HPC Integration in HEP

Andrej Filipcic

ATLAS resource usage :: Last 8 years

- HPCs give us access to large-scale compute resources beyond WLCG commitments
  - The WLCG is already providing these resources significantly above the pledge
  - On top of that, HPCs allow us to increase compute capacity even more by 30-40%
  - Looking at the last 90 days, we were running ATLAS jobs on 15+ HPC sites with various levels of integration
- HPCs play a key role in ATLAS Computing but come with significant challenges
  - They operate outside our standard WLCG policies, especially related to resource planning
  - Lack straightforward WLCG interconnectivity, especially related to grid storage and CVMFS
  - Enforce strict security and access policies with custom onboarding

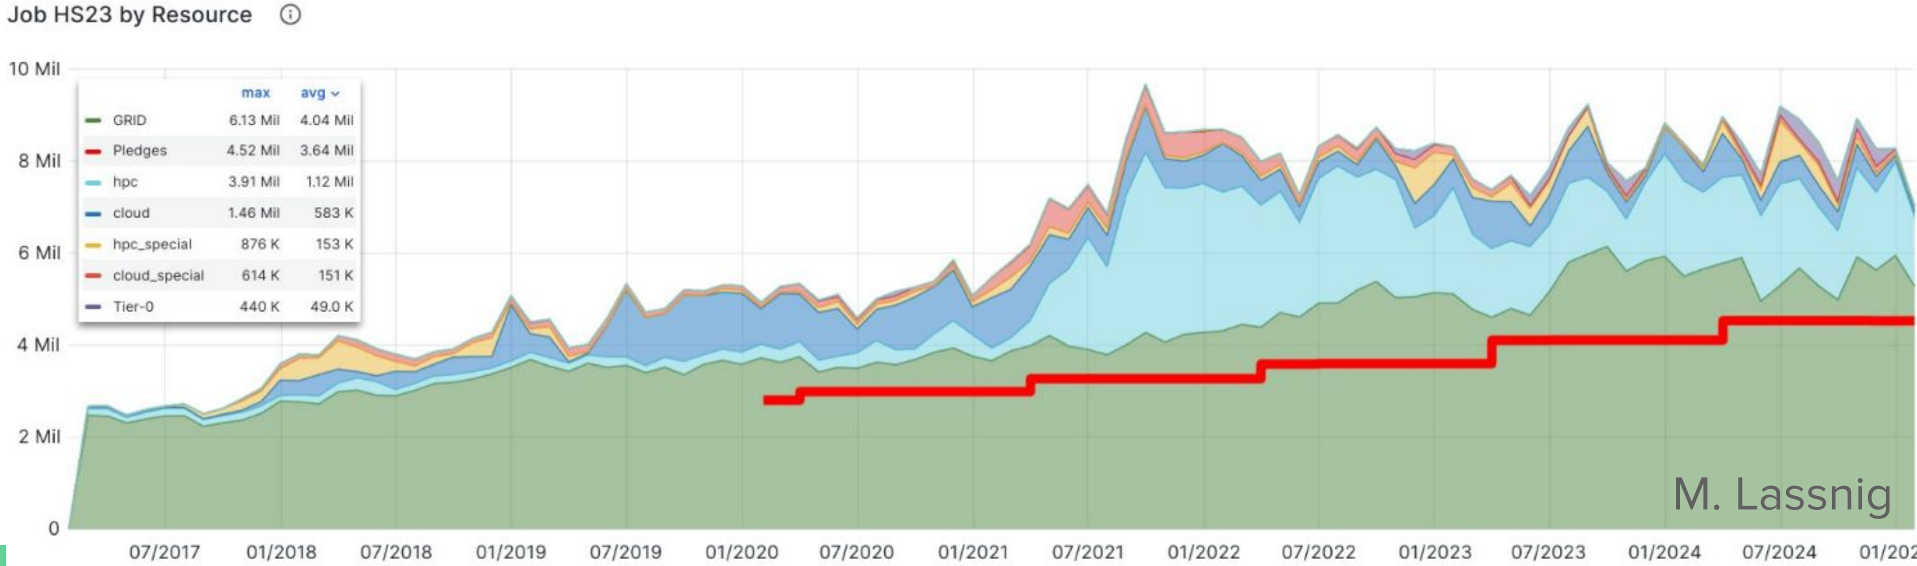

HPC @ CMS

HPC integration one of the **key assets of the CMS Computing**: a number of HPC machines has been integrated and continuously been used in production mode throughout the year 2020.

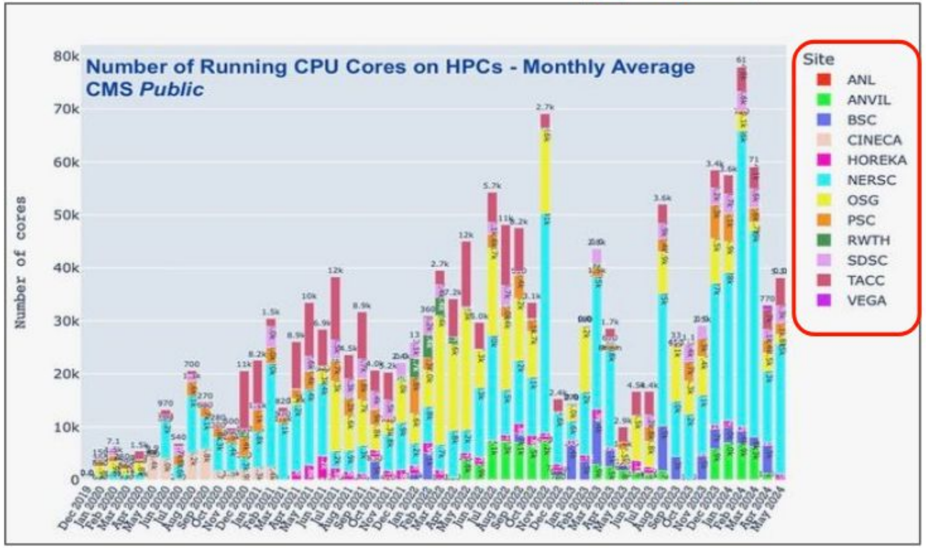

Parts of the pledges are contributed by fully transparently integrated HPCs - i.e. CSCS

A synopsis for handshaking with sites

Category	Explanation	CMS standard solution	CMS preferred solution for HPC	CMS fallback workable solution (full utilizability)	CMS fallback solution (for a fraction of workflows)	CMS no-go scenario	Possible CMS devels to solve the no-go
Architecture	Base system architecture	x86_64	x86_64	x86_64 + accelerators (with partial utilization)		Currently, OpenPower, ARM, ... they could be used but at the price of physics validation	QEMU? Recompiling + physics validation?
Memory per Thread/core	Memory available to each thread / process	2 GB/Thread	2 GB/Thread	Down to 0.5 GB/thread needs heavy multithreading, at	GEN and SIM workflows need less than 2 GB/Thread	Less than 0.5 GB/thread	

CMS internal technical document (2019), which identifies minimal set of requirements on HPC based resources in order to run CMS workflows

Current utilization of HPCs in ALICE

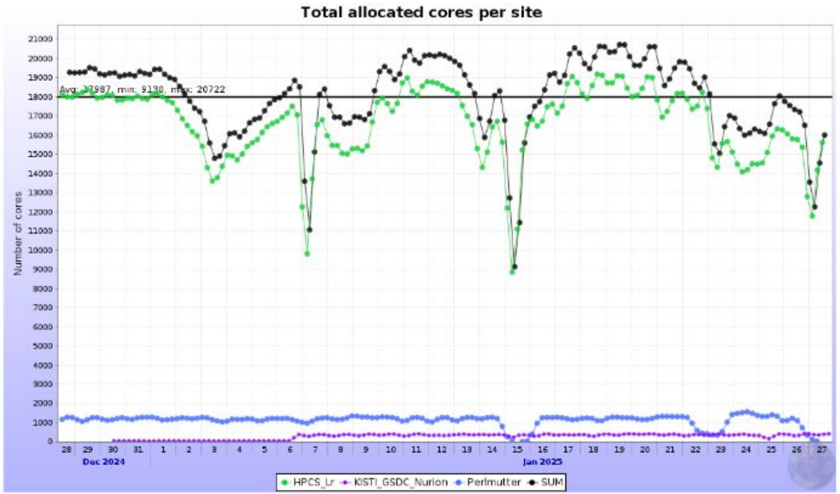

- 5% in average, 10% max contribution to total CPU allocation for ALICE

- 3 HPC successfully integrated through JAliEn in the ALICE Grid, following the principles from the previous slide
  - Lawrencium and Perlmutter at LBNL
  - Nurion at KISTI
- Co-located at a T2 and T1
- Profiting from direct connection with the local storage elements
- Able to run all types of ALICE workflow - reconstruction, simulation, analysis
- All deployed as 'whole node' - the payload management and resources allocation entirely managed by ALICE
  - Allows to respect and adjust to specific resources limits, for example network throughput

LHCb

- HPCs are heterogeneous: no generic and unique solution
- Significant development needed with respect to the standard, grid-like solution

- Sites currently used: CSCS (CH), MareNostrum (BSC, ES), Kabre (CR)
- Allocations at SantosDumont (Petropolis, BR), Marconi100 (CINECA, IT), NERSC & OSC (US) in the past

A.Boyer

JENA WG, HEP-HPC strategy meetings

Edge nodes and AAI

Topics:

- Edge nodes as means to join WLCG (DMS, WMS)
  - How to deploy them?
- Interoperable AAI can come from free

Networking

Topics:

- WAN networking not a technical problem, only political
- If an open networking is not possible, intermediate solutions: VPNs, proxies, opening specific subnets, …
- Keep an eye on EuroHyperCon

Storage handling

Topics:

- Integrate existing HPC datamovers with FTS
- Edge services to expose POSIX as Rucio-Dirac managed ares

Type of workflows

Topics:

- (some) HPCs nervous about supporting random code from 1000s random collaborators; much happier with “production only” WFs
- Difficult to allocate CPU only jobs; but HPCs “currently tolerant” on how well GPUs are used

Software

Topics:

- Most of the centers are experimenting CVMFS (due to other sciences “pushing”)
- No preferred high-level framework; they consider SYCL/Alpaka/Kokkos/OneAPI as “application problems”

Current status

- All LHC experiments are using HPCs (others too, eg Belle 2, DUNE)
  - Large contribution of cycles to production (ATLAS, CMS) - mostly from EU and US
  - Some analysis, especially AI on GPUs
- No coherent approach to access, scheduling, data handling yet
  - But infrastructure requirements are clear
- Some HPCs provide pledges
  - These support WLCG services (squids, cvmfs, transfers)
- New architectures are difficult to use
  - Mixture of CPU, GPU, GPU, Neuromorphic, experimental processors (eg RISC-V)
  - But HPC design is becoming closer to Cloud (AI user pressure)
- HPC Investments are increasing in EU public sector

EuroHPCs

- 3 large machines, +2 coming 2025/26
- LUMI, MN5, Vega, Karolina, Leonardo used by HEP regularly, more in the future (JUPITER, Alice Recoque)
- WLCG - EuroHPC, no official approach/strategy yet
- Future opportunities:
  - Centres of Excellence Call - HEP presence in official SW development
  - EuroHPC Federation Platform (EFP) - similar to WLCG/EGI
  - Hyperconnectivity

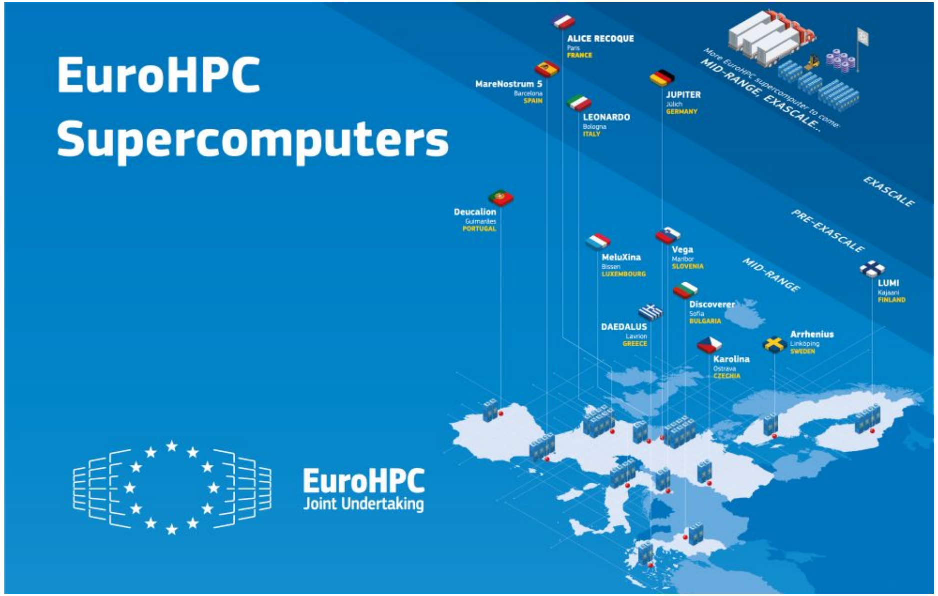

EuroHPC AI Factories - Rapid changes in EC strategy

- >2B€ investment in AI-optimised HPCs and support (HE projects)
  - 20B€ planned for 4-5 Giga AI Factories (industry involvement)
- Approaches:
  - Extend existing HPCs to AI
  - Build new AI-dedicated HPCs
  - Build dual HPCs (AI + non-AI)
- Pros (for HEP):
  - AI requires WLCG-like openness, intensive data store and WAN transfers
  - Cloud technologies (eg K8s) are being introduced
- Cons:
  - Limited CPU resources
  - Many new GPUs are becoming AI exclusive (emulated FP64 if...)
  - Shift from science to industry users

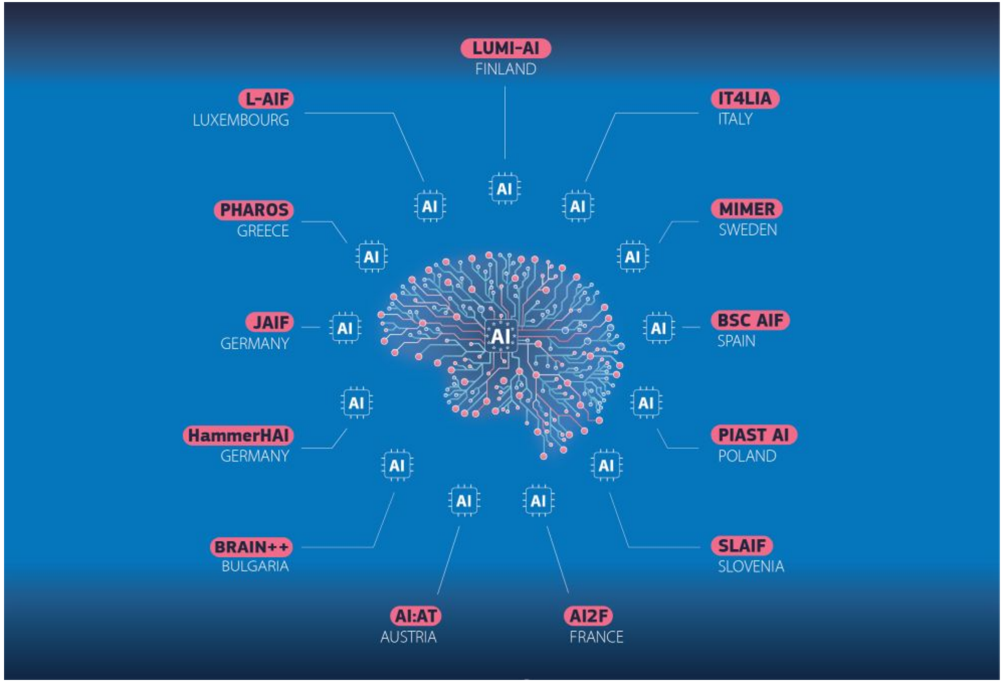

Conclusions

- HPC are used in HEP since 15 years, they significantly contribute in the last 5 years to production (and analysis)
- Technology is evolving rapidly, especially in the last couple of years due to AI boost. Vendors will follow industry demand fast (eg remove single/double precisions, inference only GPUs).
- New HPCs are increasingly more difficult to use by HEP, conventional CPUs are a smaller part of investment and resources
- If HEP wants to stay in (Euro)HPC, some strategic decisions should be made:
    - Software re-engineering for rapid technology evolution
    - Stakeholder position in EuroHPC to influence the investments
    - Participate in European Commision IT Strategy on compute, storage, open-data infrastructure...
    - Education and Training