

AI Operations – Public Training

OUTLINE

Training Overview

Running AI in the data center is an essential skill set for individuals involved in managing, deploying, and utilizing AI technologies.

This course covers essential aspects of operating AI data centers, including provisioning and management, running AI workloads, and implementing AI virtualization.

Participants will acquire expertise in operating essential AI building blocks such as GPUs, CPUs, and DPUs, empowering them to effectively operate an AI data center while ensuring optimal performance and uninterrupted operation.

Training Delivery Method

Instructor-led remote training sessions via NVIDIA Teams platform. Hands-on lab exercises focused on the AI data center operations.

Target Audience

The course is intended for administrators, DevOps professionals, and data scientists who aim to acquire the skills required to effectively run AI in the data center in the AI-driven era.

Training Duration

Remote | 4 sessions of 5 hours

Prerequisites

- Knowledge of core networking concepts and principles, including the TCP/IP model, Ethernet standards, basics of routing and switching, common network topologies, and IP addressing schemes.
- 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.
- 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
- NVIDIA AI Enterprise
- Operating AI Compute Platforms
- NVIDIA GPU Containers
- Monitoring and Troubleshooting NVLink with DCGM
- Troubleshooting Compute Systems Using NVSM
- Practice: Installing and Running AI Software
- Practice: GPU Driver Installation
- Practice: Monitoring and Troubleshooting NVLink with DCGM

Networking for AI

- AI Data Centers Networks

BlueField Networking Platform for AI

- BlueField Overview and Uses Cases
- BlueField Management
- Running DOCA Services on BlueField
- Practice: Running the BlueMan Service

Storage for AI

- Storage Requirements for AI data Centers
- Storage Configuration, Optimization and Testing
- Practice: Mounting Storage and Testing Performance

AI Data Center Management

- AI Data Center Management Overview
- AI Infrastructure Management and Provisioning with Base Command Manager
 - Overview, components, installation, cluster management tools and troubleshooting
- Scheduling AI Workloads with Slurm
- AI Cluster Orchestration with Kubernetes
- Run:ai Administration
- Practice: Operating and managing an AI cluster with BCM
- Practice: Scheduling AI Workloads with Slurm
- Practice: Orchestrating Containerized AI Applications with Kubernetes

NVIDIA AI Software

- Using NGC Containers
- NGC CLI Installation on the Host
- NVIDIA AI Enterprise Software Suite