

Accelerating CUDA C++ Applications with Multiple GPUs (ACCAMG)

ID ACCAMG Price on request Duration 1 day

Prerequisites

- Professional experience programming CUDA C/C++ applications, including the use of the nvcc compiler, kernel launches, grid-stride loops, host-to-device and device-to-host memory transfers, and CUDA error handling
- Familiarity with the Linux command line
- Experience using makefiles to compile C/C++ code

Suggested resources to satisfy prerequisites: Fundamentals of Accelerated Computing with CUDA C/C++, Ubuntu Command Line for Beginners (sections 1 through 5), Makefile Tutorial (through the Simple Examples section)

Course Objectives

- Use concurrent CUDA streams to overlap memory transfers with GPU computation
- Utilize all available GPUs on a single node to scale workloads across all available GPUs
- Combine the use of copy/compute overlap with multiple GPUs
- Rely on the NVIDIA Nsight™ Systems Visual Profiler timeline to observe improvement opportunities and the impact of the techniques covered in the workshop

Accelerating CUDA C++ Applications with Multiple GPUs (ACCAMG)

Training Centres worldwide



Fast Lane Institute for Knowledge Transfer (Switzerland) AG

Husacherstrasse 3
CH-8304 Wallisellen
Tel. +41 44 832 50 80

info@flane.ch, <https://www.flane.ch>