



NSIGHT ECLIPSE EDITION

DG-06450-001 _v5.0 | October 2012

Getting Started Guide



TABLE OF CONTENTS

Chapter 1. Introduction.....	1
1.1 About Nsight Eclipse Edition.....	1
Chapter 2. Using Nsight Eclipse Edition.....	2
2.1 Installing Nsight Eclipse Edition.....	2
2.1.1 Installing CUDA Toolkit.....	2
2.1.2 Mac OS X Additional Notes.....	2
2.2 Running Nsight Eclipse Edition.....	3
2.3 Creating a New Project.....	3
2.4 Importing CUDA Sample.....	4
2.5 Debugging CUDA Applications.....	5
2.6 Profiling CUDA applications.....	6
2.7 More Information.....	7
Appendix A. Host Platform Requirements.....	8
A.1 Linux.....	8
A.2 Mac OS X.....	8
A.3 GPU Requirements.....	8
Appendix B. Known Issues.....	10

LIST OF FIGURES

Figure 1	Nsight main window after creating a new project.....	4
Figure 2	Debugging CUDA application.....	6
Figure 3	Profiling CUDA Application.....	7
Figure 4	JRE install prompt on Mac OS X platform.....	11

Chapter 1.

INTRODUCTION

This guide introduces Nsight Eclipse Edition and provides instructions necessary to start using this tool. For a detailed description of Nsight features consult the integrated help available from inside Nsight.

1.1 About Nsight Eclipse Edition

Nsight Eclipse Edition is an integrated development environment (IDE) for developers building CUDA application on Linux and Mac OS X platforms. It is designed to help developers on all stages of the software development process. Some of the principal features are:

- ▶ Source editor with extended support for CUDA C and C++ syntax
- ▶ Projects and files management with version control management system integration. CVS and Git are supported out of the box with integrations for other systems available separately as IDE plug-ins.
- ▶ Configurable makefile-based NVCC build integration
- ▶ Graphical user interface for debugging heterogeneous applications
- ▶ Visual profiler with source code correlation for optimizing GPU code performance

Nsight Eclipse Edition is based on the popular Eclipse Platform and supports a wide range of the third-party extensions and plug-ins including:

- ▶ Version control management systems support
- ▶ Compiler integrations
- ▶ Language IDEs
- ▶ Application lifecycle management and collaboration solutions

For more information about Eclipse Platform, visit <http://eclipse.org>

Chapter 2.

USING NSIGHT ECLIPSE EDITION

2.1 Installing Nsight Eclipse Edition

Nsight Eclipse Edition is installed as a part of the CUDA Toolkit package.

2.1.1 Installing CUDA Toolkit

To install CUDA Toolkit:

1. Visit the NVIDIA CUDA Zone download page:
http://www.nvidia.com/object/cuda_get.html
2. Select appropriate operating system. Nsight Eclipse Edition is available in Mac OS X and Linux toolkit packages.
3. Download and install the CUDA Driver.
4. Download and install the CUDA Toolkit.
5. Follow instructions to configure CUDA Driver and Toolkit on your system.

2.1.2 Mac OS X Additional Notes

Nsight Eclipse Edition requires Java Runtime Environment (JRE) to be available on the local system. Compatible JRE is included in CUDA Toolkit package on Linux platform.

Mac OS X users will be offered to install JRE by downloading it from Software Update servers unless JRE is already installed.

Nsight Eclipse Edition relies on command-line tools like *make* and *ar* to compile CUDA applications. On Mac OS X 10.8 (*Mountain Lion*) these tools are no longer installed as a part of the Xcode package.



Upgrading to Mac OS X 10.8 from earlier Mac OS X versions will remove these command-line tools.

To install command-line tools on Mac OS X 10.8:

1. Install Apple Xcode from the Mac App Store
2. Run Xcode and go to the **Xcode > Preferences...** from the main menu.
3. Open the *Components* tab of the *Downloads* sheet
4. Install *Command Line Tools* package

2.2 Running Nsight Eclipse Edition

To run Nsight from the command line type

```
nsight
```



On Linux systems, Nsight is also available from Gnome, KDE and Unity desktop menus.

On the first run Nsight will ask to pick a workspace location. The workspace is a folder where Nsight will store its settings, local files history and caches. An empty folder should be selected to avoid overwriting existing files.

The main Nsight window will open after the workspace location is selected. The main window is divided into the following areas:

- ▶ *Editor* - displays source files that are opened for editing.
- ▶ *Project Explorer* - displays project files
- ▶ *Outline* - displays structure of the source file in the current editor.
- ▶ *Problems* - displays errors and warnings detected by static code analysis in IDE or by a compiler during the build.
- ▶ *Console* - displays make output during the build or output from the running application.

2.3 Creating a New Project

1. From the main menu, open the new project wizard - **File > New... > CUDA C/C++ Project**
2. Specify the project name and project files location.
3. Select **CUDA Runtime Project** to create a simple CUDA runtime application.
4. Specify the project parameters on the next wizard page.



By default Nsight will automatically detect and target CUDA hardware available locally. Nsight will default to SM 1.0 if no CUDA hardware is detected.

5. Complete the wizard.
The project will be shown in the **Project Explorer** view and source editor will be opened.
6. Build the project by clicking on the hammer button on the main toolbar.




Figure 1 Nsight main window after creating a new project

2.4 Importing CUDA Sample

CUDA samples must be installed locally. Visit http://www.nvidia.com/object/cuda_get.html to download latest CUDA samples package.

1. From the main menu, open the new project wizard - **File > New... > CUDA C/C++ Project**
2. Specify the project name and project files location.
3. Select **Import CUDA Sample** under **Executable** in the **Project type** tree.
4. On the next wizard page select project sample you want to import. Press **Next...**
5. Specify the project parameters on the next wizard page.
6. Complete the wizard.
The project will be shown in the **Project Explorer** view and source editor will be opened.
7. Build the project by clicking on the hammer button on the main toolbar.

2.5 Debugging CUDA Applications

Nsight must be running and at least one project must exist.



GPUs used to run X11 (on Linux) or Aqua (on Mac) cannot be used to debug CUDA applications in Nsight Eclipse Edition. Consult *cuda-gdb* documentation for details.

1. In the **Project Explorer** view, select project you want to debug. Make sure the project executable is compiled and no error markers are shown on the project.
2. On the main window toolbar press **Debug** button (green bug).
3. You will be offered to switch perspective when you run debugger for the first time. Click "Yes".
Perspective is a window layout preset specifically designed for a particular task.
4. Application will suspend in the *main* function. At this point there is no GPU code running.
5. Add a breakpoint in the device code. Resume the application.

Debugger will break when application reaches the breakpoint. You can now explore your CUDA device state, step through your GPU code or resume the application.




Figure 2 Debugging CUDA application

2.6 Profiling CUDA applications

Nsight must be running and at least one project must exist. Profiler cannot be used when debugging session is in progress.

Nsight Eclipse Edition profiling features are based on the NVIDIA Visual Profiler (*nvvp*) code. These two tools provide same features and have same user interface.

1. In the **Project Explorer** view, select project you want to profile. Make sure the project executable is compiled and no error markers are shown on the project.
2. On the main window toolbar press the **Profile** button.
3. Press **Yes** when Nsight prompts to switch to the **Profile** perspective.

Nsight will switch to the **Profile** perspective and will display application execution timeline.




Figure 3 Profiling CUDA Application

2.7 More Information

Some of the topics not covered in this guide are:

- ▶ Navigating debugged application state
- ▶ Enabling *cuda-memcheck* integration when debugging GPU applications
- ▶ Attaching to running applications
- ▶ Measuring GPU code efficiency
- ▶ Refactoring CUDA C/C++ source code
- ▶ Accessing CVS and Git repositories
- ▶ Installing Nsight Eclipse Edition plug-ins

More information about these and other topics is available in the Nsight built-in documentation. To access Nsight documentation select *Help->Help Contents* from the Nsight main menu.

More information about CUDA, CUDA Toolkit and other tools is available on CUDA web page at <http://developer.nvidia.com/cuda>

Appendix A.

HOST PLATFORM REQUIREMENTS

The general platform and GPU requirements for running Nsight Eclipse Edition are described in this section.

A.1 Linux

Nsight Eclipse Edition is supported on the following Linux distributions:

- ▶ Red Hat Enterprise Linux 5.5+ (64-bit only)
- ▶ Red Hat Enterprise Linux 6.x
- ▶ Ubuntu 10.04, and 11.10
- ▶ Fedora 16
- ▶ OpenSuse 12.1
- ▶ Suse Linux Enterprise Server 11.1 and 11 SP2

A.2 Mac OS X

Nsight Eclipse Edition is supported on the following Mac OS X versions:

- ▶ Mac OS X 10.7
- ▶ Mac OS X 10.8

A.3 GPU Requirements

A CUDA-capable GPU is not required for writing and compiling your CUDA application using Nsight Eclipse Edition. A CUDA-capable GPU is required for debugging and profiling CUDA applications.

Debugging is supported on all CUDA-capable GPUs with a compute capability of 1.1 or later. *Compute capability* is a device attribute that a CUDA application can query about; for more information, see the latest *NVIDIA CUDA Programming Guide* on the NVIDIA CUDA Zone Web site: <http://developer.nvidia.com/object/gpucomputing.html>

These GPUs have a compute capability of 1.0 and are *not supported* :

- ▶ GeForce 8800 GTS
- ▶ GeForce 8800 GTX
- ▶ GeForce 8800 Ultra
- ▶ Quadro Plex 1000 Model IV
- ▶ Quadro Plex 2100 Model S4
- ▶ Quadro FX 4600
- ▶ Quadro FX 5600
- ▶ Tesla C870
- ▶ Tesla D870
- ▶ Tesla S870

A GPU that is running X11 (on Linux) or Aqua (on Mac) cannot be used to debug a CUDA application and will be hidden from the application ran in the debugger. Such GPU can still be used for profiling GPU applications.

Appendix B.

KNOWN ISSUES

Executable must exist in order to start debug session for the first time

Nsight will not automatically perform build when starting debug session for a given project for the first time. Build must be invoked manually. Nsight will automatically rebuild executable when starting subsequent debug sessions.



To manually build the project, select it (or any file within the project) in a **Project Explorer** view and click hammer icon on the main window toolbar.

Source editors may show error markers on a valid code for the files in newly created projects.

These markers will be cleared after Nsight indexes included header files.

Conditional breakpoints in the device code will not be hit if set before the CUDA kernel is running.

Nsight will show a warning saying "Breakpoint attribute problem: installation failed". To reinstall the breakpoint:

1. Break in the same CUDA kernel on any line prior to the conditional breakpoint line.
2. Disable and reenale the conditional breakpoint
3. Ensure that breakpoint no longer has warning icon and that checkmark was added to the breakpoint icon meaning the breakpoint was properly set on the target.
4. Resume the application.

Mac OS X users may be prompted to install Java Runtime Environment (JRE) when running Nsight Eclipse Edition for the first time.

Nsight Eclipse Edition requires functioning Java Runtime Environment to be present on the local system to run.




Figure 4 JRE install prompt on Mac OS X platform

Security updates for Java Runtime Environment will be provided by Apple using *Software Update* application.

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.

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

© 2007-2012 NVIDIA Corporation. All rights reserved.