

NVIDIA - Data Parallelism: How to Train Deep Learning Models on Multiple GPUs

Download Whitepaper: Accelerate Your Modernization Efforts with a Cloud-Native Strategy
Get Your Free Copy Now

Course Number: NV-DP-GPU

Duration: 1 days

Overview

Course Description

Modern deep learning challenges leverage increasingly larger datasets and more complex models. As a result, significant computational power is required to train models effectively and efficiently. Learning to distribute data across multiple GPUs during deep learning model training makes possible an incredible wealth of new applications utilizing deep learning.

Additionally, the effective use of systems with multiple GPUs reduces training time, allowing for faster application development and much faster iteration cycles. Teams who are able to perform training using multiple GPUs will have an edge, building models trained on more data in shorter periods of time and with greater engineer productivity.

This workshop teaches you techniques for data-parallel deep learning training on multiple GPUs to shorten the training time required for data-intensive applications. Working with deep learning tools, frameworks, and workflows to perform neural network training, you'll learn how to decrease model training time by distributing data to multiple GPUs, while retaining the accuracy of training on a single GPU.

Skills Gained

- 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

Prerequisites

Experience with deep learning training using Python

Audience

Course Details

Course Outline

- Introduction
- Stochastic Gradient Descent and the Effects of Batch Size
- Training on Multiple GPUs with PyTorch Distributed Data Parallel (DDP)
- Maintaining Model Accuracy when Scaling to Multiple GPUs
- Workshop Assessment
- Final Review
- Next Steps