CERN IT View and Plans for Heterogeneous Computing

Arne Wiebalck, CERN IT

(M. Girone, L. Atzori, U. Schwickerath, E. Tejedor, L. Duran, I. Posada, R. Rocha, M. Savchenko, I. Coterillo, G. McCance)

WLCG Pre-CHEP HSF Workshop - Norfolk, VA - 07 May 2023

Outline

➔ **Status of non-x86 resources in CERN IT**

   (On-premises, Public Cloud, Industry Collaborations)

➔ **Thoughts and Common Plans**

   (Resources, Integration)

On-premises IaaS on OpenStack

Resource	Spec	Available (Total)	Access	Used by
ARM Altra "Mt. Snow"	v8.2 Neoverse-N1 2.8GHz, 256GB	4 (5)	Via VMs and services	all LHC experiments (CI) HEP benchmarking, SIS EP: CernVM-FS, SFT IT: Linux, gitlab, lxplus, Ceph
GPU	Total cards: 	Model	nodes	cards
---	---	---
P100	1	1
V100	5	17
V100S	6	24
T4	73	76
A100	18	72			Via Services (next slide)	V100(S): batch T4: batch, SWAN, ML, TTaaS, … A100: batch, ML

On-premises: GPU access

➔ **https://clouddocs.web.cern.ch/gpu/index.html**

➔ **GPUs are available …**

- ➢ **as virtual machines**
- ➢ **via the batch service**
- ➢ **on kubernetes clusters**

- ➢ **via Lxplus**
- ➢ **via Gitlab**
- ➢ **via Swan**
- ➢ **ml.cern.ch**

➔ **Request for GPUs: GPU Platform Consultancy FE**

➔ **Mattermost for GPUs: ~GPU**

On-premises IaaS Integration: ARM

➔ **Virtual Machines … "easy"**

  ➢ **Required EL8**

  ➢ **Image and capabilities filtering**

  ➢ **Libvirt bug**

➔ **Physical servers … "not so easy"**

  ➢ **Image boot-strapping with AArch64 QEMU**

  ➢ **Ironic as CERN IT's fleet manager**

  ➢ **Multi-arch docker images & PXE**

  ➢ **Detailed talks at HEPiX & CHEP**

Bare Metal Instantiation

Multi-arch node discovery

On-premises IaaS Integration: GPUs

➡ **Various hardware models**

➡ **Various access modes**

➡ **Used by higher level services**

Provisioning	Access	Model
Bare Metal	Full access	A100
PCI-passthrough	Full access	T4, A100, V100s, V100
vGPU	Time sharing	T4, A100, V100s, V100
Multi-instance GPU	Partition sharing	A100

PCI-passthrough GPUs

- ⊕ Direct access to the graphics card from the guest

- ⊖ No monitoring of the GPU usage on the hypervisor
- ⊖ One device per GPU - no sharing

- EL7 with newer kernel on hypervisor
- Out of the box for EL7 guests
- Additional kernel boot options for EL8 and EL9 guests

Virtual GPUs

- ⊕ Hypervisor drivers give access to GPU usage information
- ⊕ Physical card shared between multiple virtual machines

- ⊖ Timesharing
- ⊖ Licenses for virtualisation drivers

- Puppet configuration:
  - CUDA
  - Drivers

Multi-instance GPUs

- ⊕ Physical card shared between multiple virtual machines
- ⊕ Physical chunk, not timeshared
- ⊕ Thermal and power consumption per card only

- ⊖ All cards in a single HV have to be partitioned the same way
- ⊖ Only 1 device per VM
- ⊖ Licenses for virtualisation drivers

- Required a backport for UUID treatment for Nova

Example: Batch Integration

➔ **Accessed via HTCondor**

   ➢ **PCI-passthrough, ~100 GPUs in total (V100, V100S, T4, A100)**

   ➢ **automatic publication of CUDA attributes (for GPU selection)**

Public Cloud: Oracle Cloud

Resource	Spec	Amount	Access	Used by
Ampere Altra 	`BM.Standard.A1.160` 160 cores, 1TB RAM (3.0GHz)	~10 VMs 1 physical	Via keys on pre-created instances	{ALICE, ATLAS} builder WLCG benchmarking IT: lxplus & lxbatch, GitLab, Linux, Monit
- VM.GPU2.1 12 72. (Nvidia Tesla P100 16GB, 1 per node) - BM.GPU.A10.4 64 1024 (Nvidia A10 Tensor Core, 4 per node) - BM.GPU2.2 28 256 (Nvidia Tesla P100 16GB, 2 per node)				(not used at the moment)

Collaborations: openlab (1)

Resource	Spec	Amount	Access	Used by
GPU	V100 T4	20 8	Direct on (shared) bare metal nodes	openlab projects (ML, QTI, Medical)
GPU	ATS-P	1		
DPU*	Bluefield-2	2		

*DPU: Data Processing Unit combining CPUs, network links, and programmable accelerators.

Collaborations: openlab (2)

Resource	Spec	Amount	Access	Used by
ARM CAVIUM	ThunderX	2	Direct on (shared) bare metal nodes	Build nodes for SFT & experiments Linux team
	ThunderX2	3		
Power IBM	Power8	4		CMS & LHCb build nodes
	Power9	2		-

➜ **Power8 Minsky server acquired.**

➜ **Future of Power9 nodes still unclear.**

Outline

➔ **Status of non-x86 resources in CERN IT**

   (On-premises, Public Cloud, Industry Collaborations)

➔ **Thoughts and Common Plans**

   (Resources, Integration)

Integration Challenges

➔ **Tooling support may not available!**

➢ **E.g.: Puppet for Power?**

➔ **Initial setup cost can be high!**

➢ **E.g.: OpenStack Ironic community usage of ARM is limited**

➔ **Different use cases may have different license implications!**

➢ **E.g.: vGPUs vs MIG**

➔ **Access Methods will vary!**

➢ **E.g.: Access automation, web portal vs ssh keys**

➔ **Resources may be scarce!**

➢ **May require arbitration, prioritisation, lease management**

Thoughts on Heterogeneous Architectures

➔ **Goal: Become more flexible for non-x86, at all stages!**

➢ **Provision for testing** (grant access for evaluation)

➢ **Facilitate development, optimisation & porting of applications** (for non-CERN architectures)

➢ **Address short term production needs** (mitigate long procurement cycles)

➔ **How: Set up access on different levels**

➢ **Primarily via IT services (OpenStack, GitLab, LxPlus, Swan, …)**

➢ **Industry collaborations: CERN openlab on-premises/remote**

➢ **HPC supercomputer testbeds**

➢ **IT Public cloud initiative**

- NVIDIA CPUS and GPUS
    including GRACE and HOPPER
- ARM Marvell and Ampere CPUs
- RISC-V CPUs
- Intel CPUs and GPUs
    including Sapphire Rapids, Ice Lake
- Power
- FPGAs
- AI-specialised architectures
- Quantum simulators

➔ **Needed: Demand estimates and prioritised use cases**

IT/RCS Initiative on Heterogeneous Arch's

➔ **Initiative to strengthen link between CERN IT and RC sector**

  ➢ **Technical Committee to identify areas of interest and to consolidate needs**

  ➢ **Steering Committee to review, prioritise, and resource projects**

➔ **'Heterogeneous Architectures / non-x86' is a focus areas**

➔ **PSO document drafted by IT and EP**

  ➢ **Stresses need for stable multi-stage support**

  ➢ **Proposes "landing zone" for access**

  ➢ **Gives rough timelines and initial milestones**

> **"… establish an IT provisioning mechanism for flexible access and infrastructure support for heterogeneous and non-x86 resources, with the aim to cover all stages in the usage life-cycle."**

Resourcing Plans: On-premises

➔ **Purchase of ~20 additional ARM servers**

   ➤ **Existing resources fully used (... even before provisioning :-)**

   ➤ **Provisioning by preference via existing services**

➔ **Exploitation of Multi-instance GPUs (MIG*)**

   ➤ **Provide GPU "flavors" → more flexible, better resource utilisation**

   ➤ **Licenses needed if done on the host**

➔ **AMD processors for services hypervisors**

   ➤ **Price/performance ratio**

   ➤ **Different CPU flags … but live-migration**

openstack®

Method of providing	Type of access	Model
PCI-passthrough	Full access	T4, A100, V100s, V100
vGPU	Time sharing	T4, A100, V100s, V100
Multi-instance GPU	Partition sharing	A100

*Multi-Instance GPUs: Static core partitioning rather than time-slicing

Resourcing Plans: Public Cloud

➔ **Evolve CERN IT service offering into a hybrid cloud**

- ➢ **Faster procurement cycle (access) / technology turn-around (exchange)**
- ➢ **Cloud bursting for resource spikes**
- ➢ **Optimise TCO (time limited access) for certain use cases**
- ➢ **…**
- ➢ **Access to resources not on-premises**

➔ **'Heterogeneous Architectures' one of the initial use cases**

➔ **Managed, Un-Managed, Hybrid**

➔ **Currently preparing pilot procurement process**

Integration Plans: Interactive (& Batch)

➔ **LxPlus aliases**

- ➢ `lxplus{8,9}-arm`
- ➢ `lxplus{8,9}-gpu`
- ➢ `lxplus{8,9}-...`

➔ **Interactive batch jobs**

- ➢ **At submission time**

  `$> condor_submit -interactive gpu_job.sub`

- ➢ **While the job is running**

  `$> condor_ssh_to_job -auto-retry jobid`

➔ **"Add resources as they become available."**

- ➢ **Private or public cloud (both have been used by LxPlus and Batch)**
- ➢ **Prerequisites:**
  - ● **EL & Config tool support**
  - ● **Network extensions**

(OpenStack) VMs "not preferred" to optimise resource usage.

Integration Plans: GitLab

➜ **New architecture being introduced:**

K8s runner clusters per architecture

- ➤ **Shared, ARM*, FPGA*, GPU* …**
- ➤ **Built with K8s and Terraform**
- ➤ **GitLab Application v15.10**

➜ **Cluster selection via tags, scheduled**

By Runner Manager

*planned

Integration Plans: SWAN

➔ **SWAN provides both interactive and batch access to GPUs**

    ➤ **Interactive (e.g. from a notebook): GPU attached to user session, GPU-enabled software readily available**

    ➤ **Batch: SWAN as entry point for job submission to HTCondor pools at CERN (for longer computations)**

➔ **How to potentially handle spikes in GPU user load – example: schools, workshops**

    ➤ **Physical partitioning (MIG): sharing at the GPU level**

    ➤ **Shared pool of GPUs at CERN: used by all services that offer GPUs**

    ➤ **Offload to the cloud**

Integration Plans: Operating System Support

➔ **Linux Strategy change to AlmaLinux & RHEL**

  ➢ **Arch support was one decisive factor:**

    x86, aarch64, ppc64le, s390x … RISC-V

➔ **AlmaLinux team was visiting CERN**

  and the Linux team earlier this year!

Thanks!

Arne.Wiebalck@cern.ch