

Distributed Computing: An Insight Into (our) Current Developments and Future Challenges

alternatively...

How we are try to connects the dots!

Or

Perhaps we're just spoiling SOSC2025?

Lucio Anderlini
Diego Ciangottini and
Daniele Spiga

SOSC 2024 Sixth International School on Open Science Cloud

(Some of) The dots

Distributed computing

- The use numerous computing resources in different operating locations for a “single” computing purpose.



Cloud and cloud-native, aka architectures are built using a series of elements such us:

- Microservices: Applications are broken down into smaller, independent services.
- Containerization: Uses technologies like Docker for packaging and managing applications.
- DevOps, Automation: automated testing, deployment, and operations.
- Orchestration i.e. k8s.

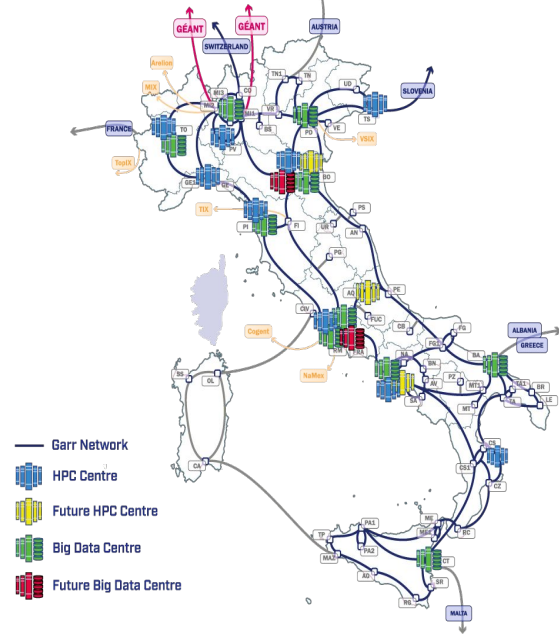
Machine Learning and AI

Specialized Hardware (GPU, FPGA...) exploitation

Distributed Resources In Italy a recap!

0

SUPERCOMPUTING CLOUD INFRASTRUCTURE



Finanziato dall'Unione europea NextGenerationEU | Ministero dell'Università e della Ricerca | Italiadomani

At the Bologna Technopole

The INFN Tier1
Inaugurated May 10, 2024

CINECA n. 9 Top500 supercomputer in the world

Leonardo

ICSC Italian Research Center on High-Performance Computing, Big Data and Quantum Computing

Missione 4 • Istruzione e Ricerca

Press Release | 19 September 2024 | European High-Performance Computing Joint Undertaking | 2 min read


New Procurement Call to Upgrade LEONARDO, the EuroHPC supercomputer located in Italy

The European High Performance Computing Joint Undertaking (EuroHPC JU) has launched a new procurement call for the acquisition, delivery, installation and maintenance of the hardware and software of LISA, the upgrade to LEONARDO Supercomputer.

How we envision to deal with this heterogeneity

Seamlessly use the most suitable (piece of) computing hardware for each given task within a pipeline

- Not really trivial for end user actually. We don't even expect users should need to learn how to do it.
 - So what, **the underlying system can do it**

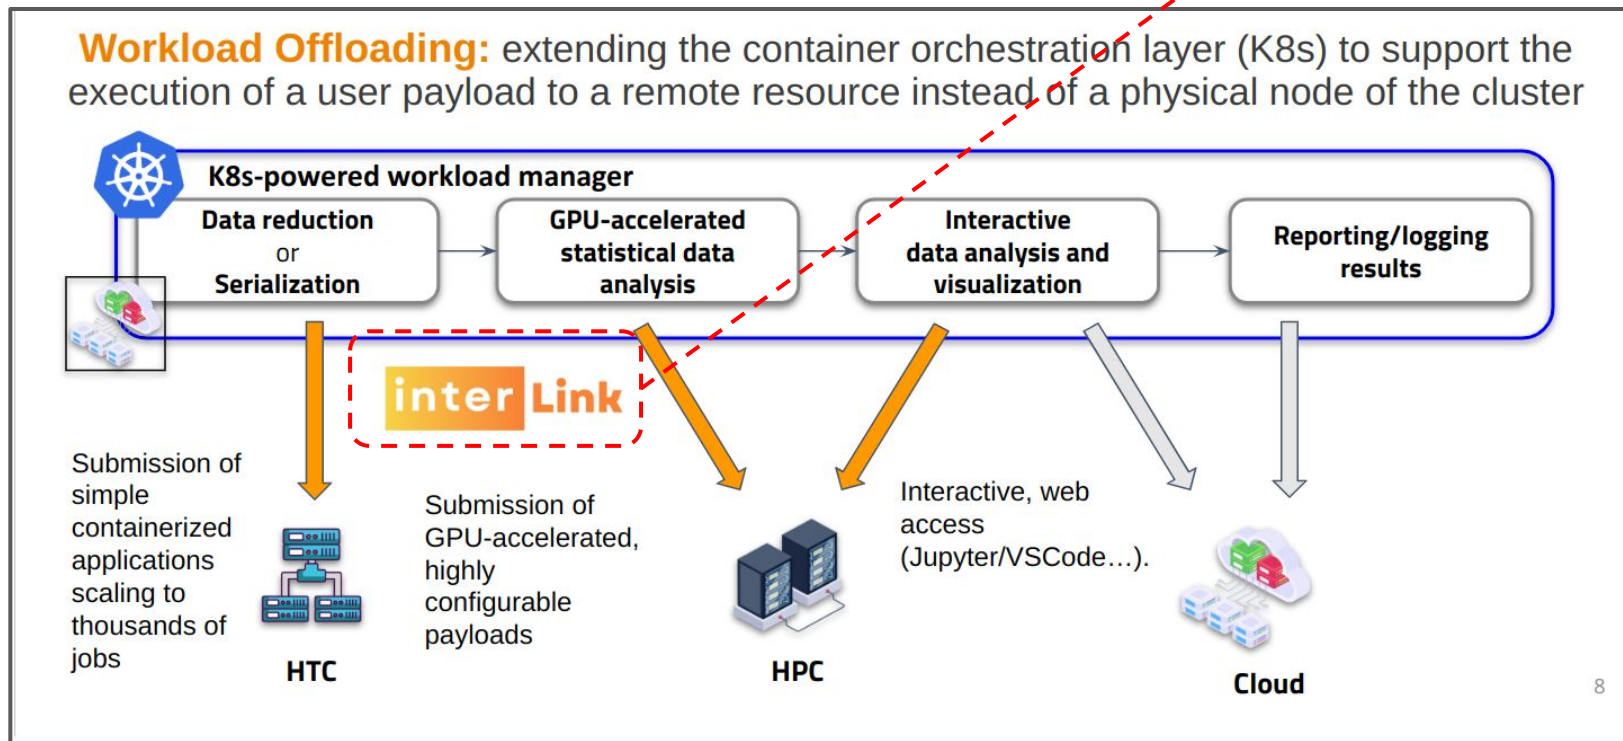


Adopt industry standards and just develop the needed glue in order to

- **To manage payload** or to prepare task to execute
- **To provision computing capacity**
- **To interface services and providers:** Aims to have unique interface to many execution hosts (API)

What are we exploring right now

Example of developed glue




Under the hood

Few keywords

- Container
- Orchestration
- Rest API

Remember Tuesday??



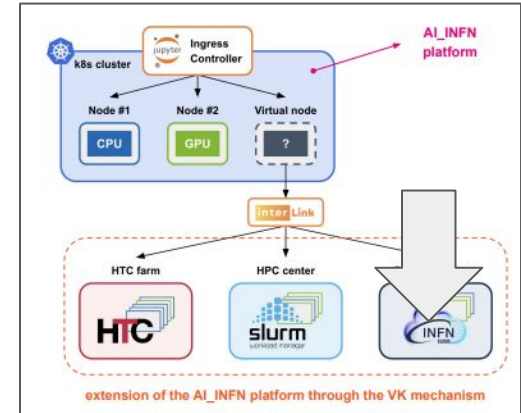
AI_INFN example: Artificial Intelligence Technologies for INFN research



AI_INFN is an initiative funded by INFN CSN5 to promote the adoption of machine learning and artificial intelligence techniques for fundamental scientific research



- ❑ Facilitating access to HPC and GPU resources
- ❑ Organizing educational and training events
- ❑ Fostering the AI community within INFN
- ❑ Conducting R&D to integrate new technologies (e.g., FPGA...) into the platform



Once AI models are developed, researchers often seek to **scale them beyond development-dedicated resources**

- The AI_INFN platform is exploring a solution to transparently extend the resource pool accessible to Kueue using the Virtual Kubelet (VK) mechanism

How to choose where to run a job

The system detects
(from user requirements)

Requires GPU?	Offloadable?	<i>First choice</i>	<i>Fallback</i>
No	No	Local resources	
No	Yes	CNAF Tier-1	Local resources
Yes	No	Local resources	
Yes	Yes	CINECA Leonardo	Local resources

User define

Managed by the system,
transparently to the user