Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

ICSC
Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Artificial Intelligence in the ICSC and Terabit projects

Lucio Anderlini – INFN Firenze, ICSC

Digital Twins for Nuclear and Particle physics - NPTwins 2024

Museo Diocesano di Genova – 16-18 December 2024

INFN FIRENZE

The National Center for HPC, Big Data and Quantum Computing

- Italy has funded, with NRRP (pandemic recovery) funds, 5 large National Centers, for a total of 1.6 Beur over 3 years
- One of them, coordinated by INFN, focuses on modern IT technologies, with the final goal to deploy a long-term (>>3 years) distributed infrastructure available to the reseach and the industrial ecosystems in Italy
- Sept 2022 to Aug 2025

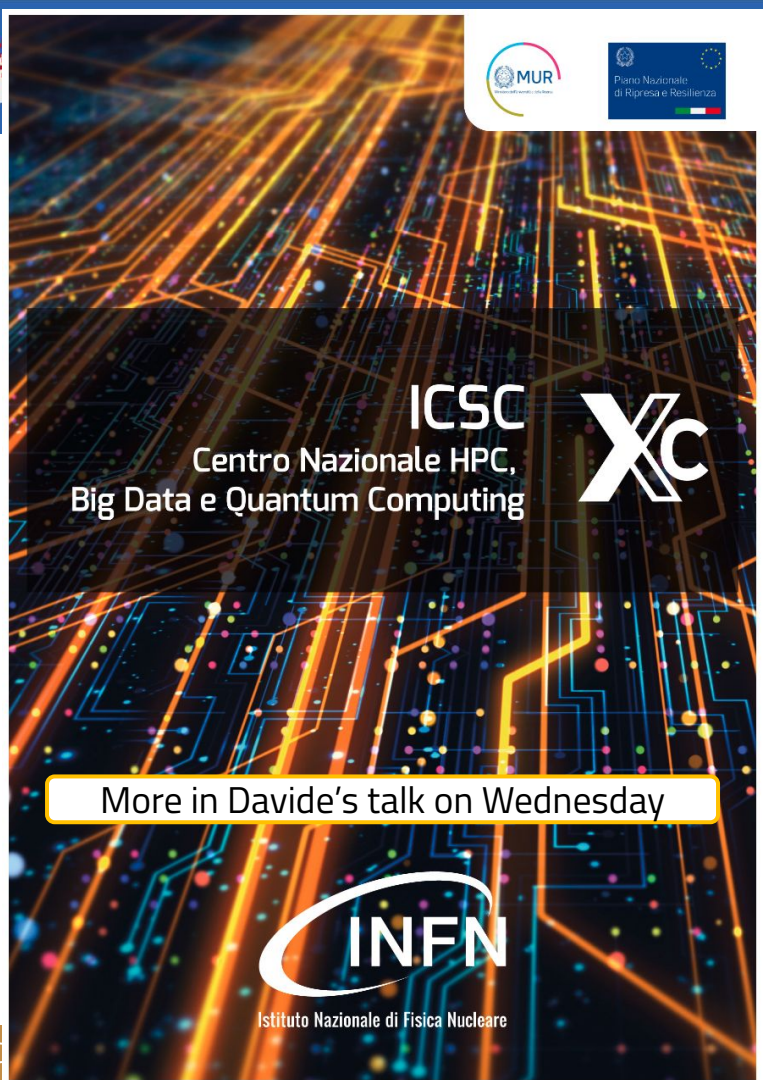

The three pillars of ICSC and Terabit projects

- **GARR**: Tbps-level connectivity between all public data centers

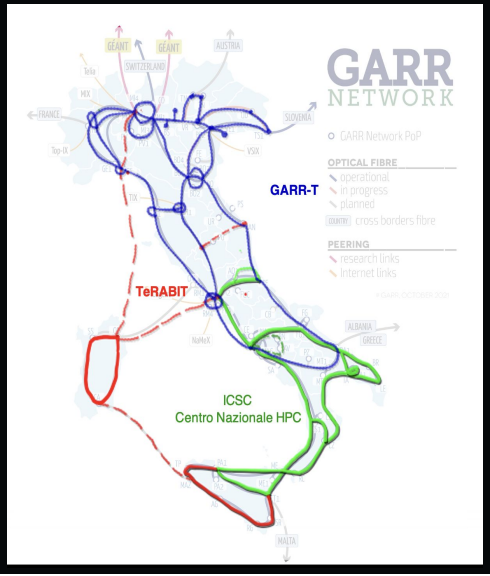

- **CINECA**: expansion of Leonardo (HPC#6 on top500.org) with Lisa, and deployment of a production level Quantum Computer

GPU Module Lisa	CPU Module Lisa
500 compute nodes 4x GPUs / node >100 PF HPL	1000 compute nodes 2x CPUs + HBM / node >6 PF HPL

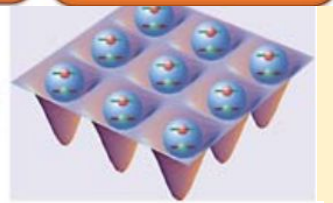

PASQAL IQM AQT

- **INFN**: strengthening of the WLCG infrastructure (1 Tier1 – 9 Tier-2s); acquisition of Cloud resources; implementation of the datalake middleware, based on INFN-Cloud

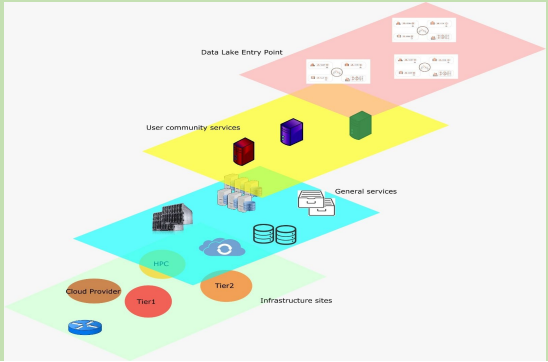

L'ICSC include
10 Spoke tematici
e
1 Spoke infrastruttura

0 SUPERCOMPUTING CLOUD INFRASTRUCTURE

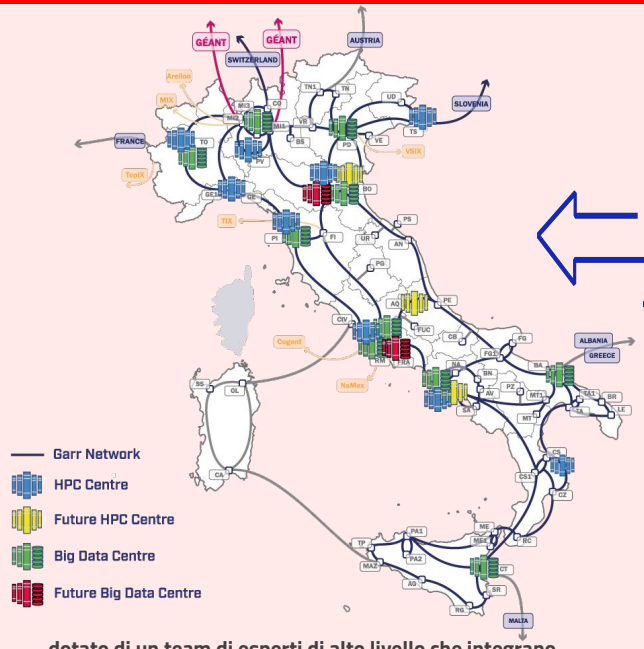

dotato di un team di esperti di alto livello che integrano i gruppi di lavoro degli Spoke (gruppi misti trasversali)

ISTRUZIONE E FORMAZIONE, IMPRENDITORIALITÀ, TRASFERIMENTO DI CONOSCENZE, POLICY, OUTREACH

1 FUTURE HPC & BIG DATA
2 FUNDAMENTAL RESEARCH & SPACE ECONOMY
3 ASTROPHYSICS & COSMOS OBSERVATIONS
4 EARTH & CLIMATE
5 ENVIRONMENT & NATURAL DISASTERS
6 MULTISCALE MODELING & ENGINEERING APPLICATIONS
7 MATERIALS & MOLECULAR SCIENCES
8 IN-SILICO MEDICINE & OMICS DATA
9 DIGITAL SOCIETY & SMART CITIES
10 QUANTUM COMPUTING

Why our communities need an ICSC

- Since at least 2 decades, research at the fundamental frontier is heavily "computing bound"
  - Latest (and next) generation experiments collect data at the Exabyte and beyond
  - Simulation efforts in theoretical and experimental physics are at the Exascale
  - We have been forced to develop in-house solutions when nothing was available, with a good success. But it is now due time to evolve to a shared infrastructure model
    - The Web, the GRID, …
- Examples:
  - LHC has has already surpassed the global scale of several Exabytes of Data, and more than 1 Million CPU cores
  - Lattice QCD simulations are, with Meteo, the main driver and benchmark for HPC systems

The Worldwide LHC Computing GRID

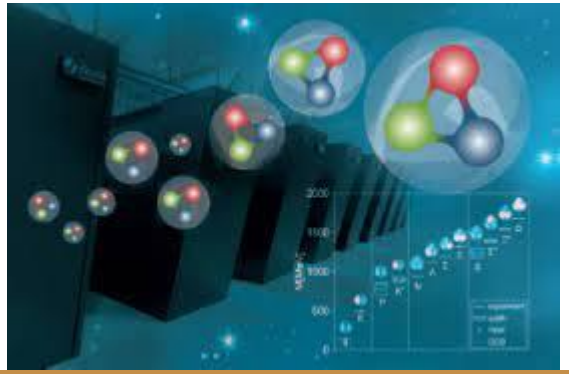

Lattice QCD on HPC

From High Performance Computing to Artificial Intelligence

T. Boccali, 1st Hackathon of Machine Learning of AI_INFN

HPC is designed for Simulation

↓

Artificial Intelligence requires data

- Leonardo (current machine @ CINECA):
  - Not designed for AI, but with standard HPC use cases (lattice QCD, Meteo, computational chemistry, …)
    - Not al lof them can use GPUs → 2 partitions
  - Leonardo Booster:
    - 3456 nodes, each 3xA100 and 32 CPU cores
  - Leonardo GP:
    - 1536 nodes, each 112 CPU cores
  - Fast interconnect between nodes and fast storage
- INFN has a large quota; "ask us!"

Italy: current and next machine for large (as in LLM?) AI

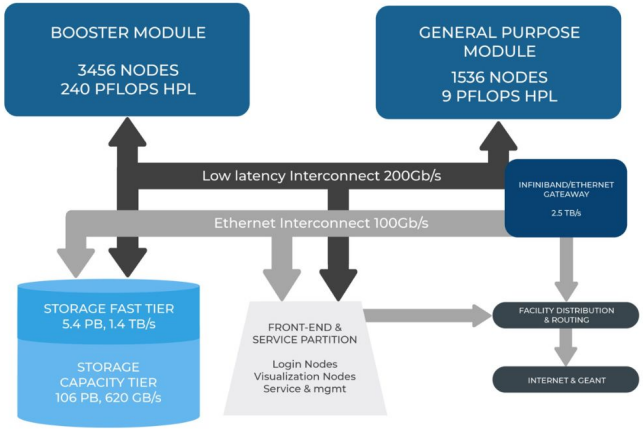

"It happened" it is also a quite good machine for AI; not automatic!
#4, Fugaku, is a terrible AI machine: no GPUs, low RAM

Testbed: what come next (early 2024)

D. Spiga (INFN)
ICSC Annual meeting

Data Management Services: data orchestration and data transfers

Establishing the DataLake testbed:

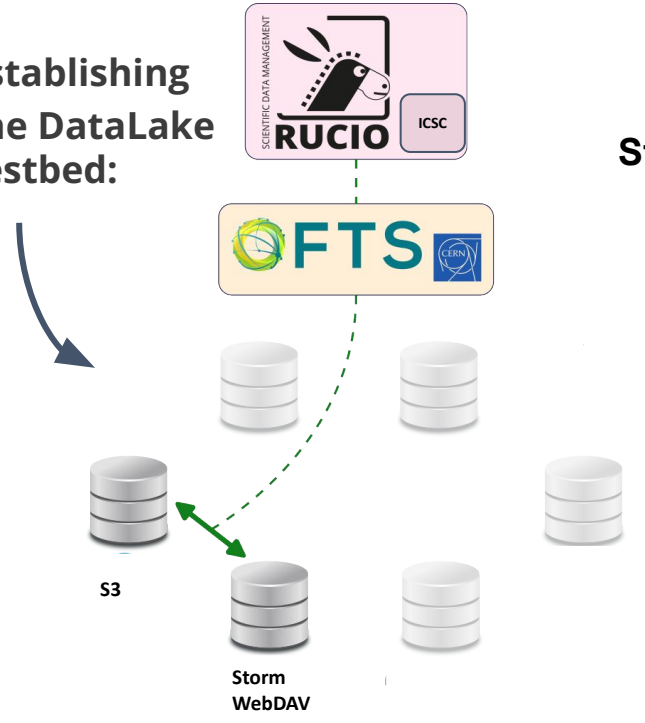

Start integration with science use-cases

- Collect exemplar datasets
- The datasets will be injected in the Datalake building science driven playground

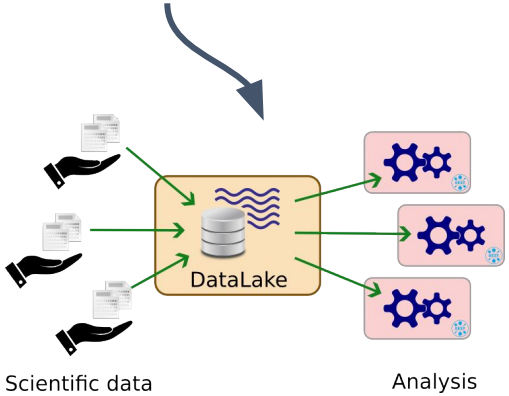

Innovation Grant (Leonardo)

WP3 and WP1 use cases not yet deeply discussed

PRESS RELEASE | 10 December 2024 | European High-Performance Computing Joint Undertaking | 9 min read

Selection of the First Seven AI Factories to Drive Europe's Leadership in AI

The EuroHPC Joint Undertaking (EuroHPC JU) selected the sites that will host the first European AI Factories, set to be deployed next year across Europe: in Finland, Germany, Greece, Italy, Luxembourg, Spain and Sweden.

EuroHPC JU

Just published

FIRST 7 EuroHPC AI Factories will provide supercomputing capabilities, specifically tailored to meet the security, confidentiality and data integrity needs of European **industrial users**.

Project IT4LIA

CINECA hosting entity

- Premise: Tecnopolo di Bologna

420M €, of which 30M € for the AI Factory (co-financing 50% EuroHPC JU – 50% MS)

https://eurohpc-ju.europa.eu/selection-first-seven-ai-factories-drive-europes-leadership-ai-2024-12-10_en

Slide from Alessia Dorazio, ICSC Status report, Catania, 2024-12-10

From Distributed to Heterogeneous computing

Heterogeneity among computing sites:
different hardware, operating systems, batch systems, data infrastructure, networking

The name of the game is build the ability to transparently move applications from site to site.

- Uniform AuthN/Z;
- non-locality of data (DataLake), with local caches
- Containerization and distribution of the applications

Lingua franca #1

Kubernetes for containerized applications

Lingua franca #2

Singularity/Apptainer as runtime in user space

D. Spiga (INFN)
ICSC Annual meeting

Further activities in the pipeline

Transparently extend analysis testbed to run "*any* **application** *anywhere*"

- **To federate (highly) heterogeneous and disparate ICSC providers**
  - **enabling a "transparent payload offloading"**

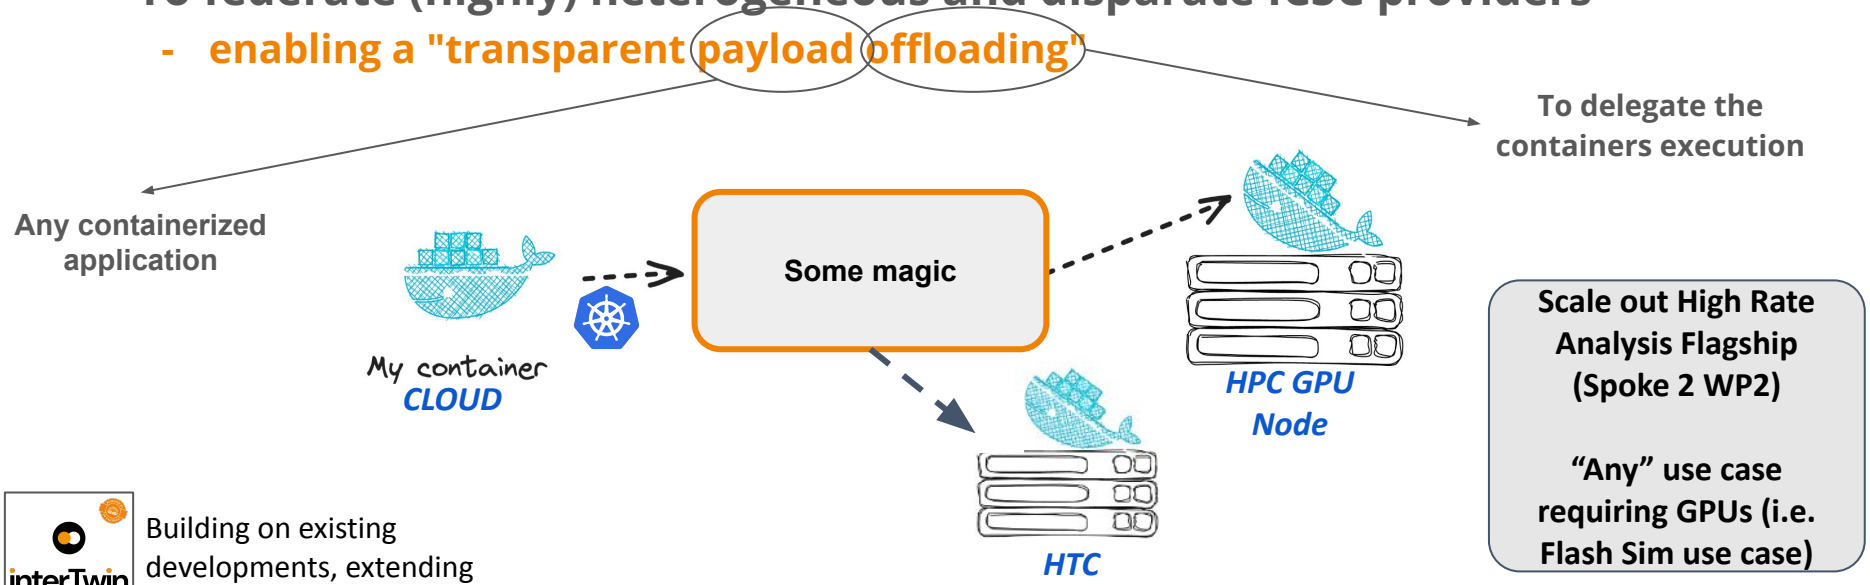

Scale out High Rate Analysis Flagship (Spoke 2 WP2)

"Any" use case requiring GPUs (i.e. Flash Sim use case)

interTwin

Building on existing developments, extending and enhancing

A more "Cloud-native" approach to distribution: *InterLink*

D. Ciangottini, ICSC, CHEP 2024
L. Anderlini, ICSC, SIF 2024

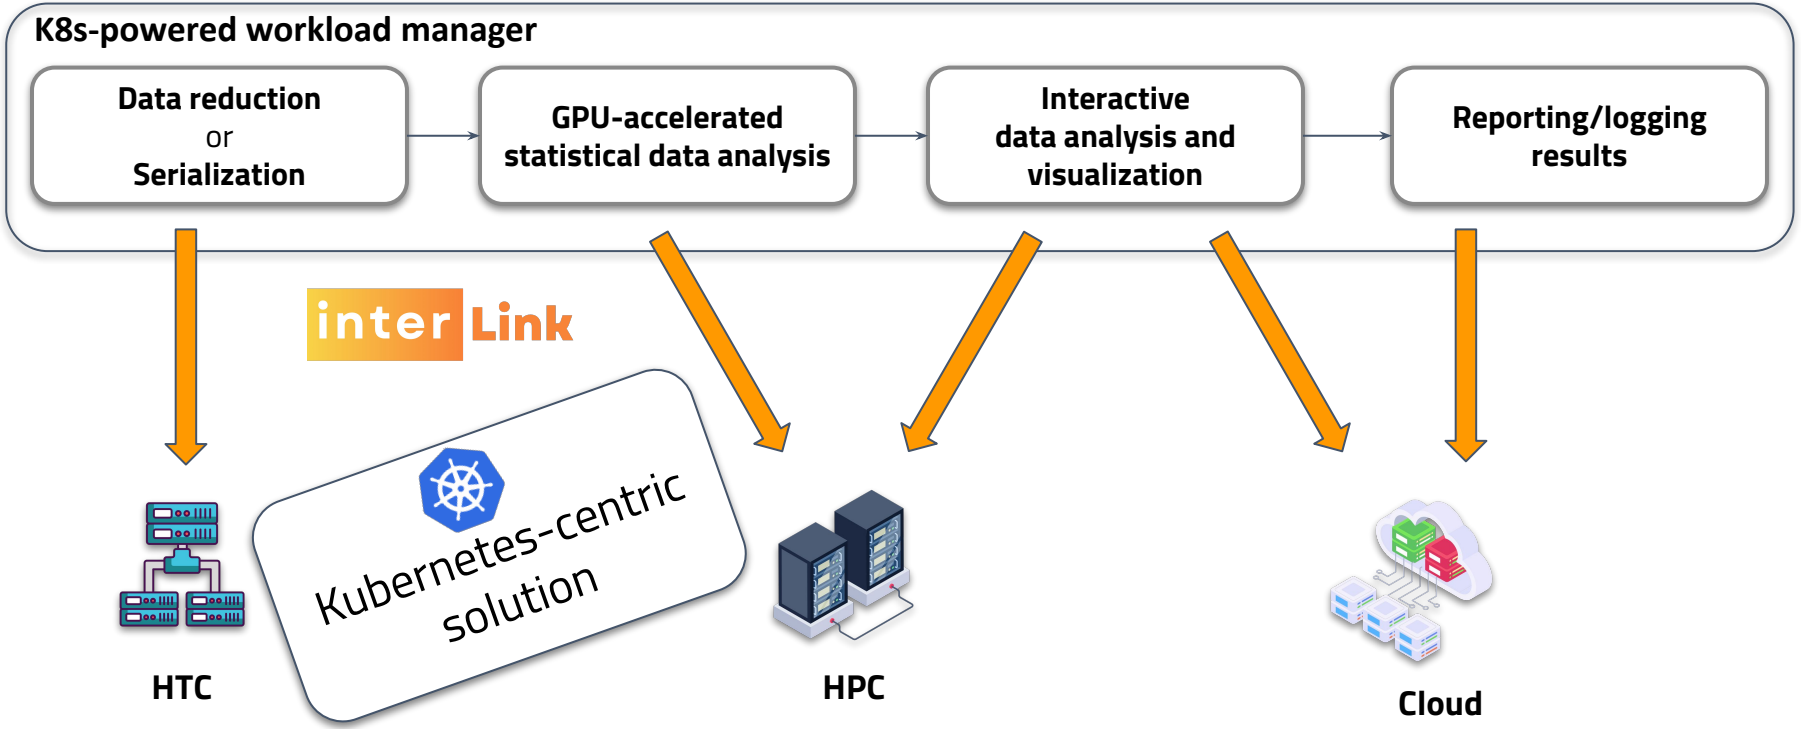

The AI_INFN Platform as a bridge from the Cloud to ICSC

Federated bare-metal resources

Computing resources available to AI_INFN are located at Bologna Technopole within the new CNAF Data Center facility, and managed through a **virtualization layer** (OpenStack of Cloud@CNAF) in **INFN Cloud**:

- 4× servers:
  - 1× 64 CPU cores with 750 GB RAM
  - 3× 128 CPU cores with 1024 GB RAM
- Total local storage: 60 TB of **NVMe disk**
- GPU cards:
  - 8× NVIDIA **Tesla T4**
  - 5× NVIDIA **RTX 5000**
  - 1× NVIDIA **A30**
  - 4× NVIDIA **A100**, potentially served as 4×7 MIG slices
- FPGA boards:
  - 2× AMD Xilinx **Alveo V70**
- 10 GbE connection to CNAF resources

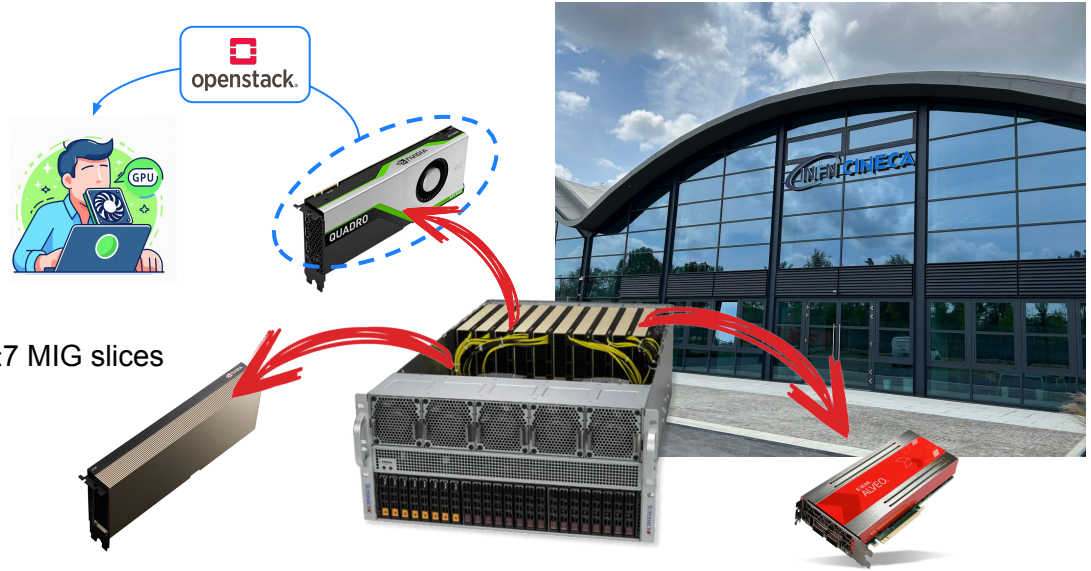

The ML_INFN platform architecture

The **ML_INFN** outcome:

> *Sharing precious GPUs through the Cloud is feasible and effective!*

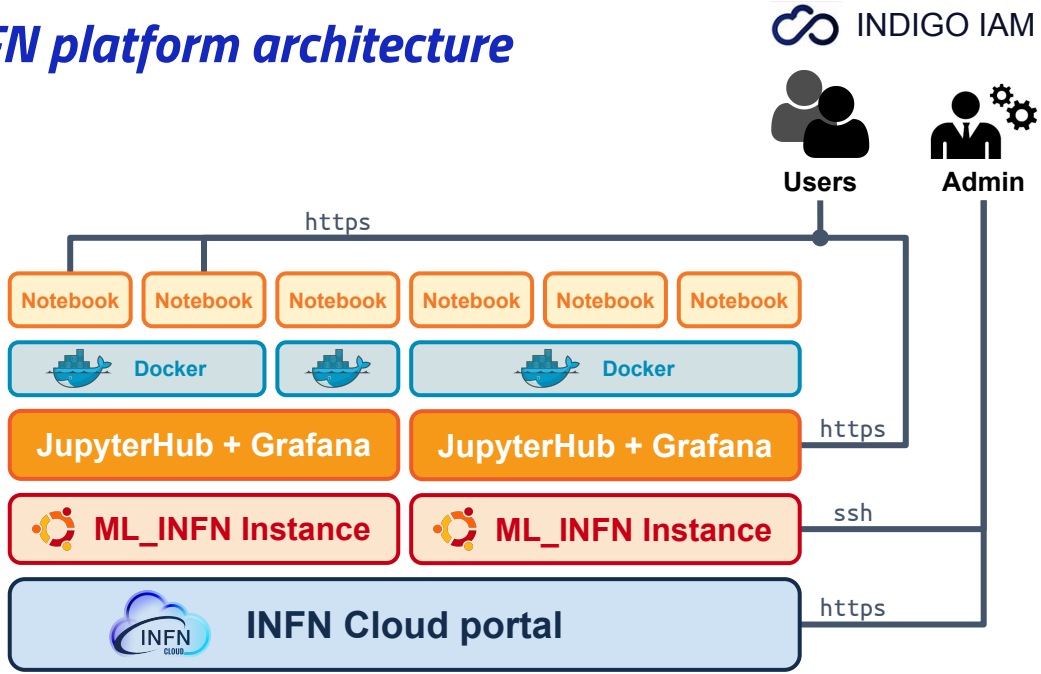

The AI_INFN platform architecture

The **ML_INFN** outcome:

> *Sharing precious GPUs through the Cloud is feasible and effective!*

AI_INFN improves the sharing capabilities:

- addition of an abstract and elastic overlay powered by **Kubernetes**
  - login via AAI → **INDIGO IAM**
  - distributed filesystem
  - managed environments for ML
  - monitoring & accounting
- **data decoupled from computing resources** with a filesystem shared across the VMs
- adding and removing VMs enables manual **horizontal scaling**

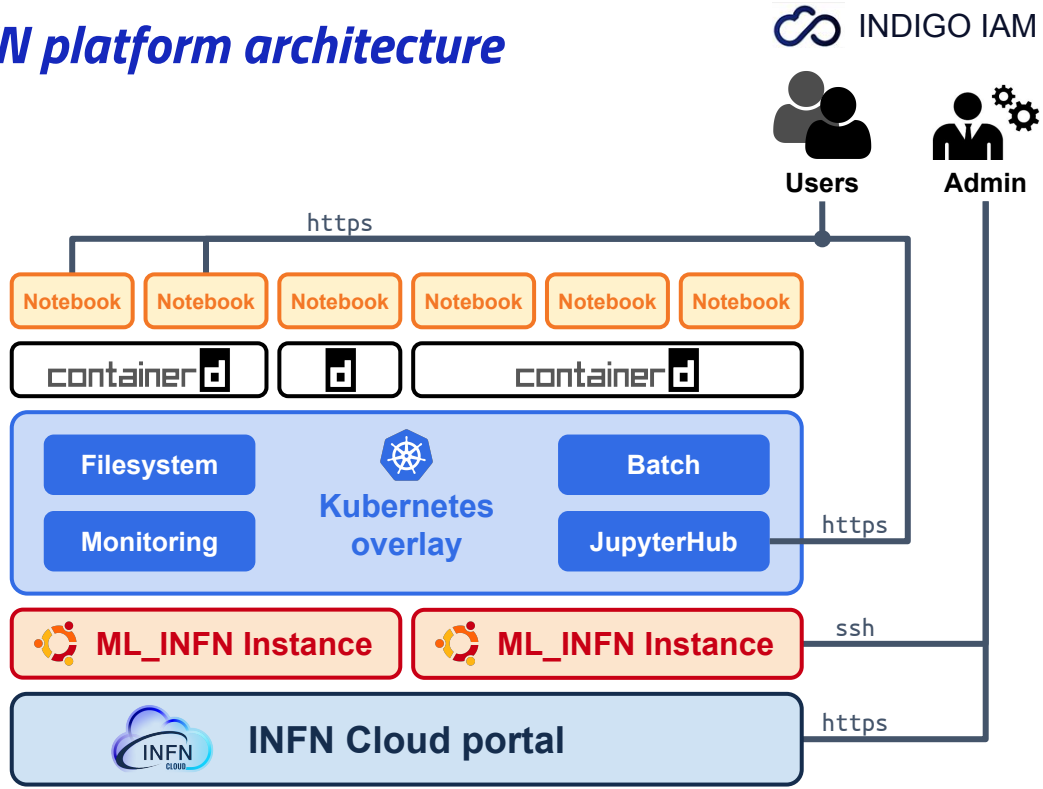

Kueue: k8s-native batch system

Kueue offers a set of APIs and dedicated controllers to simplify and enhance **job queue management** in Kubernetes clusters for batch processing, HPC, AI/ML, and similar applications:

- **Queue management.** Provides a robust infrastructure for job queue management, ensuring reliable and scalable job execution within the Kubernetes cluster
- **Integration with Kubernetes resources.** Kueue integrates natively with Kubernetes resources and functionalities, leveraging the cluster's orchestration and management capabilities
- **Monitoring and Scalability.** With dedicated controllers, Kueue simplifies job state monitoring and enables automatic resource scaling based on workload demands

vk-dispatcher

vkd provides an **authenticated delegation layer** between JupyterHub and Kubernetes, enabling the translation of a user's interactive session into a Kubernetes job

The **Kubernetes job** are submitted to queues managed by **Kueue** that may be enabled for specific projects through the JupyterHub groups

Enabling offloading using interLink as Virtual Kubelet provider

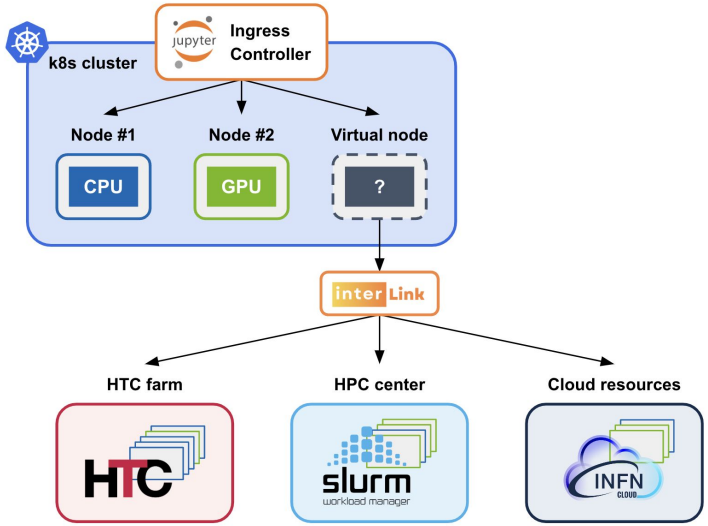

extension of the AI_INFN platform through the VK mechanism

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:

- VKs provide k8s cluster with **"*Virtual Computing Nodes*"** that have no networking towards the API server or other services
- VKs are **ideal for batch processing**, where the connection between the cluster and the working node is only needed at job submission and retrieval

The interLink protocol offers a batch-system native backend for Virtual Kubelets (e.g., SLURM, HTCondor, or other Kueue instances)

Developed as part of the Flagship Activities of ICSC Spoke 2

First workflow with INFN Tier-1, CINECA Leonardo & AI_INFN resources

On October 25th, we run a first workflow combining CPU resources from **CNAF Tier-1**, **Leonardo** and a **local node**.

Test limited to CPU-only payloads.

Offloading to Leonardo booster (with GPU payloads) coming soon.

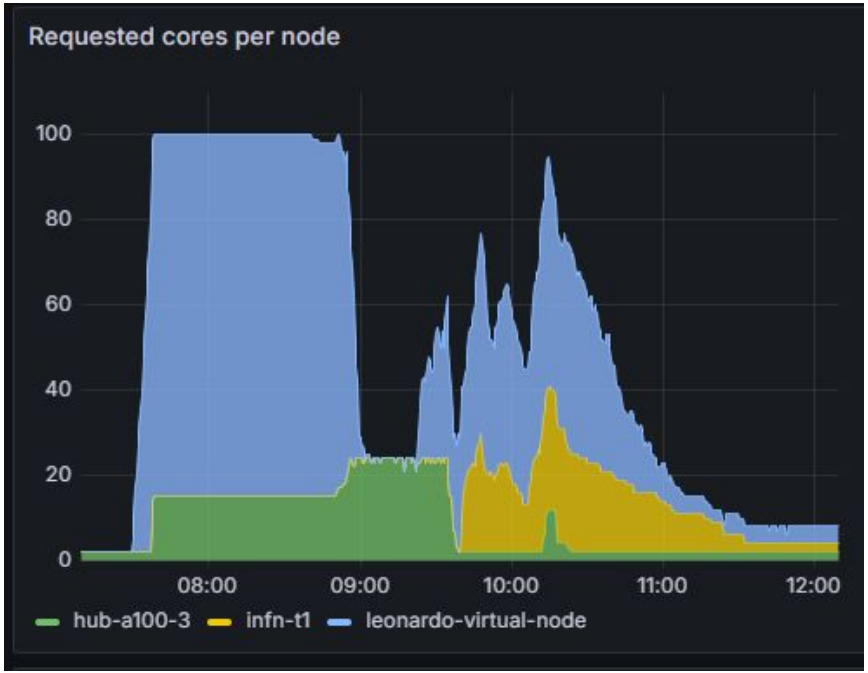

Some(!!!) applications of machine learning studied within ICSC

- Fast Reconstruction of the ATLAS experiment
- Flash Simulation of the CMS experiment
- Lamarr, Flash Simulation of the LHCb experiment
- Fast Simulation of cascades from cosmic rays in atmosphere
- Physics Informed Machine Learning for AI-based simulations
- FADeR, Real-time track reconstruction of the LHC experiment
- Machine Learning for Predictive maintenance
- SAIFIN, Satellite data and Artificial Intelligence for FINtech
- Event reconstruction in Super-Kamiokande
- Fast detector simulation for 3D Silicon Pixel detectors
- Super-resolution surrogate model for accelerated Geant4 simulations
- Satellite data management for advanced environmental applications
- Particle ID reconstruction for Future Collider Experiments
- Searches for low-frequency Gravitational Waves
- Model sources of gravitational waves
- Processing of 3D Medical Images (e.g. CT Scans)

Quantum machine learning applications

- **First applications of Quantum Computing to LHCb data**
- Jet identification, Flavour Tagging, calorimeter simulation
- **Variational Quantum Circuit**: parameters determined with minimization of a loss function
  - Evaluation of the circuit on quantum hardware (IBM)
- **Fast simulation** of ECAL module with quantum methods:
  - Quantum Born Machines (QBM)
  - Quantum Generative Adversarial Networks (qGAN)

Variational Quantum Circuit: Angle embedding

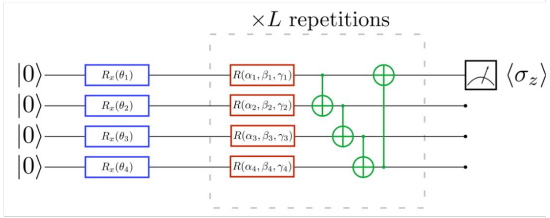

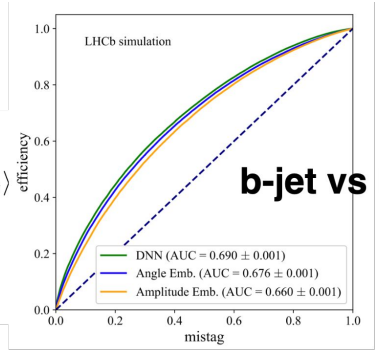

Original dataset: 100 GeV electrons impinging on the module

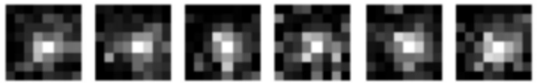

Dataset generated by the trained quantum circuit (7 qubits)

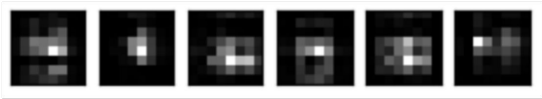

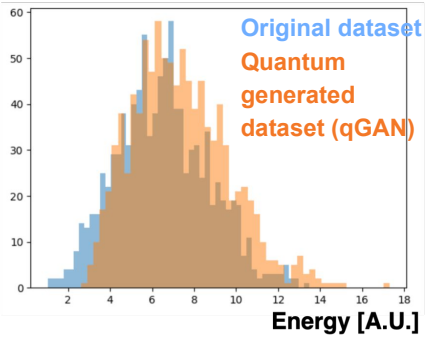



Advanced Hackathon of Machine Learning (AI_INFN)

Bari 2022, Pisa 2023, Padova 2024

School of Open Science in the Cloud (SOSC)

Perugia 2022, Perugia 2023, Bologna 2024


Seminar on Software for Nuclear, Subnuclear and Applied Physics

Alghero 2022, Alghero 2023, Alghero 2024

Conclusions?

ICSC/Terabit are entering its third (last NRRP-funded) year of scientific activity.

It fostered a **huge effort on digital technologies**, and in particular on Machine Learning, enabling access to an unprecedented amount of computing resources.

The call for sustainability of the action is now pressing (see Davide's talk).

The recent selection of Bologna for one of the AI factories confirms that Machine Learning will remain central.