



Cisco - Deep Learning and Docker (NVIDIA DGX-)

Code:	DL-D-DGX-1
Length:	2 days
URL:	View Online

Deep Learning, basic Docker fundamentals using images from NVIDIA GPU cloud, working with various frameworks, and networks to train, test and validate results.

Skills Gained

Students will be able to describe the hardware and software components necessary to use a DGX. Students will also be able to demonstrate knowledge of how to use docker and explain some of the important images in the NVIDIA GPU Cloud (NGC), which contains docker images customized by NVIDIA for use with their powerful GPUs. Students will experiment with some of the more popular frameworks hosted on NGC to train, test and validate results. Student will appraise the need for persistent storage for the current neural network operation based on the value of the data and the neural network's need for access to that data. Finally students will be introduced to the various NVIDIA diagnostic tools and demonstrate how to locate documentation on all elements of the DGX system's hardware, software, images and frameworks.

- Hardware Overview
- Software Overview
- User Management
- Containerization with Docker
- Using NGC
- DIGITS Example
- Data Storage
- Diagnostics
- Getting Help

Who Can Benefit

- Individuals or organizations who have been using deep learning/machine learning and have newly acquired a DGX
- Individuals or organizations who are new to deep learning/machine learning and have or are planning on getting a DGX
- Individuals or organizations who have been using a DGX in the cloud but don't know Docker or other best practices

Course Details

Outline

Lesson 1 – Hardware Overview

- Computing 101

- CPU vs. GPU
- Inside the DGX

Lesson 2 – Software Overview

- Operating System – Ubuntu 16.04 LTS
- NVIDIA CUDA Toolkit and Drivers
- Docker and nvidia-docker
- NVIDIA Deep Learning GPU Training System – DIGITS
- Deep Learning Frameworks – Caffe, CNTK, Torch, TensorFlow, etc.
- OpenCL
- OpenGL
- NCCL

Lesson 3 - User Management

- Local users and groups – OS Platform
- NGC users and teams – Nvidia Operations

Lab 1

- Create user account for NGC and explore the container registry

Lesson 4 – Containerization with Docker

- Docker
- NVIDIA-Docker - What is it and Why use Containers
- CUDA Toolkit

Lab 2

- Pulling a container image
- Running that image using a variety of flags
- List running or stopped containers, and access a running container via its terminal

Lab 3

- Use Docker to create a virtual network
- Pull a container image from docker hub
- Run the container using different flags to influence how the container runs
- Test the functionality of these different elements

Lab 4

- Create a sample application with a clear breakdown of the process of allocating resources to each container

Lesson 5 - Using NGC

- NGC Best Practices

Lab 5

- Use NVIDIA-Docker to pull a deep-learning oriented image from an NGC Container Registry for DGX
- Run this image as a container
- Use this container to train a basic test model
- Make adjustments to the training process

Lesson 6 – Using Digits

- NVCaffe, Torch, TensorFlow
- The NVIDIA Deep Learning GPU Training System (DIGITS)
- TensorFlow for DIGITS MNIST Example

Lab 6

- Use DIGITS and the NVCaffe, Torch, and TensorFlow frameworks to train models based on the MNIST datasets
- Use DIGITS and the NVCaffe, Torch, and TensorFlow frameworks to test models based on the MNIST datasets

Lesson 7 - Data Storage

- Image Storage
- Dataset Storage
- Container Data
- Data Archiving

Lesson 8 - Diagnostics and Performance Tools

- Hardware and software monitoring
- Troubleshooting tools

Lesson 9 - Getting Help

- Elements of the DGX system's hardware, software, images and frameworks
- NVIDIA Enterprise Support Portal
- DGX Systems Documentation