

TRAINING OVERVIEW

The course dives deep into the DGX SuperPOD infrastructure and revolves around InfiniBand and Ethernet networking to get a holistic understanding of the SuperPOD storage, compute, and management fabrics.

The learner will be introduced to a variety of tools, built on NVIDIA's field best practices, and get the necessary knowledge and skills to fully administrate, troubleshoot, and maintain all DGX SuperPOD components and services, including login and authentication, monitoring, provisioning, workload management and container management.

TRAINING DELIVERY METHOD

Instructor-led remote or onsite training sessions via NVIDIA Academy WebEx platform. Hands-on lab exercises focused on the DGX System and network infrastructure.

TARGET AUDIENCE

This course is aimed at Administrators and data center professionals to be able to successfully administrate DGX SuperPOD configurations.

TRAINING DURATION – 4 sessions of 4 hours

TRAINING OUTLINE

Session 1

Unit 1: Introduction to GPU Computing

- GPU Computing Overview
- GPU Overview
- Networking Requirements for GPU-Computing
- NVIDIA Networking Technologies

Unit 2: DGX SuperPOD Overview

- Scalable supercomputing infrastructure
- Scalable Unit (SU)
- Rail-optimized compute InfiniBand fabric
- Storage InfiniBand fabric

Unit 3: DGX System Overview

- Compute Building Block
- Hardware Architecture
- Out-of-band management

Unit 4: DGX OS & Software & First Boot

- OS, Drivers, Software Components
- Initial System Configuration
- DGX-OS Update

Unit 5: DGX System Storage

- System Memory, NVMe Drives, Storage Fabric
- Fast Access to Storage with GPUDirect Storage

Unit 6: System Debugging with NVIDIA NVSM

- NVSM Overview
- NVSM Architecture and APIs
- NVSM CLI – Commands and Use Cases
- DGX System Health Check with NVSM

Unit 7: GPU-related Debugging with NVIDIA DCGM

- DCGM Overview
- DCGM CLI Usage and Examples

Unit 8: NVIDIA GPU Containers

- Container Overview
- Container Cache
- NVIDIA NGC Repository Registry
- Docker and Container Management

Unit 9: Running a Stress Test and Performance Validation

- Performance Testing Overview
- NVSM Stress-Test
- GPU Bandwidth Test
- Running a Job with Jupyter Notebook

Unit 10: Multi-Instance GPU

- Multi-Instance GPU Overview
- MIG Setup and Usage

Session 2: Networking and InfiniBand

Unit 1: Introduction to InfiniBand

- InfiniBand Overview
- InfiniBand Key Features
- InfiniBand Fabric Components

Unit 2: InfiniBand Architecture and Management

- InfiniBand Network Stack
- InfiniBand Architecture
- Subnet Manage
- Fabric Addressing and Segmentation
- OFED and OFED Utilities

Unit 3: DGX SuperPOD Fabric Topology

- Overview of Network Topology

Unit 4: Ibdagnet

- Ibdagnet Overview
- Fabric Debug
- Link Speed Verification

Unit 5: Unified Fabric Monitor

- UFM Overview
- Key Features
- Architecture
- Operational Dashboard
- Fabric Health & Logging

Session 3 and 4 : Base Command (Bright Cluster)

Unit 1: Base Command Software Platform Overview

- DGX SW Stack Overview
- DGX OS
- Magnum IO
- Base Command Manager
- NVIDIA AI Enterprise

Unit 2: Base Command Manager

- Overview
- Base Command Manager Components
- Cluster Management Tools
- Node Provisioning
- Software Images
- Node Categories
- User Management
- Workload Management
- Health Monitoring
- Jupyter Notebooks
- GPU Configuration

Unit 3: Slurm Overview

- Slurm Overview
- Slurm Configuration

Unit 3a: Kubernetes Overview

- Kubernetes Overview
- GPU & Network Operator
- MLOps Tools