

This workshop covers the fundamental tools and techniques needed for accelerating C/C++ or Fortran applications to run on massively parallel GPUs with OpenACC. You will learn how to write code, configure code parallelization with OpenACC, optimize memory movements between the CPU and GPU accelerator, and implement the workflow that you have learned on a new task—accelerating a fully functional, but CPU-only, Laplace Heat Equation code for massive performance gains. At the end of the workshop, you will have access to additional resources for creating new GPU-accelerated applications on your own.

Learning Objectives

By participating in this workshop, you'll have an understanding of the fundamental tools and techniques for GPU-accelerating C/C++ or Fortran applications with OpenACC and be able to:

- > Profile and optimize your CPU-only applications to identify hot spots for acceleration.
- > Use OpenACC directives to GPU-accelerate your codebase.
- > Optimize data movement between the CPU and GPU accelerator.

Workshop Information and Prerequisites:

Duration:	8 hours
Price:	Contact us for pricing
Prerequisites:	> Basic C/C++ or Fortran competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. > No previous knowledge of GPU programming is assumed.
Tools, libraries, and frameworks:	NVIDIA® Nsight™, OpenACC
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

Workshop Outline

Introduction (15 mins)	> Meet the instructor. > Create an account at courses.nvidia.com/join
Introduction to Parallel Programming (50 mins)	Learn about parallelism in a conceptual way, as well as how to express it with OpenACC. Topics that will be covered are as follows: > Introduction to parallelism > The goals of OpenACC > Basic parallelization of code using OpenACC
Break (10 mins)	
Profiling with OpenACC (50 mins)	Learn how to build and compile an OpenACC code, the importance of profiling, and how to use the NVIDIA Nsight Systems profiler. Topics that will be covered are as follows: > Compiling sequential and OpenACC code > The importance of code profiling > Profiling sequential and OpenACC multicore code > Technical introduction to the code used in introductory modules
Break (45 mins)	
Introduction to OpenACC Directives (50 mins)	Learn how to parallelize your code with OpenACC directives and understand the differences between parallel, kernel, and loop directives. Topics that will be covered are as follows: > The Parallel directive > The Kernels directive > The Loop directive
Break (10 mins)	
GPU Programming with OpenACC (50 mins)	Learn about the differences between GPUs and multicore CPUs, and manage memory with CUDA Unified Memory. Topics that will be covered are as follows: > Definition of a GPU > Basic OpenACC data management > CUDA Unified Memory > Profiling GPU applications
Break (15 mins)	
Data Management with OpenACC (50 mins)	Learn how to explicitly manage data movement with OpenACC data directives to reduce data transfers. Topics that will be covered are as follows: > OpenACC data directive/clauses > OpenACC structured data region > OpenACC unstructured data region > OpenACC update directive > Data management with C/C++ Structs/Classes
Break (10 mins)	

Loop Optimizations with OpenACC (50 mins)	Understand the various levels of parallelism on a GPU and learn about ways to extract more parallelism with OpenACC by optimizing loops in your code. Topics that will be covered are as follows: > Seq/Auto clause > Independent clause > Reduction clause > Collapse clause > Tile clause > Gang, Worker, Vector
Break (10 mins)	
Final Review (60 mins)	> Review key learnings and answer questions. > Complete the assessment and earn a certificate. > Complete the workshop survey.
Next Steps	Continue learning with these DLI trainings: > Fundamentals of Accelerated Computing with CUDA C/C++ > Scaling Workloads Across Multiple GPUs with CUDA C++ > High Performance Computing with Containers

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 server 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 DLI 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 nvdl@nvidia.com

© 2021 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, and Nsight are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. All other trademarks and copyrights are the property of their respective owners. AUG21

