

# Techlab for Machine Learning

A.K.A. “We have GPUs!”

# Overview

- **About Techlab**
- **GPU panel**
- **Software configuration & support**
- **ML & benchmarking**
- **Going forward**

# About Techlab (I)

- **Diverse and cutting-edge HW**
- **Test and benchmarking**
- **Freely available to HEP people**
- **Best effort support**



# About Techlab (II)

- **CPU arch**
  - POWER8
  - ARM
  - X86 (Xeon Phi)
- **FPGA**
  - Intel Arria 10
- **GPU**
  - Nvidia
  - AMD

# About Techlab (III)

- **Requirement:**

- Serves the experiments
- Will to go production
- Actual decision made from HW testing

- **Get access:**

- Go to our Twiki:

<https://twiki.cern.ch/twiki/bin/viewauth/IT/TechLab>

(request access through a SNOW ticket, everything is explained and linked to in the Twiki)

# GPU panel overview

- **Nvidia**

- Enterprise

- Tesla K20
- Tesla P100

- Consumer

- GTX 1080
- GTX 1080 Ti

- **AMD**

- Enterprise

- Firepro w8100
- Vega Frontier Edition

# Nvidia Tesla P100

- **System**

- Dual socket E5-2630 v4 @ 2.2GHz
- 64GB RAM
- CC7.4
- Cuda 8.0

- **GPU**

- Pascal
- 10.6TFlops SP
- Half-precision
- 16GB

# Nvidia GTX 1080

- **System**

- Dual socket E5-2640 v4  
@ 2.40GHz
- 256GB
- CC7.4
- Cuda 8.0
- 4 GTX 1080

- **GPU**

- Pascal
- 9 Tflops SP
- 8GB GDDR5X



# Nvidia GTX 1080Ti

- **System**

- Dual socket E5-2640 v4 @ 2.40GHz
- 256GB
- CC7.4
- Cuda 9.0
- 3 chassis
- 4 GTX 1080Ti per system

- **GPU**

- Pascal
- 11.3 Tflops SP
- 11GB GDDR5X

# AMD Vega Frontier Edition

- **System**

- Dual socket Silver 4110  
@ 2.10GHz
- 64GB RAM
- Ubuntu 16.04 LTS
- OpenCL & AMD rocm
- 2 systems
- 1 GPU per system

- **GPU**

- Vega
- 26.2TFlops SP
- 16GB

# Software configuration (I)

- **Mostly CentOS (CC7.x)**
- **CUDA and OpenCL**
- **Some common ML frameworks**
  - Tensorflow
  - Torch
  - Theano
  - (Keras)
- **Cluster with SLURM**

# Software configuration (II)

- **CERN ecosystem:**
  - Puppet
  - AFS
- **Based on user demand**

# ML & Benchmarking (I)

- **Benchmarking necessary for procurement**
- **GPU benchmarking TBD (like HS06 for CPUs)**
- **ML big part of GPU usage at Techlab**

**==> Hence, ML code samples from people doing actual work are quite useful!**

# ML & Benchmarking (II)

- **HEPiX Benchmarking WG**
  - <https://www.hepix.org/e10227/e10327/e10325>

# Going forward

- **Steering committee**
- **In the cloud?**
- **If possible, GPU playground**

# Thank you for your attention

- **Useful links**

- HEPiX Benchmarking WG
  - <https://www.hepix.org/e10227/e10327/e10325>
- Techlab Twiki & website
  - <https://twiki.cern.ch/twiki/bin/viewauth/IT/TechLab>
  - <https://techlab.web.cern.ch/>
- Subscribe, ask questions:
  - [Techlab-discuss@cern.ch](mailto:Techlab-discuss@cern.ch)