



DGX SuperPOD Administration

OUTLINE

Training Overview

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

The learners 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 onsite or remote training sessions via NVIDIA Academy Teams 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

Remote: 4 sessions of 4 hours

Session 1

Introduction to GPU Computing

- GPU Computing Overview
- NVIDIA H100 GPU Overview
- Networking Requirements for GPU-Computing
- NVIDIA Networking Technologies

DGX SuperPOD Overview

- Scalable Supercomputing Infrastructure
- Scalable Unit (SU)
- Rail-optimized Compute InfiniBand Fabric
- Storage InfiniBand Fabric

DGX System Overview

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

DGX System Storage

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

System Debugging with NVIDIA NVSM

- NVSM Overview
- NVSM Architecture and APIs
- NVSM CLI – Commands and Use Cases
- DGX H100 System Health Check with NVSM

GPU-related Debugging with NVIDIA DCGM

- DCGM Overview
- DCGM CLI Usage and Examples

NVIDIA GPU Containers

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

Running a Stress Test and Performance Validation

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

Multi-Instance GPU

- Multi-Instance GPU Overview
- MIG Setup and Usage

Session 2: Networking and InfiniBand

Introduction to InfiniBand

- InfiniBand Overview and Key Features
- InfiniBand Fabric Components

InfiniBand Architecture and Management

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

DGX SuperPOD Fabric Topology

- Overview of Network Topology

Monitoring the Fabric with ibdiagnet Utility

- ibdiagnet Overview
- Fabric Debug
- Link Speed Verification

Monitoring the Fabric with Unified Fabric Monitor (UFM)

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

Session 3 & 4: NVIDIA Base Command

Base Command Software Platform Overview

- DGX Software Stack Overview
- DGX OS
- Magnum IO
- Base Command Manager

NVIDIA AI Enterprise Overview

- NVIDIA AI Enterprise Catalog
- AI Use Cases, Frameworks, and Pretrained Models
- AI Development
- Sharing Access to NVIDIA AI Enterprise Software
- Using the NGC Private Registry

Base Command Manager

- Base Command Manager Overview and Components
- Cluster Management Tools
- Node Provisioning
- Software Images
- Node Categories
- User Management
- Workload Management
- GPU Configuration

Slurm Overview

- Slurm Overview
- Slurm Configuration

Kubernetes Overview

- Kubernetes Overview
- GPU & Network Operator
- MLOps Tools

Monitoring with Base Command Manager

- Data Producers
- Measurables
- Consolidators
- Actions
- Triggers
- Health Status and Health Checks
- Job Monitoring
- Dashboards