SPECTRUM and JENA

Concezio Bozzi (INFN)*
HEP/HPC Strategy Meeting - All Regions
CERN, January 30th 2025

* Kindly acnowledging: Tommaso Boccali (SPECTRUM), Gonzalo Merino & Johan Messchendorp (JENA WG1)

OUTLINE

- The evolution of HEP computing calls for ever-increasing resources
- Aspects of immediate interest:
  - some Funding Agencies (FAs) expect that a fraction of that comes from HPC systems
  - more (scientific) domains will have sizeable computing requests

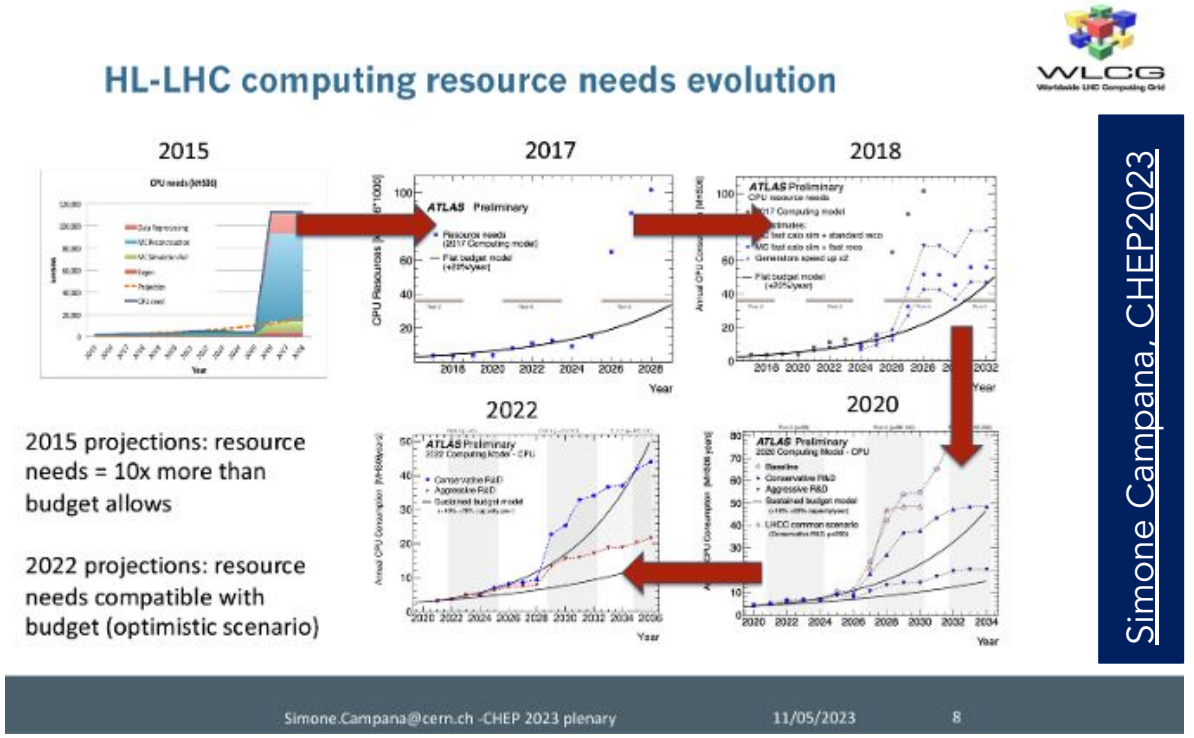

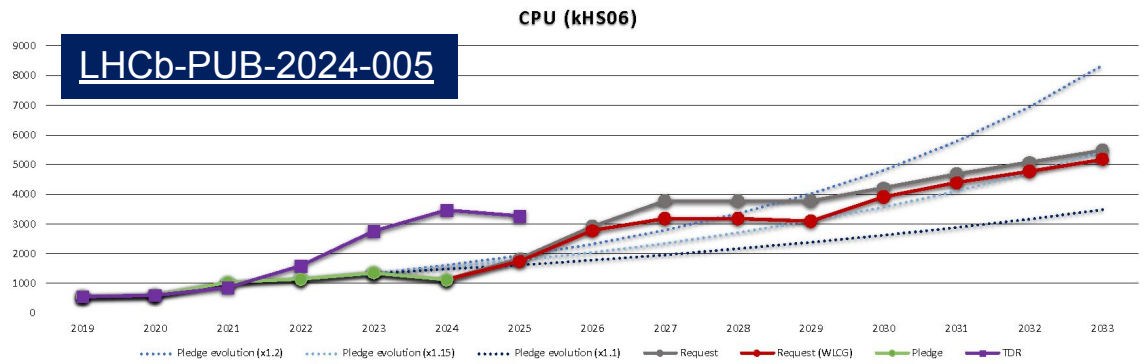

THE HPC SCENARIO(s)

- Already a sizeable part of LHC computing
- Pressure from some FAs to move a part of the processing to HPC (often when the same FA pays for HTC and HPC)
- HPC → more and more accelerators

- Era change from Accelerator Programming (CUDA, ….) to "AI factories" not helping us
-

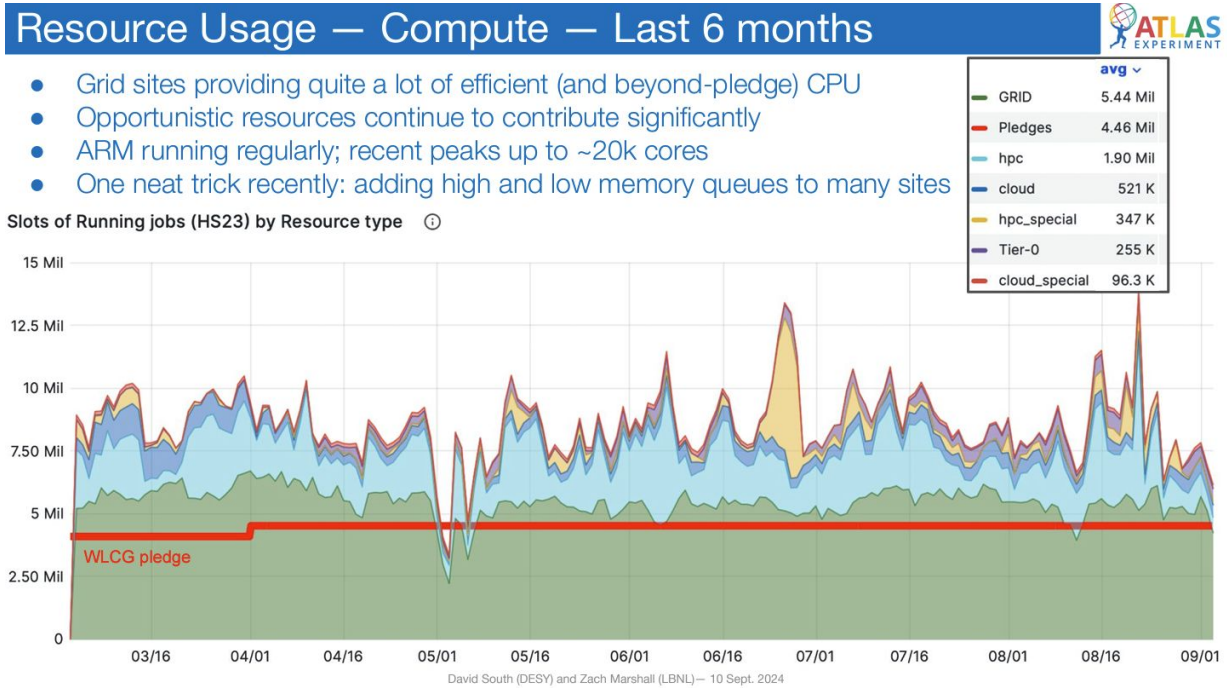

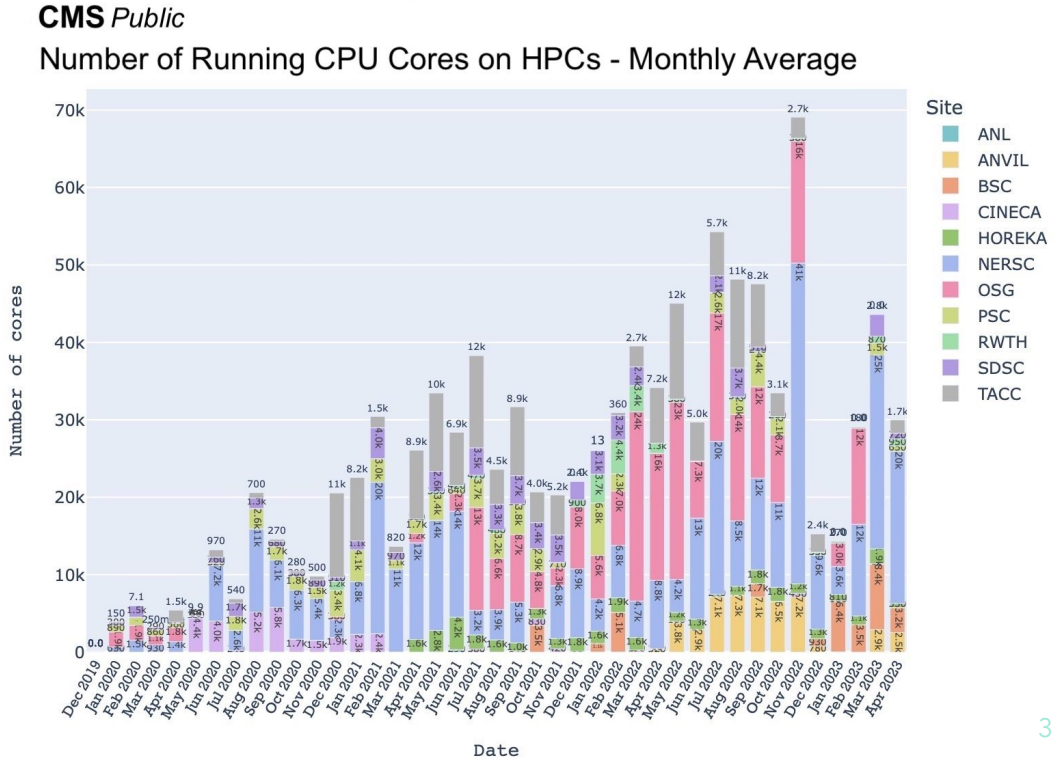

Other (not so different) scientific domains: the RA/SKA example

Delivery to users with SKA Regional Centres

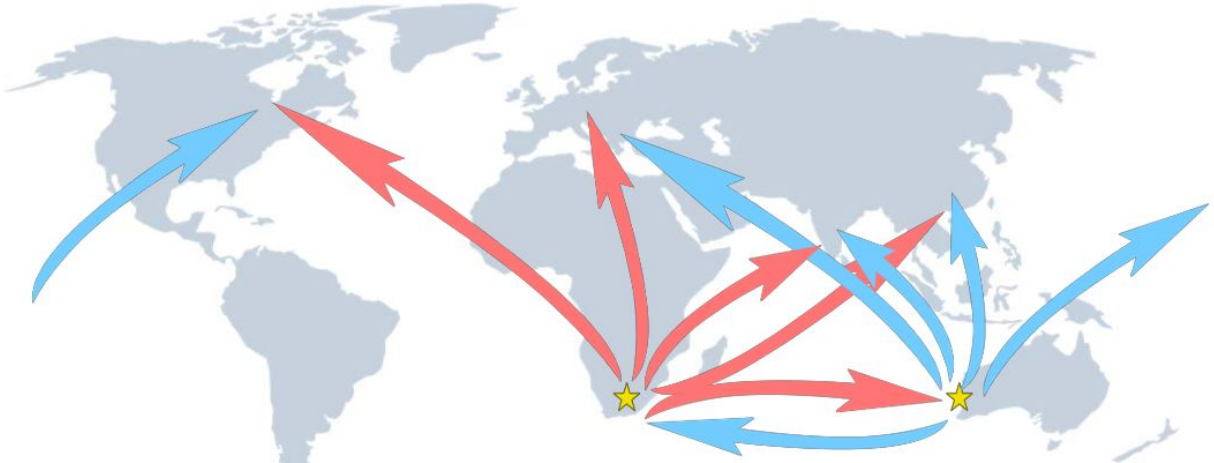

Observatory Data Products flow from the Science Data Processors in Perth and Cape Town to SKA Regional Centres around the globe

2019

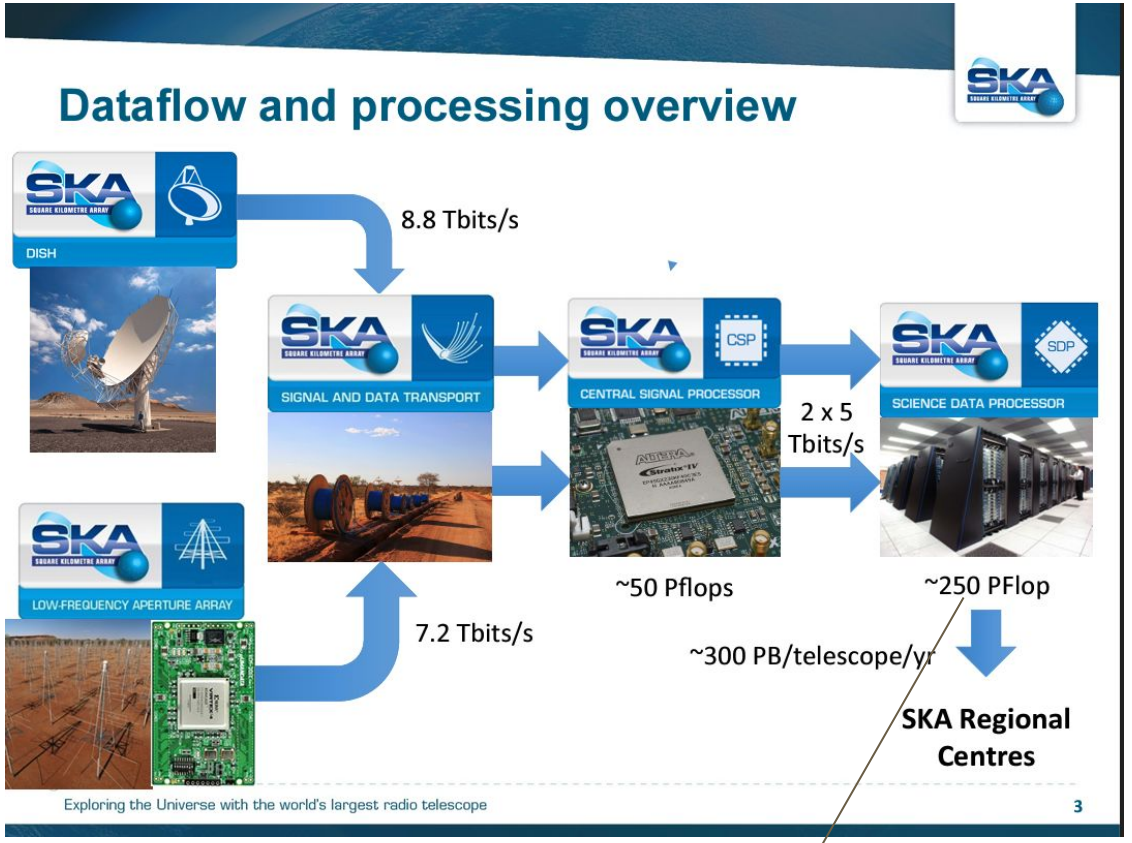

2018

250 PFlops = today's #9 HPC on top500.org

THE SPECTRUM PROJECT

- SPECTRUM is an EU project (link, link) responding to the call: "*Preparation of common strategies for future development of RI technologies and services within broad RI communities … and … Developing, consolidating and optimising the European research infrastructures landscape, maintaining global leadership.*"
- It is NOT a R&D project: the expected outcomes are
  - The preparation of a common strategy between some Scientific Domains with the EuroHPC
  - The realization of a "Community of Practice" of people discussing common items
  - Technically, it is a CSA: Coordination and Support Action: outcome == Word Documents
- Scientific domains of reference: HEP and RA (RadioAstronomy)
  - → HL_LHC + SKA (and precursors)

SPECTRUM

- **Research Infrastructure representatives**
  - LHC: CERN, INFN
  - SKA: CNRS/OCA
  - LOFAR: NWO-I through ASTRON
- **e-Infrastructure representatives**
  - FZJ (HPC Exascale and quantum computing)
  - CINECA (HPC & Quantum)
  - SURF (HTC, HPC, Cloud)
  - (also EGI Foundation, INFN)

SPECTRUM (2)

- Expected (real) outcomes:
  - the project reports to European Commission (EC) and its various branches (including EuroHPC, the EOSC, …)
  - wishlist:
    - convince EC that HEP should be (for example) in the list of EuroHPC golden use cases
    - convince EC to emit calls (2025+) to fund R&D to execute the program
  - learn from others (programming, strategies, solutions, …)

- SPECTRUM is a 30 months project, in which:
  - months [1,15]: bottom-up phase. Probing the community, gathering documents, forming discussion fora, getting the “facts”
    - we formed a “Community of Practice”
  - months [16,30]: elaborate the gathered info, write a Blueprint and a Strategic Research, Innovation and Deployment Agenda (SRIDA)

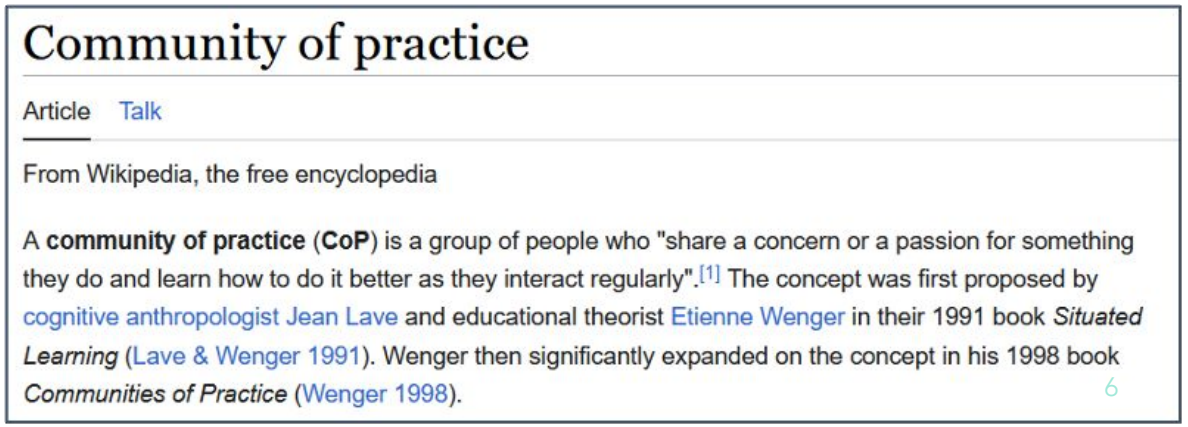

Community of practice

Article Talk

From Wikipedia, the free encyclopedia

A **community of practice** (**CoP**) is a group of people who "share a concern or a passion for something they do and learn how to do it better as they interact regularly".[1] The concept was first proposed by cognitive anthropologist Jean Lave and educational theorist Etienne Wenger in their 1991 book *Situated Learning* (Lave & Wenger 1991). Wenger then significantly expanded on the concept in his 1998 book *Communities of Practice* (Wenger 1998).

The Community of Practice, the Working Groups, the Survey

- Six (6) Working groups defined
- Call for participation issued in March
  → O(80) persons
- Activities:
  - ~30 meetings so far (level of attendance medium, but we are a busy community)
  - task before the summer → prepare a survey
  - task now: analyze the survey results
  - task later: continue the info retrieval and reciprocal understand via presenting use cases

WG1: Data Management and Access

Chair: BAGNASCO Stefano (INFN, ET, Virgo)

- Data Management
- Data Access Protocols
- Data Archiving
- Security

WG2: Workflow management and organization

Chair: DELL'AGNELLO Luca (INFN CNAF-T1)

- Resource Discovery and Workflow Submission
- Resource Allocation
- Complex Workflows

WG3: Compute Environment

Chair: BOZZI Concezio (INFN, LHCb, JENA-HPC)

- Expected Tools and Services
- Facility Expectations
- Edge Services
- Library Provisioning

WG4: SW tools

Chair: SWINBANK John (Astron)

- Machine Learning Frameworks
- Multithreading Frameworks
- Multi-Node Tools
- Compilers, toolchains, ...
- Quantum computing tools and frameworks
- Code Management Practices

WG5: Scientific Use cases

Chairs: FERRARI Chiara (CNRS, OCA), GIRONE Maria (CERN)

- Typical Use Cases
- Requirements and Needs
- Best Practices Collection
- Data Fluxes and Paths

WG6: Facilities

Chair: HOPPE Hans-Christian (Jülich)

- HPC Centers
- Access to Quantum Computing Hardware
- Access to Commercial and Public Clouds
- Sustainability
- Security

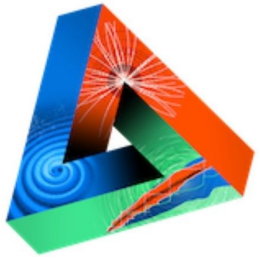

JENAA

Joint ECFA-NuPECC-APPEC Activities

ECFA

European Committee for Future Accelerators

European Committee for Future Accelerators

NuPECC

Nuclear Physics European Collaboration Committee

Astro Particle Physics European Consortium

APPEC

THE JENA(A) PROCESS

- **JENA(A)**: Joint **ECFA**,**NuPECC**, **APPEC** (Activities)
- During the 2a JENA Symposium (Madrid, 2022), it was decided to give a particular focus to (shared) computing perspectives in the three communities (strongly desired by the FAs)
- A specific Workshop was organized (Bologna, 2023) to discuss computing for future initiatives in the three domains (with an EU focus but not only)
- Scope is larger than SPECTRUM (for example, the inclusion of NP; or a specific focus on AI)
- Decisions:
  - organize Working Groups to discuss inter domains; eventually submit a survey
  - prepare a white paper / document on the timescale of 2024, for the next JENA Symposium (RAL, 8-10 April 2025)

JENA Computing Initiative: WG1 (HPC)

JENA Computing Initiative: WG2 (Software)

JENA Computing Initiative: WG3 (Data)

JENA Computing Initiative: WG4 (AI)

JENA Computing Initiative: WG5 (TDE)

JENA Motivations

- Nuclear physics and astroparticle are increasing their data rates and definitely enter an era of “big science”
  - FAIR experiments, CTA, KM3NeT, …
- All the communities face very similar issues
  - How to use HPCs practically and effectively
  - How to have software that runs well on HPCs and other facilities
    - Which means accelerator programming
    - Understand how we use AI in our workflows and what infrastructures we need for them
  - Common data handling layers that can be shared (promoted by ESCAPE)
  - How to train software and computing experts and give a viable career path

- Process of communities involvement in the WGs and community endorsement for the conclusions - draft document being circulated
- Outcome: send a clear message to the funding agencies about ***common needs*** **and** *how to fund them*
  - It helps a lot when we can demonstrate that we can work together coherently (when it makes sense, which is not the same as one-size-fits all!)

SPECTRUM + JENA

- submitting one survey is annoying enough → let's merge them
- Also the timescale is surprisingly aligned:
  - SPECTRUM: internal doc end-of-october, deliverable March 2025
  - JENA: white paper by end-of-2024
- and on top of all: CERN called for the European Particle Physics Strategy Update 2024 – 2026 (here) (process similar to Snowmass in the US)
  - It happens every 5-10 years; It was initially scheduled in >2026. Now to end in early 2026.
  - Will accept community contributions by March 2025 → same time scale as SPECTRUM, and also a target for JENA
  - We think it is a good opportunity to show that our environment (the enlarged domain) is cohesive and able to propose common solutions to (quasi) common problems
- Solution
  - intertwine SPECTRUM and JENA: common survey, presence in each other WGs

The SPECTRUM Survey

- Online since Jul 18th, as “SPECTRUM and JENA”
- Official closing Sept 30th .. but in practice we decided not to close to have the possibility for a top-up later in time
- A complex survey: do not ask everything to everyone. Ask for “your roles” and switch on/off sections accordingly

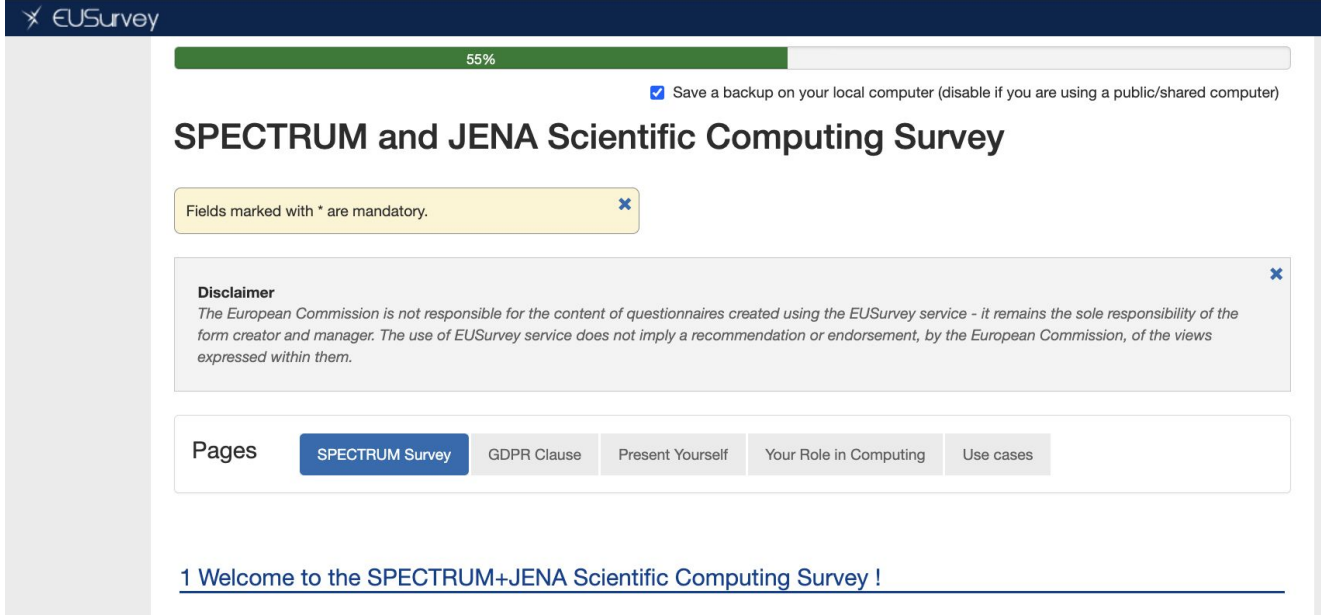

Authentication and Authorization

Processing Needs

Data Management Needs

Training and careers

Expected Computing Environment

Software development, distribution, policies

e-Infrastructures

80 answers before the official closing, a few more afterwards…

The Survey (2)

- 80+ answers so far

Which are the categories which better describe your role(s)?

		Answers
Researcher / user of scientific computing resources (doing analysis, R&D, operations, ...)		65
Manager of a scientific initiative (for example, an experiment, an instrument, an observatory)		15
Manager of an e-Infractructure (for example a computing centre, a storage facility, a distributed computing facility)		19
Research software engineer (writing, testing and managing code for an initiative)		20
Other (please specify)		4
No Answer		0

Which is your role?

Which is/are your scientific domain(s) of expertise (if applicable)?

		Answers
Experimental High Energy Physics (HEP)		31
Theoretical High Energy Physics (HEP)		18
Observational Radio Astronomy (RA)		8
Experimental Gravitational Waves (GW)		6
Theoretical Gravitational Waves (GW)		3
Experimental Nuclear Physics (NP)		9
Theoretical Nuclear Physics (NP)		12
Observational Astroparticle (not RA or GW)		9
Other physics related domains (please specify below)		6
Other non-physics related research domains (please specify below)		4

Which is your domain?

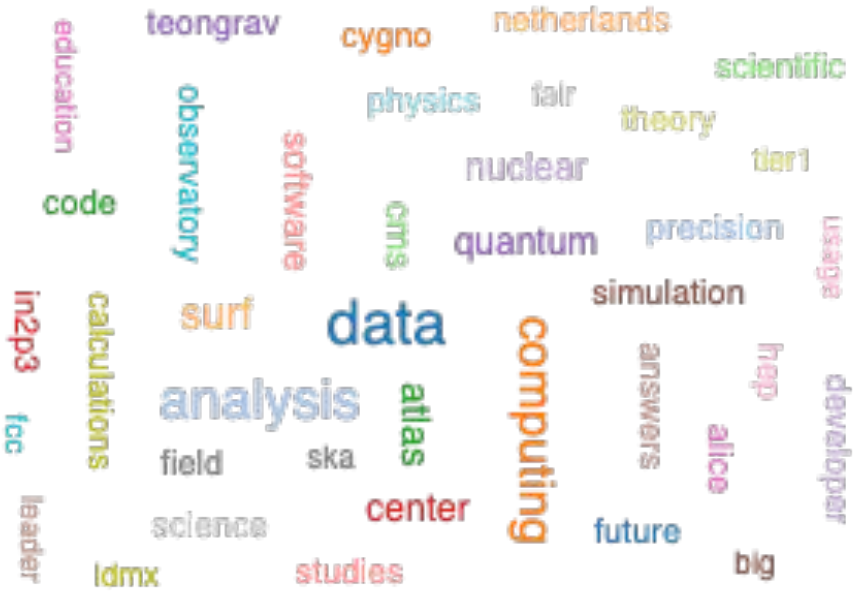

most present words

Please select your areas of expertise, for which you can answer technical questions:

	Answers
Authentication and Authorization	11
Workload Management	20
Data Management	21
Computing Environment	25
Software, software development and software management	52
Training and careers	19
No Answer	14

Which are the main mechanisms to get access to your infrastructure?

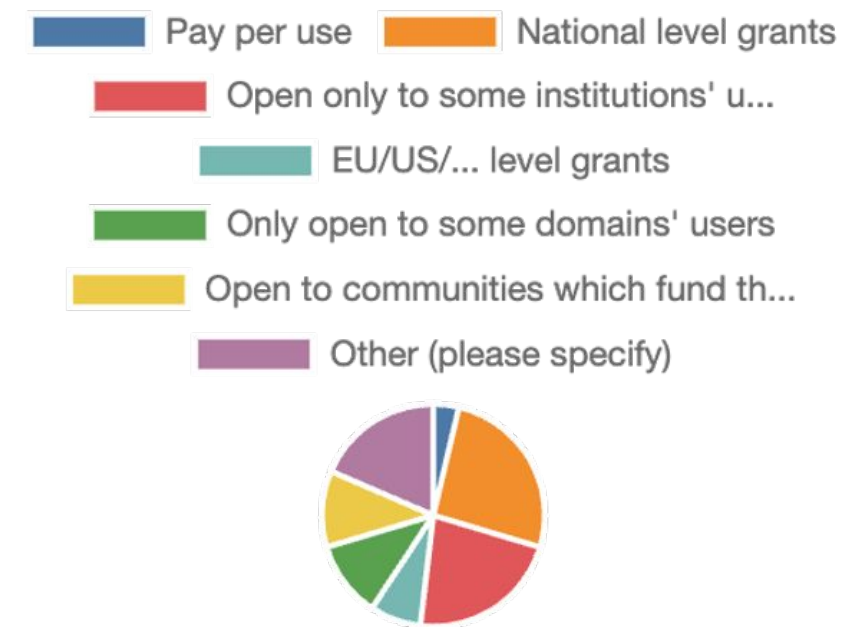

Which type(s) of resources does your e-Infrastructure provide today?

	Answers	Ratio
HPC (CPU only)	5	7.14 %
HPC (CPU and GPU)	9	12.86 %
Cloud (CPU only)	3	4.29 %
Cloud (CPU and GPU)	6	8.57 %
Quantum Emulators (classical systems emulating quantum hardware)	4	5.71 %
Quantum Hardware	0	0.00 %
Other (FPGAs, TPUs, ...) - please specify	2	2.86 %
HTC CPU only (Grid-style, batch, ...)	6	8.57 %
HTC CPU+GPU (Grid-style, batch, ...)	8	11.43 %
Storage systems (disks)	13	18.57 %
Storage systems (tapes)	10	14.29 %

What is the driving force behind the architecture / accelerators decisions?

	Answers	Ratio
Cost/performance optimization	21	30.00 %
Capability to execute of more centres	8	11.43 %
Need for more performance than simple CPUs could provide	12	17.14 %
Previous experience in the developers' base	7	10.00 %
External decision (funders, reviewers , ..)	4	5.71 %
Decision coming from resources deployed in the centre we need to use	8	11.43 %
Other	1	1.43 %

A few emerging messages (more details in backup)

Technical areas were identified where work will be needed to enable the effective integration of HPC centers and experimental facilities:

- *edge services,*
- *federated access and AAI,*
- *workflow management,*
- *wide-area networking,*
- *data management,*
- *software deployment,*
- *programming models.*

Organisational areas were identified as well, in particular:

- *Resource access* and
- *training and talent retention*

Preliminary recommendations of JENA WG1

- A **coordinated effort** is **needed** to submit **access requests** to EuroHPC and other regional/national/international initiatives from the ENA communities to **develop a coherent roadmap** for efficiently utilizing large HPC systems.
  - establish a roadmap for addressing technical gaps
  - to provide a consistent picture of global needs in preparation for e.g. an EuroHPC Strategic Access application.
  - Ultimate goal: be recognized as a strategic activity, enabling multi-year allocations of compute and storage at HPC facilities, incorporate them into the long-term capacity planning of research groups and collaborations.
- Securing adequate **funding** for **software development** to effectively integrate HPC resources into research workflows
  - optimize scientific codes for new architectures
  - develop workflow and data orchestration software
- Establishing of **monitoring** mechanisms to regularly **assess** the evolution of the **CPU and GPU needs** of the scientific collaborations and their impact on the computing models and on the distributed computing infrastructure.
  - Incorporate the results of these assessments into the roadmap for integrating HPC resources.

Preliminary recommendations of JENA WG1 (cont.)

- Collaboration mechanisms, both at the technical and political level, should be identified in order to **ensure that the future global scientific computing landscape builds on the strengths of HPC and HTC systems**
  - Evolution of EuroHPC JU towards a more interconnected and data-centric architecture (Federation Platform, AI Factories)
  - Leverage valuable experience gained in e.g. WLCG in distributed computing and data analysis.
  - Ensure the needs of ENA communities are considered in the design and procurement of HPC systems at regional, national, European levels

- Sustained **investment in JENA computing and data infrastructures** tailored to the **specific scientific requirements** of these research communities **remains essential**.
  - Risk that priorities of HPC systems might not necessarily align
  - Datasets are unique asset

CONCLUSIONS

- Many initiatives trying to tackle the HPC (and not only) utilization in HEP - here a view mostly on EU side
  - at moments the effort seems a bit chaotic, but (as a gas) it is evidently generating some pressure on the “decision levels”
  - this is vital part of a general roadmap for scientific computing, which the FAs push for
- “common/global” approach to the interactions with HPC centers worldwide should be pursued
  - Well-identified issues to be tackled in technical and organisational areas

Backup: technical areas

- *Edge services* should be exploited to address some typical HPC issues, such as: the capability to receive external workloads (and bridge them to the internal batch system), and to insert (some parts of) the center storage into a global data management systems; bridging external AAIs with the specific users at a center; serving as proxies for services like CVMFS and direct data management; routing a part of internal traffic towards the general internet (via NAT and routing mechanism).
- *Federated Access, AAI:* the use of “special user”, or possibly service accounts, in conjunction with edge services, would facilitate handling complex services such as the execution of generic user jobs and the enabling of writing on the filesystems by data management systems.

Workflow management:

- workflows, especially the experimental HEP ones, are not the typical ones used at HPC centers, due to: very complex code bases and dependencies; late binding; the need for external services not necessarily known before run time; varying needs in terms of RAM, I/O, external connection and relative bandwidth.
- “Production workflows”, i.e. those submitted at large scale, after thorough testing for undesired behaviour, by a limited number of users (typically computing operators for the experiments), are generally preferred to “chaotic” workflows, i.e. those submitted by any users with limited scrutiny of their behaviour.
- Some variety in production workflows in terms of memory, I/O, network, CPU intensity might drive the ability to direct them to specific HPC centers that have compatible infrastructure and services.
- JENA workflows are also making a limited usage of GPUs. This does not guarantee the saturation of the hardware capabilities of HPC centers, which might refrain access to their resources. Some threshold of acceptance should be defined by e.g. looking at the hardware cost.

Backup: technical areas (2)

- *Global networking:* workflows typically need the capability to open outgoing connections from the compute node to the global internet. There are technical difficulties for HPC sites to allow for this (infiniband-to-ethernet conversions at the Tbps level and bandwidth towards general internet and NRENs), but also security aspects (HPC centers are sensitive targets against DDOSs, hackers breaking attempts, and so on), hence strict access policies and robust firewall rules. A range of solutions have proven to be workable, with different viabilities for HEP workflows, ranging from NATted connections using a few edge services (optionally limited to a list of accepted subnets), to VPNs (handled by the center or a user level), to proxies deployed service by service.
- *Software deployment:* CVMFS usage is increasing in various scientific domains and it can therefore be thought of as a possible solution for HPC centers.
- *Data management:* edge services, offering POSIX access/write on a shared file system, would be sufficient to integrate access via remote protocols (WebDAV, XrootD) and data management via higher-level tools that maintain catalogs of the files and their physical replicas, and perform data transfers and deletion. An integration of the existing HPC datamovers with services such as FTS could be the fastest and least intrusive way to connect the centers with central data management systems.
- *Programming models:* the gap between workflows and accelerated hardware provided by the HPC centers should be addressed by high-level frameworks (Kokkos, Alpaka, SYCL,...). HPC sites regard this as an application issue. Support in training of people / testing/ optimization would be beneficial.

Backup: organizational areas

- *Resource access:* the lack of long-term HPC resource allocations affects the reliability of critical workflows. A solution could involve collaboration with funding agencies and HPC providers to explore more stable allocation models.
- *Training and retaining talent:* Attracting and retaining young software developers skilled in accelerator programming and portable software design is a major challenge in today's competitive market. These developers are crucial for improving scientific software in a sustainable way, and focused effort is needed to address this issue.