

TRAINING OVERVIEW

In today's AI-driven era, the ability to deploy AI clusters efficiently and effectively is crucial for organizations.

NVIDIA Spectrum™ -X networking platform enhances network performance by 1.6X, accelerating the processing, analysis, and execution of AI workloads and, in turn, the development and deployment of AI solutions.

In this course, participants will gain essential skills to deploy and manage AI clusters using the NVIDIA Spectrum-X networking platform.

Through instructor-led sessions and hands-on labs in the NVIDIA Air environment, participants will learn to provision and manage AI clusters.

Key topics include advanced networking for AI, Spectrum-X deployment, and real-time network monitoring using NetQ and Cumulus Linux.

TRAINING DELIVERY METHOD

- Instructor-led remote training sessions via NVIDIA Teams platform.
- Hands-on lab exercises based on the NVIDIA Air environment.

TARGET AUDIENCE

The course is designed for network administrators, DevOps professionals, and IT-related roles who want to gain the knowledge and skills necessary to deploy and maintain Spectrum-X networking platform-based AI data centers.

TRAINING DURATION

Remote | 3 sessions of 4 hours each

PREREQUISITES

- Knowledge of networking concepts and principles, including technologies used in data centers and high-performance computing 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.

Equivalent knowledge to [AI for All: From Basics to GenAI Practice](#) course.

- Practical experience in configuring and managing Cumulus Linux based network environments.
Equivalent knowledge to “Cumulus Linux Professional” course.
- Familiarity with installing [DOCA OFED](#) on the host

LEARNING OBJECTIVES

By the end of the course, you should be able to:

- Explain the fundamentals of NVIDIA Spectrum-X Networking Platform, including its architecture, key components, and reference design for AI environments.
- Gain hands-on experience with NVIDIA Air environment for simulating and testing Spectrum-X deployments.
- Deploy the Spectrum-X platform, including IP addressing, QoS configurations, routing policies, and virtualized network setup for multi-tenancy.
- Apply advanced networking concepts such as RoCE , Adaptive Routing, and Congestion Control in the context of AI workloads.
- Monitor and troubleshoot Spectrum-X fabric using NVIDIA NetQ and Cumulus Linux CLI.

TRAINING OUTLINE

Day 1

Introduction to Spectrum-X Networking Platform

- Spectrum-X networking platform overview
- Architecture overview
- Reference architecture
- NVIDIA digital twins with Air environment
- Hands-on – Accessing the Air environment

Day 2

Spectrum-X Platform Deployment

- Deployment Guide
 - IP Addressing Overview

- QoS: RoCE, Adaptive Routing and Congestion Control
- Routing Policies
- Underlay Network
- Virtualized Network and Multitenancy
- Hands-on - Deploying the Spectrum-X Platform

Day 3:

Monitoring and Troubleshooting

- Spectrum-X Fabric Telemetry with NetQ
 - NetQ features
 - Installing and configuring the NetQ agent
 - Validation checks for network health
 - Fabric monitoring methods:
 - ASIC monitoring tools
 - OTLP (Open Telemetry)
 - DTS – DOCA Telemetry Service
- Hands-on - Managing fabric telemetry with NetQ
- Hands-on - Troubleshooting Spectrum-X platform deployment