PROFESSIONAL COURSE OUTLINE

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

NVIDIA

Program aligned

Explores how to train deep learning models across multiple GPUs using data-parallel training approaches.

AI | Advanced | Distributed Deep Learning Frameworks

PROGRAM CODE	DELIVERY	DURATION	CERTIFICATION
Program aligned	Virtual, On-site, or Hybrid	8 hours	Available on request

AUDIENCE PROFILE

Who This Program Is For

Built for teams moving into faster and larger-scale deep learning training.

PROGRAM SUMMARY

What This Course Covers

Official NVIDIA DLI workshop on data parallelism for multi-GPU deep learning training.

TAILORED DELIVERY

Adapt the program around your team.

This outline can be adapted for virtual, on-site, or hybrid delivery, with emphasis adjusted for yourteam's platform priorities, role mix, and implementation goals.

Enterprise-ready delivery format

VNode ITeS can align labs, examples, delivery pace, and assessment checkpoints to the required audience profile while preserving the official program sequence where applicable.

COMPLETE MODULE SEQUENCE

Module Flow and Topic Coverage

The structure below presents the current delivery flow for this program, including the associated topics covered under each module.

1

MODULE 1

Scale training across GPUs

Learn data-parallel approaches for training deep learning models more efficiently on multiple GPUs.

- Multi-GPU training basics
- Data-parallel scaling patterns

VNode ITeS

Plan the next session

We can tune this outline around your delivery goals and team mix.

TALK TO US

info@vnodeites.com

+91 9419 11 4792

+91 9419 11 4792

Serving enterprise clients across India and global markets

OFFICE

Wework, DLF Forum Infinity Tower C, Gurugram, INDIA

Serving enterprise clients across India and global markets