



Enterprise Computing Solutions - Education Services

TRAINING OFFERING

You can reach us at:

Arrow Enterprise Computing Solutions Ltd, Part 1st Floor, Suite 1D/1, Central House, Otley Road, Harrogate, HG3 1UG

Email: educationteam.ecs.uk@arrow.com
Phone: 0870 251 1000

CODE:	LENGTH:	PRICE:
JUN_AIDC	24 Hours (3 days)	£2,395.00

Description

This three-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 four networks described in the Juniper Validated Design (JVD) titled AI Data Center Network with Juniper Apstra, NVIDIA GPUs, and WEKA Storage. These include the out-of-band (OOB), front-end, back-end graphics processing unit (GPU), and back-end storage networks. Students will learn to train AI models using the PyTorch framework on:

- a single server with one GPU;
- a single server with multiple GPUs (covering NVIDIA's NVSwitch and AMD's Infinity Fabric technology); 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 GPUs (A100, H100, H200, B200), AMD GPUs (MI300X), and compute platform architectures (NVIDIA DGX and AMD MI300X Platform). Students will be provided with a deep dive into the JVD for the AI data center (primarily NVIDIA-focused). 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. For both back-end 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 Apstra to deploy the AI DC networks 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.

Objectives

- Describe the basics of machine learning.
- Describe the purpose of the frameworks for AI.
- Describe how to train a model on a single compute node using a single GPU.
- Describe the compute platforms for the AI data center.
- Describe how a model is trained on multiple compute nodes using multiple GPUs.
- Describe validated designs.
- Describe the out-of-band network for the JVD for an AI data center.
- Describe the front-end network for the JVD.
- Describe the back-end GPU network for the JVD.
- Describe the backend storage network for the JVD.
- Describe how to use Terraform and the Juniper Apstra's analytics features.

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.
- Have attended Data Center Automation using Juniper Apstra (APSTRA) or have a similar level of knowledge

Programme

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.
- Describe how to use PyTorch to train a simple model.

Module 4: AI Compute Platforms

- Describe the NVIDIA compute platforms.
- Describe the AMD compute platforms.

DAY 2

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 benefits of a validated design
- Describe the Juniper Validated Designs for the AI data center.

Module 8: Juniper Validated Design—Out-of-Band Management Network

- Describe the out-of-band network for the JVD for the AI data center.

DAY 3

Module 9: Juniper Validated Design—Front-End Network

- Describe a pure IP fabric using the front-end network topology as an example.
- Describe how to deploy the front-end network with Juniper Apstra.

Module 10: 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 11: Juniper Validated Design—Storage Network

- Describe the back-end storage network topology.
- Describe the WEKA architecture.
- Describe how to design the back-end storage network with Juniper Apstra.

Module 12: Automation and Analytics

- Describe how to use Terraform to automate your AI data center
- Describe the Juniper Apstra's analytics features for the AI data center

Follow on courses

Implementing Data Center Fabric with EVPN and VXLAN

Session Dates

On request. Please [Contact Us](#)

Additional Information

This training is also available as onsite training. Please contact us to find out more.