DGX System Administration

OUTLINE

Training Overview

The DGX System Administration course provides an overview of the NVIDIA DGX System, tools for in-band and out-of-band management, NGC, the basics of running workloads, and specific management tools and CLI commands. In addition, this course includes content on multi-instance GPU, managing storage, performance validation, and other system management tools and concepts. The related instruction and guidance are based on NVIDIA's best practices and cover the critical knowledge and skills to administer and manage your DGX system.

Target Audience

This course is aimed at IT and data center professionals to be able to successfully administrate standalone NVIDIA DGX systems or small DGX clusters.

Training Duration

Remote: 4 sessions of 4 hours

Training Delivery Method

Instructor-led onsite or remote training sessions via the NVIDIA Academy Teams platform. Hands-on lab exercises focused on the DGX system and Base Command Manager.

Training Outline

Session 1: DGX System Overview

Introduction to GPU Computing

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

DGX System Overview

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

DGX Software Overview and First Boot

- First Boot
  - Baseboard Management Controller connection
  - Initial Settings
    - Configuring a User Account
    - Configuring Network Settings
  - DGX OS Update

DGX Firmware Updating

- Firmware Retrieval
- Firmware Update

DGX System Storage

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

NVIDIA System Management (NVSM)

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

NVIDIA Data Center GPU Management (DCGM)

- DCGM Overview
- DGCM CLI Usage and Examples

NVIDIA GPU Containers

- Container Overview
- 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

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 Overview

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

Kubernetes Overview

- GPU & Network Operator
- MLOps Tools

Slurm Overview

- Slurm Overview
- Slurm Configuration

Monitoring with Base Command Manager

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