Implementing AI Data Center Networks One-Day Workshop

COURSE OVERVIEW

This one-day, intermediate-level workshop provides students with knowledge that might be helpful when building and working with Juniper Apstra™ in an artificial intelligence data center (AI data center). This workshop will provide attendees with the background knowledge necessary to understand the usage of the backend graphic processing unit (GPU) network described in the Juniper Validated Design (JVD) titled AI Data Center Network with Juniper Apstra, NVIDIA GPUs, and WEKA Storage. Students will learn to train AI models using the PyTorch framework on:

- a single server with multiple GPUs (covering NVIDIA's NVSwitch); and
- multiple servers with each having multiple GPUs.

Students will gain familiarity with network interface cards (NICs) for AI (NVIDIA ConnectX-7 and Broadcom P2200G), Nvidia H100 GPUs, and a compute platform architecture (NVIDIA DGX H100). Students will be provided with an overview of the NVIDIA-focused JVD for the AI data center. In the case of the back-end GPU network, students will learn that using NVIDIA Collective Communication Library (NCCL), remote direct memory access (RDMA) over Converged Ethernet (RoCEv2), and a rail-optimized network design ensures an optimal communication path for the collective operations of NCCL. Students will learn how to use both data center quantized congestion notification (DCQCN) and dynamic load balancing (DLB) to ensure lossless data transfer over an Ethernet-based network. Students will learn how to use Apstra to deploy the back-end GPU network as well as orchestrate the training cluster using Slurm.

Through lectures only, students will gain knowledge in deploying and training AI models in a DC based on the JVD titled *AI Data Center Networks with Juniper Apstra, NVIDIA GPUs, and WEKA Storage*.

Note: Want more? A three-day version of this workshop is available, with expanded coverage of topics such as front-end and storage networks, automation with Terraform, and validated designs. View details

COURSE LEVEL

Intermediate

AUDIENCE

- Individuals who want an understanding of how to train a machine learning model in a data center that is optimized for AI model training.
- Individuals that will manage and operate a data center that is optimized for AI model training.

PREREQUISITES

- Strong background in network design and operations
- Understanding of a Clos IP fabric
- Knowledge of basic automation design and workflow
- Background in Linux, Python, and Junos-based class-of-service
- Completion of the *Data Center Automation using Juniper Apstra* (APSTRA) course or equivalent knowledge

RELATED JUNIPER PRODUCTS

Juniper Apstra, QFX Series, PTX Series

RECOMMENDED NEXT COURSE

Implementing Data Center Fabric with EVPN and VXLAN

OBJECTIVES

- Describe the basics of machine learning
- Describe the purpose of the frameworks for AI
- Describe a compute platform for the AI data center
- Describe how a model is trained on multiple compute nodes using multiple GPUs
- Describe validated designs

© 2025 Juniper Networks, Inc. Course content subject to change. See www.juniper.net/courses for the latest details.

Juniper Public

- Describe the back-end GPU network for the JVD
- Describe the Juniper Apstra's analytics features for the backend GPU network

COURSE CONTENTS

DAY 1

Module 1: What Is Machine Learning?

- Describe the various forms of machine learning
- Describe the reasoning behind building an AI data center
- Describe machine learning data types and operations
- Describe the process of training an AI model

Module 2: Machine Learning Stack

- Describe a compute platform built for artificial intelligence
- Describe the machine learning stack

Module 3: Machine Learning—Single GPU

- Describe the CUDA processing flow.

Module 4: AI Compute Platforms

- Describe the NVIDIA DGX H100 compute platforms

Module 5: Machine Learning—Multiple GPUs

- Describe the processing flow of CUDA and NCCL within a single compute node
- Describe how to use PyTorch to train a simple model using multiple GPUs

Module 6: Machine Learning—Multiple Nodes

- Describe the processing flow of CUDA and NCCL between multiple compute nodes
- Describe how to use Slurm to orchestrate parallel tasks
- Describe how to use PyTorch and Slurm to train a simple model using multiple nodes

Module 7: Reference Designs

- Describe the Juniper Validated Designs for the AI data center

Module 8: Juniper Validated Design—Compute Network

- Describe the back-end GPU network topology
- Describe how to design the back-end GPU network with Juniper Apstra
- Describe lossless Ethernet using DCQCN
- Describe dynamic load balancing

Module 9: Automation and Analytics

- Describe the Juniper Apstra's analytics features for the AI data center

© 2025 Juniper Networks, Inc. Course content subject to change. See www.juniper.net/courses for the latest details.

Juniper Public