

Scaling CUDA C++ Applications to Multiple Nodes (SCCAMN)

ID SCCAMN Price on request Duration 1 day

Prerequisites

Intermediate experience writing CUDA C/C++ applications.

Suggested materials to satisfy the prerequisites:

- Fundamentals of Accelerated Computing with CUDA C/C++
- Accelerating CUDA C++ Applications with Multiple GPUs
- Accelerating CUDA C++ Applications with Concurrent Streams
- Scaling Workloads Across Multiple GPUs with CUDA C++

Course Objectives

By participating in this workshop, you'll:

- Learn several methods for writing multi-GPU CUDA C++ applications
- Use a variety of multi-GPU communication patterns and understand their tradeoffs
- Write portable, scalable CUDA code with the single-program multiple-data (SPMD) paradigm using CUDA-aware MPI and NVSHMEM
- Improve multi-GPU SPMD code with NVSHMEM's symmetric memory model and its ability to perform GPU-initiated data transfers
- Get practice with common multi-GPU coding paradigms like domain decomposition and halo exchanges

Scaling CUDA C++ Applications to Multiple Nodes (SCCAMN)

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>