

TRAINING OVERVIEW

The course dives deep into the DGX A100 BasePOD infrastructure and revolves around InfiniBand and Ethernet networking to get a holistic understanding of the BasePOD 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 BasePOD 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 A100 System and network infrastructure.

TARGET AUDIENCE

This course is aimed at IT and data center professionals. It will enable administrators to successfully administrate DGX BasePOD configurations.

TRAINING SKU: 789-PODIV4

TRAINING DURATION – Remote: 4 sessions of 3.5 hours / 3-days onsite

TRAINING OUTLINE

DGX System Overview

Unit 1: Introduction to GPU Computing

- GPU Computing Overview
- NVIDIA A100 GPU Overview
- Networking Requirements for GPU-Computing
- NVIDIA Networking Technologies

Unit 2: DGX BasePOD Overview

- BasePOD Architecture
- Management Overview
- Compute InfiniBand fabric
- Storage InfiniBand fabric

Unit 3: DGX A100 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 A100 System Storage

- System Memory, NVMe Drives, Storage Fabric
- Storage types and considerations

Unit 6: NVIDIA System Management (NVSM)

- NVSM Overview
- NVSM Architecture and APIs
- NVSM CLI – Commands and Use Cases
- DGX A100 System Health Check with NVSM

Unit 7: NVIDIA Data Center GPU Management (DCGM)

- DCGM Overview
- DCGM CLI Usage and Examples

Unit 8: NVIDIA GPU Containers

- Container Overview
- 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

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 BasePOD Fabric Topology

- Overview of Network Topology

Unit 4: Ibdaignet

- Ibdiagnet Overview
- Fabric Debug
- Link Speed Verification

Unit 5: Unified Fabric Monitor

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

Base Command

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

- Overview
- Base Command Manager Components
- Cluster management tools
- Health Monitoring

Unit 3a: Kubernetes Overview

- Kubernetes Overview
- GPU & Network Operator
- MLOps Tools

Unit 3b: Slurm Overview

- Slurm Overview
- Slurm Configuration