FUNDAMENTALS OF ACCELERATED COMPUTING WITH CUDA C/C++

NVIDIA. | DEEP LEARNING INSTITUTE

This workshop teaches the fundamental tools and techniques for accelerating C/C++ applications to run on massively parallel GPUs with CUDA®. You'll learn how to write code, configure code parallelization with CUDA, optimize memory migration between the CPU and GPU accelerator, and implement the workflow that you've learned on a new task—accelerating a fully functional, but CPU-only, particle simulator for observable massive performance gains. At the end of the workshop, you'll have access to additional resources to create new GPU-accelerated applications on your own.

Learning Objectives

At the conclusion of the workshop, you'll have an understanding of the fundamental tools and techniques for GPU-accelerating C/C++ applications with CUDA and be able to:

> Write code to be executed by a GPU accelerator

> Expose and express data and instruction-level parallelism in C/C++ applications using CUDA

> Utilize CUDA-managed memory and optimize memory migration using asynchronous prefetching

> Leverage command line and visual profilers to guide your work

> Utilize concurrent streams for instruction-level parallelism

> Write GPU-accelerated CUDA C/C++ applications, or refactor existing CPU-only applications, using a profile-driven approach.

Workshop Information and Prerequisites:

Duration:	8 hours
Price:	**Contact us for pricing**
Prerequisites:	Basic C/C++ competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. No previous knowledge of CUDA programming is assumed.
Tools, libraries, and frameworks:	NVIDIA® Nsight™, nsys
Assessment type:	Code-based
Certificate:	Upon successful completion of the assessment, participants will receive an NVIDIA DLI certificate to recognize their subject matter competency and support professional career growth.
Hardware/software requirements:	Desktop or laptop computer capable of running the latest version of Chrome or Firefox. Each participant will be provided with dedicated access to a fully configured, GPU-accelerated workstation in the cloud.
Language:	English, Japanese, Chinese

Sample Workshop Outline

Introduction (15 mins)	> Meet the instructor. > Create an account at **courses.nvidia.com/join**
Accelerating Applications with CUDA C/C++ (120 mins)	Learn the essential syntax and concepts to be able to write GPU-enabled C/C++ applications with CUDA: > Write, compile, and run GPU code. > Control parallel thread hierarchy. > Allocate and free memory for the GPU.
Break (60 mins)	
Managing Accelerated Application Memory with CUDA C/C++ (120 mins)	Learn the command line profiler and CUDA managed memory, focusing on observation-driven application improvements and a deep understanding of managed memory behavior: > Profile CUDA code with the command line profiler. > Go deep on unified memory. > Optimize unified memory management.
Break (15 mins)	
Asynchronous Streaming and Visual Profiling for Accelerated Applications with CUDA C/C++ (50 mins)	Identify opportunities for improved memory management and instruction-level parallelism: > Profile CUDA code with NVIDIA Nsight Systems. > Use concurrent CUDA streams.
Final Review (15 mins)	> Review key learnings and answer questions. > Complete the assessment and earn a certificate. > Complete the workshop survey.

Why Choose NVIDIA Deep Learning Institute for Hands-On Training?

> Access workshops from anywhere with just your desktop/laptop and an internet connection. Each participant will have access to a fully configured, GPU-accelerated workstation in the cloud.

> Obtain hands-on experience with the most widely used, industry-standard software, tools, and frameworks.

> Learn to build deep learning and accelerated computing applications for industries, such as healthcare, robotics, manufacturing, accelerated computing, and more.

> Gain real-world experience through content designed in collaboration with industry leaders, such as the Children's Hospital of Los Angeles, Mayo Clinic, and PwC.

> Earn an NVIDIA Deep Learning Institute certificate to demonstrate your subject matter competency and support your career growth.

For the latest DLI workshops and trainings, visit **www.nvidia.com/dli**

For questions, contact us at **nvdli@nvidia.com**