

Implementing AI Data Center Networks Workshop

COURSE OVERVIEW

This advanced three-day workshop provides students with advanced-level knowledge that might be helpful when building and working with Juniper Apstra in an artificial intelligence data center (AI DC). This workshop will provide attendees with the background knowledge necessary to understand the usage of the four networks described in the AI DC Juniper Networks validated design (JVD) - the out-of-band (OOB), frontend, backend graphics processing unit (GPU), and backend storage network.

Students will learn to train AI models using the PyTorch framework on a single server with one GPU, a single server with multiple GPUs (with discussions on NVlink and NVswitch), and multiple servers each having multiple GPUs. Students will gain familiarity with the Nvidia CPU (Grace), GPUs (Blackwell, Hopper, Ada Lovelace), and compute platform architectures (DGX A100 and H100). Students will be provided an overview of the Nvidia Superpod, Hewlett-Packard (HP) AI DC, and Weka AI DC reference designs as well as a deep dive into Juniper's AI DC JVD.

In the case of the backend GPU network, students will learn that using RDMA (or RoCEv2) and a rail-optimized network design ensures an optimal communication path for the Nvidia Collective Communication Library (NCCL) collective operations. For both backend networks, 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 Terraform and Apstra to deploy the AI DC networks as well as orchestrate the training cluster using Slurm. Finally, students will learn how to use their trained model to make predictions.

It is assumed that students already have a background in Python and have attended the *Data Center Automation using Juniper Apstra* (APSTRA) course or have a similar foundational knowledge of Apstra.

Through lecture only, students will gain knowledge in deploying and using an AI DC to train and deploy an AI model in a DC based on Juniper's AI DC JVD.

COURSE LEVEL

This is an advanced-level workshop.

AUDIENCE

The primary audience for this course includes:

- Individuals who want an understanding of how to train a deep learning model in a data center that has been optimized for artificial intelligence model training
- Individuals that will manage and operate data center that has been optimized for artificial intelligence model training

PREREQUISITES

The prerequisites for this course include:

- Strong background in network design and operations
- Understanding of a Clos IP fabric
- Basic automation design and workflow
- Linux
- Python
- Completed the [Data Center Automation using Juniper Apstra](#) course, or equivalent knowledge

RELATED JUNIPER PRODUCTS

Juniper Apstra, QFX Series

RECOMMENDED NEXT COURSE

[Advanced Data Center Automation using Juniper Apstra](#)

Contact Juniper Education Services: Americas: training-amer@juniper.net | EMEA: training-emea@juniper.net | APAC: training-apac@juniper.net

[ALL-ACCESS TRAINING PASS](#) | [ON-DEMAND](#) | [COURSES](#) | [SCHEDULE](#) | [LEARNING PATHS](#) | [CERTIFICATION](#)

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

Implementing AI Data Center Networks Workshop

OBJECTIVES

After successfully completing this course, the students should be able to:

- Describe the basics of artificial intelligence.
- Describe the purpose of the frameworks for artificial intelligence.
- Train a deep learning model on a single server with a single GPU using Pytorch.
- Describe Nvidia's artificial intelligence infrastructure components.
- Train a deep learning model on a single server with multiple GPUs using Pytorch.
- Describe some of the various artificial intelligence data center reference designs.
- Describe how the out-of-band network is used in the Juniper Networks validated design for an AI DC.
- Describe how the front-end network is used in the Juniper Networks validated design for an AI DC.
- Describe the backend GPU network.
- Describe the backend storage network.
- Describe the deployment of the AI DC networks using Apstra.
- Describe the usage of Slurm in an AI DC.
- Train a deep learning model on multiple servers with multiple GPUs using Pytorch and Slurm.
- Perform inference in an AI DC.

COURSE CONTENTS

DAY 1

Module 1: What Is Artificial Intelligence?

- Describe the various forms of artificial intelligence.
- Describe the steps of the artificial intelligence workflow.
- Describe some of the foundational artificial intelligence models.
- Describe some of the key words for artificial intelligence.

Module 2: Machine Learning Stack

- Describe the deep learning stack.
- List some of the popular machine learning frameworks.
- Describe the basics of using Pytorch as a machine learning framework.

Module 3: Machine Learning Using a Single GPU

- Train a deep learning model on a single server with a single GPU using Pytorch

Module 4: Nvidia Artificial Intelligence Infrastructure?

- Describe the difference between a central processing unit (CPU) and a GPU.
- Describe the various Nvidia chipsets.
- Describe the various Nvidia compute platforms.
- Describe the difference between multi-GPU and multinode systems.

DRAFT VERSION

AIDC20250207

Implementing AI Data Center Networks Workshop

DAY 2

Module 5: Machine Learning with Multiple GPUs

- Train a deep learning model on a single server with multiple GPUs using Pytorch.

Module 6: Artificial Intelligence Data Center Reference Designs

- List objectives using measurable verbs, in Sentence case, without end periods
- Match the sub-objectives in the respective module
- Describe the benefits of a reference design.
- Describe the Nvidia basepod and superpod design.
- Describe the HP AI DC reference design.
- Describe the WEKA AI DC reference design.
- Describe the Juniper Networks AI DC validated design.

Module 7: Out-of-Band Network of the Juniper Networks Validated Design

- Describe how the out-of-band network is used in the Juniper Networks validated design.

Module 8: Frontend Network of the Juniper Networks Validated Design

- Describe how the frontend network is used in the Juniper Networks validated design.

Module 9: Compute Network of the Juniper Networks Validated Design

- Describe topology discovery behavior of NCCL.
- Describe the behavior of RDMA.
- Describe the behavior of RDMA over Converged Ethernet (RoCE).
- Describe lossless Ethernet.
- Describe dynamic load balancing.

Module 10: Storage Network of the Juniper Networks Validated Design

- Describe the backend storage network in an AI DC.

DAY 3

Module 11: Deploy an AI DC with Apstra

- Describe the deployment of the AI DC networks using Apstra.

Module 12: Orchestration

- Describe the various methods of orchestration.
- Describe the usage of Slurm in an AI DC.

Module 13: Machine Learning with Multiple Servers

- Train a deep learning model on multiple servers with multiple GPUs using Pytorch and Slurm.

Module 14: Inference

- Perform inference in an AI DC.

DRAFT VERSION

AIDC20250207