

AI Infrastructure and Operations – Public Training

OUTLINE

Training Overview

In today's AI-driven era, the ability to deploy and manage AI clusters efficiently and effectively is critical for organizations seeking to leverage the power of artificial intelligence. This specialized course is designed to equip professionals with the skills and knowledge necessary to optimize efficiency, reliability, and scalability in deploying AI environments within data centers.

Covering various infrastructure and operation aspects like compute platforms, networking, and storage, the course focuses on key AI components such as GPUs, CPUs, and DPUs.

Participants will gain practical insights into provisioning and managing AI data centers, implementing AI workloads, and utilizing AI virtualization techniques. Through a blend of hands-on training and theoretical learning, participants develop the expertise needed to effectively deploy and manage AI infrastructure, ensuring optimal performance and continuous operation of AI data centers by course end.

Training Delivery Method

Instructor-led remote training sessions via NVIDIA Academy platform. Hands-on lab exercises focused on the data center infrastructure and operations.

Target Audience

The course is designed for administrators, DevOps professionals, and IT-related roles who want to gain the knowledge and skills necessary to deploy and operate AI data centers.

Training Duration

Remote | 11 sessions of 4 hours

Prerequisites

- Knowledge of networking concepts and principles, including Ethernet and InfiniBand technologies used in data centers and high-performance computing environments.
- Hands-on experience in Linux-like systems administration, such as managing users and permissions, installing software packages, configuring network settings, and troubleshooting common issues in a Linux environment.
- Basic understanding of server hardware components and their roles in a data center environment. This includes knowledge of CPUs, memory, storage devices, and networking interfaces commonly found in servers.
- Knowledge of storage concepts and principles, including different file systems and their characteristics, as well as the functioning and usage of storage protocols in data storage and retrieval.
- Basic understanding of virtualization technologies, including virtual machines (VMs) and containers. You should be familiar with VM creation, management, and the role of hypervisors in virtualized environments.
- Basic understanding of artificial intelligence (AI) concepts and terminology. This may include knowledge of topics such as machine learning, deep learning, neural networks, and common AI applications.
- Before attending the course, we recommend completing the [AI Infrastructure and Operation Fundamentals](#) self-paced course. This course will provide the foundations for this training.

Training Outline

AI in the Data Center Overview

- AI Overview
- Data Center Architecture for AI Workloads

Compute Platforms for AI

- AI Compute Platforms Overview
- Scaling AI Compute
- AI Compute Software
- Operating AI Compute Platforms
- Operating AI Compute Platforms
- NVIDIA GPU Containers
- Practice: Installing the NVIDIA GPU driver and using the nvidia-smi tool
- Practice: NVIDIA DGX System Health Diagnostics with NVSM
- Practice: Monitoring GPUs and NVLink with DCGM

Networking for AI

- Networking for AI Data Centers
- Building and Operating InfiniBand Fabrics
- Monitoring InfiniBand Fabrics with UFM
- Adapting Ethernet Networks to Run AI Workloads
- AI Data Centers Networks
- Practice: Managing and operating the InfiniBand Fabric
- Practice: Monitoring the InfiniBand Fabric with UFM

Storage for AI

- Storage Requirements for AI data Centers
- Storage Architecture for AI
- Practice: Mounting storage and testing performance

BlueField Networking Platform for AI

- BlueField Overview and Uses Cases
- BlueField Bring-up + hands-on
 - Installing DOCA
 - Firmware Upgrade
 - Management via RShim
 - BlueField Interfaces – Network Interfaces and OVS Bridges
- BlueField Management
 - Installing DOCA on BlueField
- Running DOCA Services on BlueField
- Practice: Bringing-up BlueField and running DOCA services

AI Data Center Management

- AI Data Center Management Overview
- AI Infrastructure Management and Provisioning with Base Command Manager
- Scheduling AI Workloads with SLURM
- AI Cluster Orchestration with Kubernetes
- Managing AI at the Edge with Fleet Command
 - Practice: Operating and managing an AI cluster with BCM

Virtualizing GPU Resources

- GPU Temporal Partitioning
- GPU Spatial Partitioning
- Practice: Virtualizing GPU resources with NVIDIA vGPU on VMware vSphere
- Practice: Virtualizing GPU resources using MIG

NVIDIA AI Software

- Using NGC Containers
- NVIDIA AI Enterprise Software Suite
- Practice: Running AI applications in NGC containers