Data Parallelism: How to Train Deep Learning Models on Multiple GPUs (DPHTDLM)

ID DPHTDLM **Price** on request **Duration** 1 days

Prerequisites

Experience with deep learning training using Python

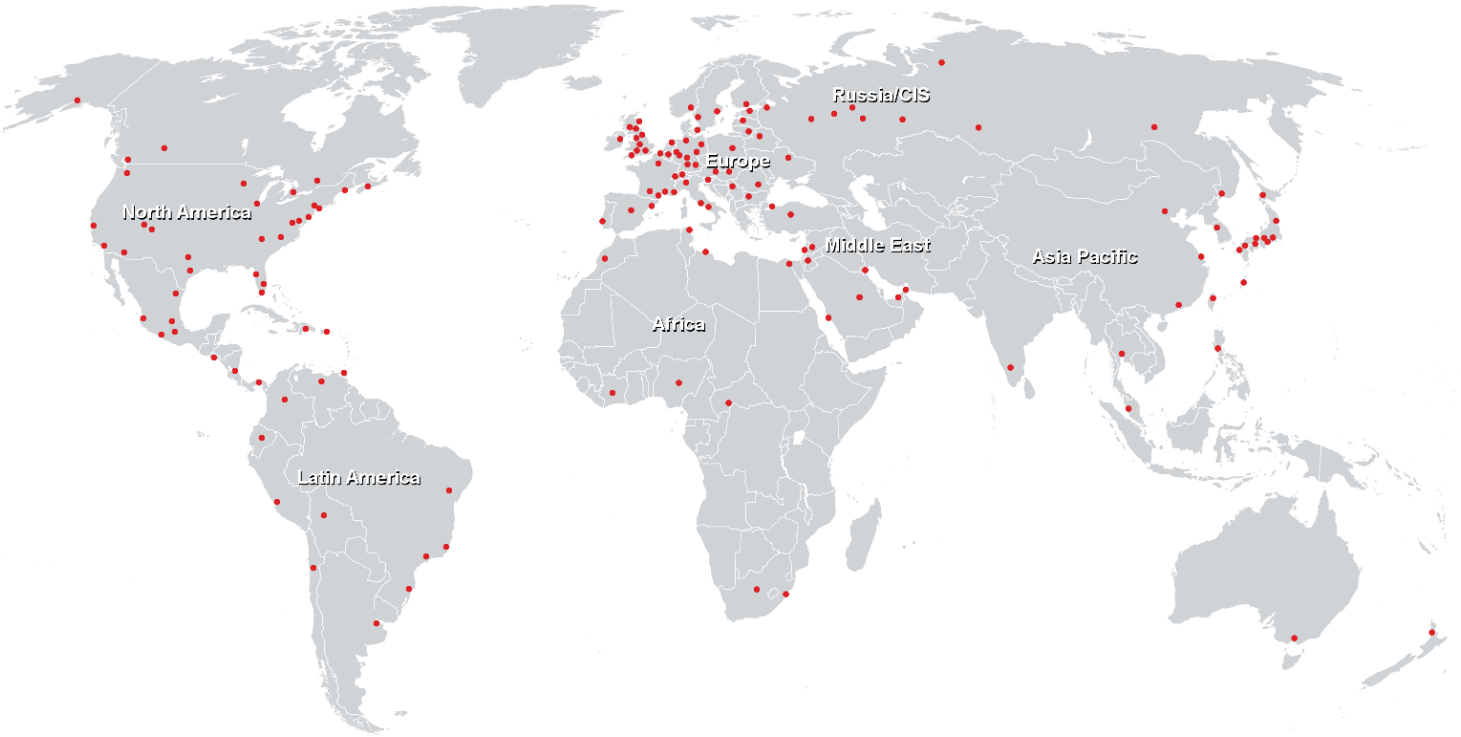
Course Objectives

By participating in this workshop, you'll:

- Understand how data parallel deep learning training is performed using multiple GPUs
- Achieve maximum throughput when training, for the best use of multiple GPUs
- Distribute training to multiple GPUs using Pytorch Distributed Data Parallel
- Understand and utilize algorithmic considerations specific to multi-GPU training performance and accuracy

Data Parallelism: How to Train Deep Learning Models on Multiple GPUs (DPHTDLM)

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