


CMS Usage and Perspectives

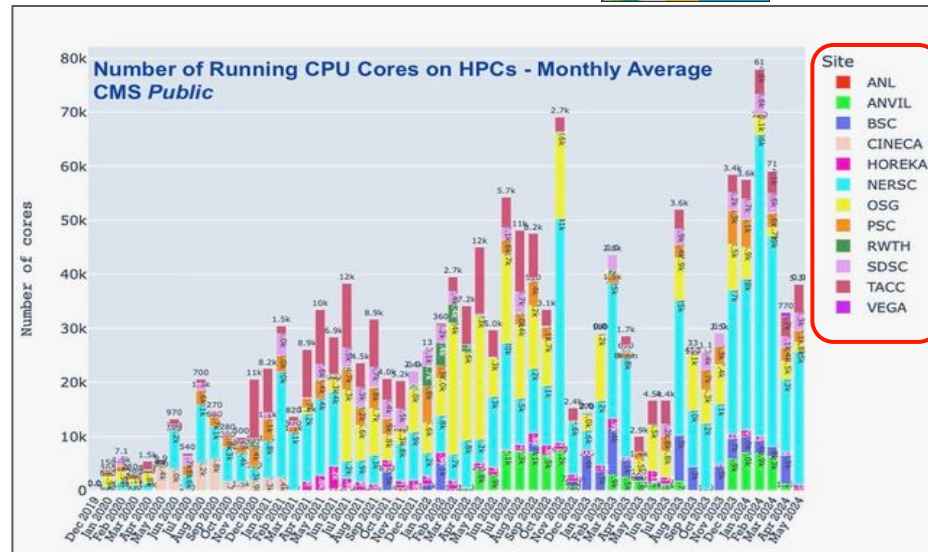
Daniele Spiga -INFN
on behalf of CMS

HEP/HPC Strategy Meeting - All Regions
CERN 30-31 Jan. 2025



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.




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 devlets 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? Reconfiguring + 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	

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



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



HPC Commissioning


Each HPC is/can be different in terms of policies and/or technical setup.

Need individual solutions → **costs**

- CMS has to provide interface between experiment and HPC.
- Ad hoc solutions needed in order to bypass some HPC design features not matching CMS computing mode.

**Distinct strategies developed and deployed.
Today a quite comprehensive portfolio:**

1. **Overlay batch model**
2. **Site Extension**
3. **HTCondor split starter mechanisms for filesystem based communication**



- US HPCs via HEPCloud
- German HPC Resource Integration using COBaID/TARDIS
- Barcelona Supercomputing Center
 - ad hoc solutions for Software; Detector conditions; input output data transfer
- Italy site extension via logical partition of a existing WLCG site Tier1 CNAF



HPC and Software Perspectives

CMS very active and quite advanced

- Investing on GPU support, porting to new architectures (POWER, ARM, Risc-V) or accelerators. HPC have a role representing a valuable “playground” to contribute to the initiatives of the CMS Software group
- **Software - Computing interplay**

Vega GPU-equipped nodes to execute online workflow

- Release Validation workflows of Alpaka-based version of the HLT.



Exploiting also AMD MI250x at LUMI-G HPC

- Alpaka performance portability library studies

Exploiting M100 to enable a multi-arch WM

MARCONI - 100

Nodes: 980


Processors: 2x16 cores IBM POWER9 AC922 at 3.1 GHz

Accelerators: 4 x NVIDIA Volta V100 GPUs, Nvlink 2.0, 16GB


Cores: 32 cores/node

RAM: 256 GB/node

Peak Performance: ~32 PFlop/s



Transnational site extension



HPC and Operations at CMS

CMS target is to make the HPC integration as **transparent as possible** and tries to run almost **all production workflows at HPC**, but often we cherry-pick the "easiest" (least demanding in terms of I/O) types → **costs**


- If we manage to increase HPC usage, is this (still) sustainable?

On the Operational model: status

- Well established for the HPC allocations in the US via HEPCloud
- Working models for some EU HPCs, based on national contacts (i.e. via WLCG site)
- For non-CMS countries that will not provide specific operation support to CMS still under definition

Investing to keep achieving experience with operational model for a long term sustainability

- See Vega, other will come



(Some of the) open points

for later discussions

Integration with CMS computing infrastructure do not represent a real problem so far, **if we can rely on enough network** (and storage)

- Is this still true when we scale up HPC use?

Actively managed storage at HPC

- Streaming is “simple” although add stress on networks and remote CMS site. **Larger storage allocation should be explored to enable a full integration with CMS Data Management**
 - Converging on a small set of approaches would help to maintain HPC integrations

Allocation process and CMS (LHC experiments) planning horizon/process mismatch


- Based on the current allocation process it would need order 1+ years to include HPC resources into CMS planning
- Current HPC use is beyond pledge and thus not an issue. It can be problematic if use increases a lot

Platforms for edge services (cloud/K8s based) at HPC

- It seems to be more popular now. These are instrumental and ease the integration with CMS infrastructure, **further R/D activities should be made**

Access policy (Identity Federation)

- Current approach is “service account based”. A federated model (based on “trust”) can be needed to further integrate HPC



Summary

Continuous effort is spent to integrate HPC resources to increase HPC contributions

- **CMS wants to further increase the HPC exploitation particularly in the EU Zone,** where CMS uses less resources compared to the US

Overall a quite comprehensive portfolio of solutions, **strategically positions CMS toward the exploitation of additional HPC resources**

- Recent success story with VEGA machine in Slovenia
- Looking forward : Lumi? Deucalion?

Further R&D activities needed, possibly guided by open issues.

- Data/storage is one of the main priorities
- Moving from single site integration toward a more “federate system”