

Update on the European HPC Strategy meeting

Maria Girone (CERN)

Tommaso Boccali (INFN Pisa)

Outline

- The process is a bottom-up attempt (yes, experiments are “bottom=users” for HPC centers) for an HPC/HEP strategy
- This talk is mostly a
 - Recap from US meeting (early 2024)
 - A small report from Lubljana EU meeting (Sept 5th+6th)
 - How to go on and some ideas

What is happening?

- It was decided that we (==HEP) would have profited from a more direct link with the HPC world
- At the **political** (==funders) and **technical** (==sysadmin) levels
- For **current** systems (==how to get access to, and how to implement our needs) and **future** systems (==how to design future systems)

Granting process, long term support, technical support, ...

How to (best) execute our workflows on system X?

Explain our (future needs) and try to negotiate systems able to satisfy them

How to implement this?

- Small groups meeting: no formal agenda (not a conference), as open as possible discussion
 - Which is risky, it needs to find persons in the proper mood – on both sides
 - ~ 1 person per Experiment
 - ~ 1 person per HPC center
 - + CERN and WLCG
 - Going more technical or more political: depends on the people, last second steering of discussion
- Know each other:
 - Who we are, what we do, what we (will) need
 - What we can gain from collaboration

US meeting – a reminder

(thanks to M. Girone,
report @ JENA)

- Held in in January 2024; present in 13
 - CMS, ATLAS, WLCG, CERN, Argonne, Berkeley, SDSC
- Discussions focused on
 - Provisioning
 - Interfaces
 - Software
 - Data Management
 - Workflows
- There is already a report available



Strategy For HPC Integration (US) Whitepaper

Authors: Debbie Bard, Doug Benjamin, Paolo Calafiura, Simone Campana, Taylor Childers, Ian Fisk, Maria Girone, Oliver Gutsche, Dirk Hufnagel, Alexei Klimentov, Eric Lançon, Verena Martinez Outschoorn, Frank Wuerthwein

Indico: <https://indico.cern.ch/event/1356688/>

Provisioning

Interfaces

Topics:

Topics:

- Adaptive
- Proposal-
- Resource

Takeaways

The software and the application development will be a critical issue

There are not show stopping technical or political issues preventing more active participation by HPC centers

There are many common issues

A common document and strategy would be beneficial

Data

Topics:

- Scalab
- Data C
- Data M

- Task Splitting for Partial Reconstructions
- Workflow Orchestration

Europe September 5th–6th Ljubljana (agenda)

- Present:
 - CMS, ALICE, ATLAS, LHCb, WLCG, CERN
 - VEGA, LUMI, CSCS, GENCI/SURJF, Juelich, Deucalion, Karolina, Flatiron for US, JENA HPC working group
- Agenda only includes
 - “know each other” – my center, my experiment
 - Open technical discussion on “items”
- Wrap up of the 1st day discussion written during the night, and further elaborated on the second day
- Report in preparation

+ remote attendance



• Items:

- Allocation policies and opportunities
- Edge nodes and AAI
- Type of workflows
- Networking
- SW
- Storage handling
- Strategy

Some points from the summaries

Allocation policies and opportunities

Review of EuroHPC allocation policies

- Nearly all systems (but CSCS) are 50% owned by EuroHPC, 50% by the hosting state; hence two paths to obtain grants
- In both cases, at least 3 types of calls
 - **R&D:** show that your code can be executed on the machine (possibly after some negotiations / changes with local admins)
 - *Only technical evaluation*
 - **Benchmarking:** show that your code can scale 1 → N nodes
 - *Only technical evaluation*
 - **Regular Access:** production level, after having shown the readiness of the codes (via the previous two); tens-hundred millions CoreH
 - *Technical + scientific evaluation*
 - **(Strategic access:** top-down, decided by EC or the EuroHPC governing board; up to 10% of the full EuroHPC capacity)
 - *Political evaluation (+ do we want it?)*

We can go one by one, or try to have an umbrella request handled for example by CERN; suggestion is to go the “national path” if there are links.
Fast proposals, 100% success rate if “reasonable”

Previously via PRACE; now handled completely within EuroHPC. Needs a full proposal, and passing scientific evaluation.
Success rate well below 100%

Less known; up to now “strategic access status” granted only to Destination Earth (climate) and Human Brain Project (brain simulation)

Edge nodes and AAI

- Possibility to deploy edge services (no root access needed as an assumption) from the user point of view (VM, container, ...)
 - Ask for the functionality, not the virtualization technology!
- Have a look at IRI proposals (CHEP 2023) – if not code, as directions (blueprint)
- Is a federated AAI possible (for example, eduGAIN profiles or similar) or even existing?
 - Expected by 2024 – in principle should be adopted by all HPC sites
 - Already now, LUMI supports eduGAIN
- Possible uses of edge nodes: storage, bridging between Condor/Dirac/Panda and SLURM, routing, SQUIDs, ...

Type of workflows

- Production only jobs vs also user jobs
 - Fewer users to consider / enable + more tested workflows
- Which is the lowest GPU utilization you tolerate?
 - One answer was as low as 10%
 - Focus on minimum speedup requested? (looking at the cost)

Networking

- LAN not a problem: these sites exceed what we need by factors (current standard is > 200 Gbps per node; AI specialized machines can exceed 1 Tbps per node)
- WAN: more a political problem than anything else; situation very diverse
 - LUMI: all is open (outgoing connections from compute nodes)
 - Juelich: all is closed
- Possible intermediate solutions
 - Networking opened towards some trusted subnets (solution at CINECA)
 - Allow to open (user level) VPNs
- Upcoming EuroHyperConn strategy: how to connect EuroHPC centers between them, and with everyone else

Proxy and fan-out....

Software

- Is CVMFS usage increasing outside HEP so that we can assume it will be available sooner or later?
 - Alternatively: is userland CVMFS (via aFuse or cvmfsexec) an acceptable option?
 - Alternatively: other solutions? An incremental CVMFS dump?
 - *Is that available? Is that a reasonable feature to ask to CVMFS?*
- Is there a site-preferred high level framework (kokkos vs sycl vs alpaka...)?
 - Not really, for them is an application problem

Storage handling

- Edge services as nodes to deploy data management (for example via services which then use internal centre large POSIX)
- The need to deploy caches: again as edge services? How else?
- Some sites have “datamovers” (globus, Fenix, Simpl, ...)
 - Evaluate enabling them to FTS

Strategy

- (my personal view): the centers consider us as “technical parties” they would like to work with
 - Much better than usual approach from other domains “we send you a tar file please find a way to execute it”
 - But what about the political level?
 - We probably need 2 different types of channels
- A possible strategy (TBC/D!)
 - Form a domain-wide portfolio of successful “Project Access”-like granted projects
 - Use high level opportunities to inform EuroHPC (or above) about us/our needs
 - The annual EC visit to CERN and similar FA-EC meetings
 - The upcoming European Particle Physics Strategy Update; via the WS and the Briefing Book
 - Initiatives like JENA and SPECTRUM
 - EuroHPC User Forum (Maria deputy-chair)

A few possible “actions” from the meeting

(to be validated by the MB)

- Write a report like that written by the US colleagues (but independent at this point)
- Have a global meeting, including US, EU and Asia
 - Now defined: @ CERN Jan 30th–31st (used to be 13th–14th but changed by popular demand)
 - At that point, aim for a global document
- Prepare a DB of existing (successful) grants
- Discuss at the level of WLCG MB a common strategy between experiments towards EuroHPC, in synergy with CERN efforts
 - Divide and conquer? Common projects?
- Continuously monitor new opportunities from the centers (early access, new systems coming online, ...)
 - We have a list from the meeting!
- Prepare a description of main WLCG use cases and inform HPC centers
- Monitor the possibility for a HEP (or more) EuroHPC CoE approval path
 - Submitted in early 2024; not discussed since EuroHPC had a change of plan. Possibly to be reconsidered in early 2025
- Monitor the submission of inputs about HPC utilization to the ESPP