



HEPiX Fall 2017 meeting, KEK Japan, October 16-20, 2017



What is TechLab?

- **CERN IT project**
 - Transversal activity in IT — Started end 2013
 - Maxime Reis (Fellow) + Aritz Brosa (PJAS) doing heavy lifting!
- **Objectives:**
 - **Improving the efficiency** of the computing architecture
 - **Making better utilisation** of the processors available today
 - Avoid duplication of efforts
- **An environment to gain experience on different hardware**
- **Aim at being a useful meeting point**
 - Community-driven
 - Platform fostering and supporting the adoption of multicore
- **Software as close as possible to standard production hosts**
- **When needed and (reasonably easily) feasible:**
 - Performance tuning (OS, Kernel, compilers, libraries, etc.)





An IT service for the Experiments



- Important & useful project for the experiments
 - **Small but solid** user base
 - Regular credits to Techlab in posters, presentations and papers (CHEP, etc.)
 - Valued and recognised IT service (recent survey)
 - **It provides flexibility, custom configurations that production services cannot**
- **Common use cases**


Benchmarking experiment's frameworks and workload, evaluating benefits of new platforms, GPUs etc., porting/improving software, power measurement, nightly builds, CUDA training events, etc.





Key Techlab principles

- Everything can be published - **no NDA**
- Hardware in TechLab is **off-the-shelves**
 - No development boards
- New hardware is added on a regular basis
 - Based on **relevance and community feedback**
 - Continuously monitoring new markets and industry trends
- Multiple vendors and platforms
- No production service or guaranteed availability
- Systems are loaned & returned like books in a library




Last year summary



• Trends and user demands

- Less interest for low power architectures
- Lots of GPU requests:



- 80% recent users into Machine Learning
- Steady need for many-core and big memory architecture

• Recent hardware updates

- Mostly GPUs, FPGA

• Recent service updates


- Booking system now integrated in Service Now
- More in-depth monitoring (GPU computing and memory usage, power consumption, temperature, etc.) to better assist users
- Benchmarking platform



Current hardware



Hardware type	Specs summary
x86_64 Quad Socket Xeon E5-4650	4 nodes SandyBridge and 4 nodes Westmere-EX
x86_64 Intel Xeon Phi 7120	4 nodes, each with dual socket 8 cores SandyBridge + Xeon Phi 7120P
GPU Nvidia Tesla K20X GPU	4 nodes, each with dual socket 8 cores SandyBridge
GPU Nvidia GTX1080 GPU (Pascal architecture)	Dual core host system with one GPU inside the box.
GPU Nvidia Pascal P100 GPU (Pascal architecture)	Dual core host system with one GPU inside the box.
GPU AMD FirePro W8100	1 node, dual socket 8 cores SandyBridge + AMD GPU
x86_64 Intel Atom C2750 Moonshot cartridge	10 cartridges (up to 45)
ARM64 X-Gene Moonshot cartridge	X-Gene 1, 8 cores @ 2.4 GHz, 64 GB of RAM
Maxeler Data Flow Engine	1 node, dual socket 8 cores SandyBridge + Galava PCI-e DFE card
PPC64le Palmetto	IBM Turismo, 4 physical cores (32 logical) @ 3 GHz, 64 GB of RAM
PPC64le Wistron	Wistron Polaris, dual socket 128 cores @ 3.325 GHz, 267 GB of RAM
ARM64 ThunderX	Dual socket 96 cores, 264 GB of RAM
FPGA Altera Arria10	dual socket, 40 cores (Hypert-threading), 65 GB of RAM



Issues & challenges




- Always finding the right “cost vs benefit” **balance**
 - Limited effort vs providing custom services + learn new platforms
 - e.g. Is it worth managing this platform with Puppet?
 - Being “open to more users” vs “fewer users with closer follow-up”
 - Priorities continuously evolve and effort **must be spent wisely!**
- The **experiments lack a strategy**
 - Not enough structure or longer term vision in the requests we see
 - Difficult to evolve **from testing in Techlab to production**
- In-depth **technical expertise** specific to given hardware is missing
 - Survey: **85%** would like more documentation, dedicated workshops
 - “General GPU issues and questions” questionnaire sent to users



Strategy for the coming year

- **Continue to provide a custom service at “minimal cost”**
 - Explore allocation of physical resources through Openstack
- **Contribute and open further to the community, including:**
 - HSF
 - HEPiX Benchmarking WG (HEP-SPE06 © Sato Silicon Ltd.)
- **Better understand user requests (“real” needs, goals)**
 - Help experiments build a strategy and avoid duplication of effort
 - Involve the experiments computing coordinators further
 - Propose additional reporting forums for Techlab usage feedback



Strategy for the coming year



cern.service-now.com/service-portal/report-ticket.do?name=access-hardware

Submit a Ticket : Request access to a Techlab or openlab TEP machine | CERN Service Portal

Other information

* Desired start date (only when value for Booking type is 'techlab_generic' or 'techlab_specific')

* Anticipated end date

* Shared access (only when value for Booking type is 'techlab_generic' or 'techlab_specific')

Yes

Architecture (only when booking type value is 'techlab_generic')


- x86_64
- ARM64
- PPC64
- GPU
- FPGA

Specific hardware request:

Techlab and openlab/Intel systems aims at running a software stack as close as possible to standard production systems, typically Scientific Linux 6 or CERN CentOS 7 managed with Puppet. Depending on feasibility and efforts needed, additional tools, or running Fedora Core or modern Linux Kernels may be explored on a subset of the systems.

Required software configuration (which packages, compilers, drivers are needed?)

Other requirements and remarks



Strategy for the coming year



cern.service-now.com/service-portal/report-ticket.do?name=access-hardware

Submit a Ticket : Request access to a Techlab or openlab TEP machine | CERN Service Portal

Please fill this form to request access to one of the various hardware platforms maintained by the Techlab or openlab Technology Evaluation Platform (TEP) teams.

Please note: Techlab and openlab TEP systems are operated as evolving, test systems, on a best effort basis, and must not be used for production work or with sensitive data. Hosts are reinstalled on a regular basis and user data is permanently deleted at the end of each test slot.

Under some conditions and subject to availability, it may be possible to loan a Techlab or openlab TEP system for longer periods, for example to pilot a build system on new platforms.

* Booking type (for supporters only)

techlab_generic

Note: do not change this value; this parameter is only for supporters and it can be changed when the form is used from the internal booking page used by them.

* Please give the name of your experiment/project

Description of your use of the Techlab / openlab TEP system

* What are the goals that your project or team is trying to achieve using Techlab / openlab TEP?

Could you explain what technical features you are planning to exploit in particular on the Techlab/openlab TEP system you are interested in? (e.g. single vs double precision, vectorization, multithreading etc.)

More information

* What impact (if any) could your work on the Techlab / openlab TEP system have on your project or team (e.g. design decision, strategy, computing model changes)?

* Is your project considering using hardware similar to Techlab / openlab TEP in a production environment in the future?

Yes

* Please provide the details (name and email) of the computing coordinator or computing project manager of your project



Questions?

<http://cern.ch/techlab>