARY_PATH`, as described earlier, so make sure it points to the right directory.

The executables need to find `libcutil.a` in `/Developer/GPU Computing/C/lib` and the corresponding include file in `/Developer/GPU Computing/C/common/inc`. They also need to access graphics libraries in `/Developer/GPU Computing/C/common/lib/darwin`. If you have installed the CUDA software as explained in this document, these locations are defaults and the programs should run without difficulty.

Once these required files are in place, go to `/Developer/GPU Computing/C/bin/darwin/release` and run `deviceQuery`. If CUDA is installed and configured correctly, the output for `deviceQuery` should look similar to Figure 2.

```

Terminal — deviceQuery — 92x31
There is 1 device supporting CUDA
Device 0: "GeForce 8600M GT"
CUDA Driver Version:            2.30
CUDA Runtime Version:          2.30
CUDA Capability Major revision number:  1
CUDA Capability Minor revision number:  1
Total amount of global memory:         268238048 bytes
Number of multiprocessors:            4
Number of cores:                     32
Total amount of constant memory:       65536 bytes
Total amount of shared memory per block: 16384 bytes
Total number of registers available per block: 8192
Warp size:                             32
Maximum number of threads per block:   512
Maximum sizes of each dimension of a block: 512 x 512 x 64
Maximum sizes of each dimension of a grid: 65535 x 65535 x 1
Maximum memory pitch:                 262144 bytes
Texture alignment:                    256 bytes
Clock rate:                            8.75 GHz
Concurrent copy and execution:         Yes
Run time limit on kernels:             Yes
Integrated:                             No
Support host page-locked memory mapping: No
Compute mode:                           Default (multiple host threads can use this
device simultaneously)

Test PASSED

Press ENTER to exit...

```

Figure 2. Valid Results from Sample CUDA deviceQuery Program

Note that the parameters for your CUDA device will vary. The key lines are the first and second ones that confirm a device was found and what model it is. Also, the next-to-last line, as indicated, should show that the test passed.

Running the **bandwidthTest** program ensures that the system and the CUDA device are able to communicate correctly. Figure 3 shows a sample of the output.

```

Terminal — bandwidthTest — 84x22
Quick Mode
Host to Device Bandwidth for Pageable memory
.
Transfer Size (Bytes)  Bandwidth(MB/s)
33554432                1127.4

Quick Mode
Device to Host Bandwidth for Pageable memory
.
Transfer Size (Bytes)  Bandwidth(MB/s)
33554432                1186.7

Quick Mode
Device to Device Bandwidth
.
Transfer Size (Bytes)  Bandwidth(MB/s)
33554432                9367.4

Test PASSED

Press ENTER to exit...

```

Figure 3. Valid Results from Sample CUDA bandwidthTest Program

Note that the measurements for your CUDA-enabled device description will vary from system to system. The important point is that you obtain measurements, and that the second-to-last line (in Figure 3) confirms that all necessary tests passed.

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

If you run into difficulties with the link step (such as libraries not being found), consult the Release Notes found in the doc folder in the SDK directory.

To see a graphical representation of what CUDA can do, run the sample particles executable.

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 `/usr/local/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, and Quadro 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.