Workshop sul Calcolo nell'I.N.F.N.
Hotel Hermitage
La Biodola - Isola d'Elba
26 - 30 maggio 2025

# INFN Infrastructure for Artificial Intelligence

*Lucio Anderlini* – INFN Firenze

AI_INFN
Artificial Intelligence technologies for INFN research

Istituto Nazionale di Fisica Nucleare
SEZIONE DI FIRENZE

# Gartner's Hype Cycle

In 1995, Gartner introduced a visual representation of the **maturity of an innovative technology**.

"*Artificial Intelligence*" is a plethora of innovative technologies, each one following this cycle.

from Wikipedia

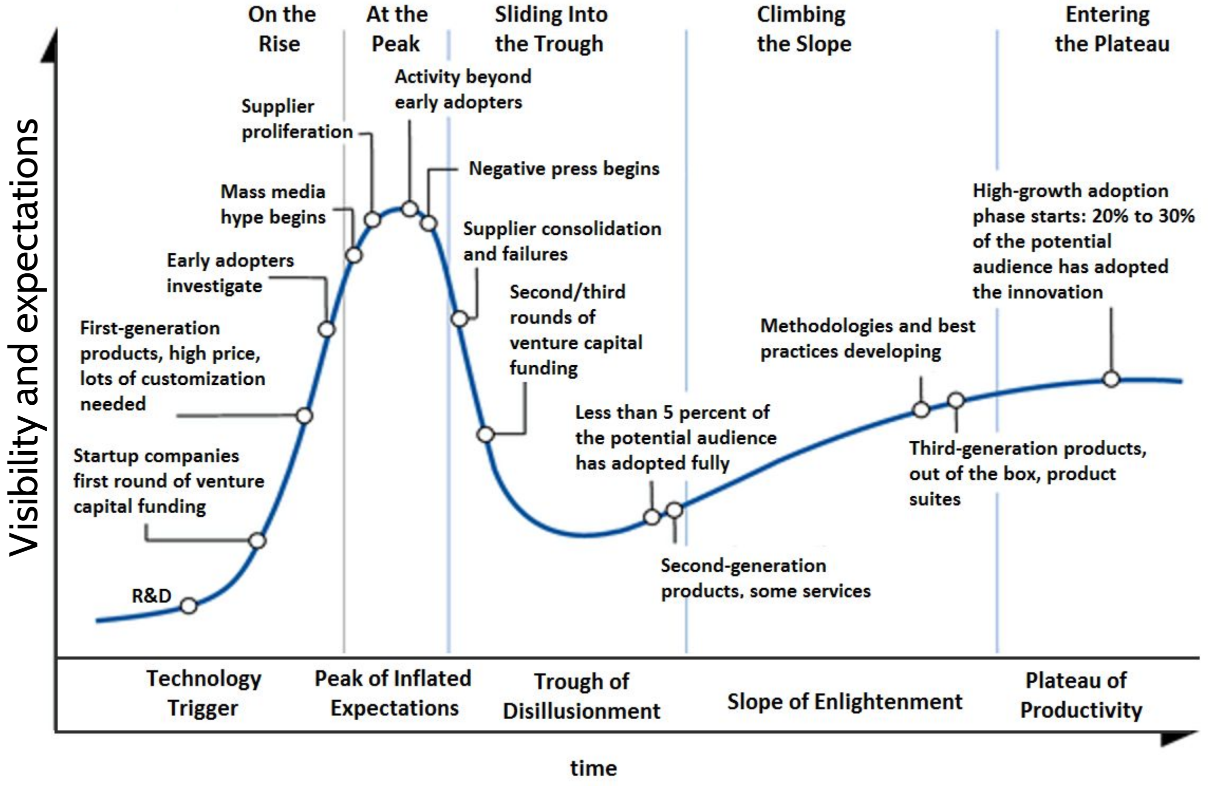

# Two examples of Artificial Intelligence technologies

## Deep Neural Networks

Major revolution in **Computer Vision** in 2012.

Requires:

- ❑ GPU acceleration to develop (and sometime operate)
- ❑ Strong domain knowledge

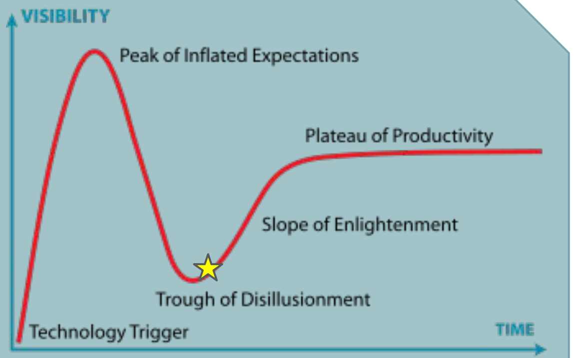



## Generative Pretrained Transformers

Major revolution in **Natural Language Processing** in 2023.

Requires:

- ❑ Several of GPUs to develop and operate
- ❑ Humongeous datasets (*e.g.* the whole Internet)

# Two examples of Artificial Intelligence technologies

## Deep Neural Networks

Major revolution in **Computer Vision** in 2012.

Requires:
- ❑ GPU acceleration to develop (and sometime operate)
- ❑ Strong domain knowledge

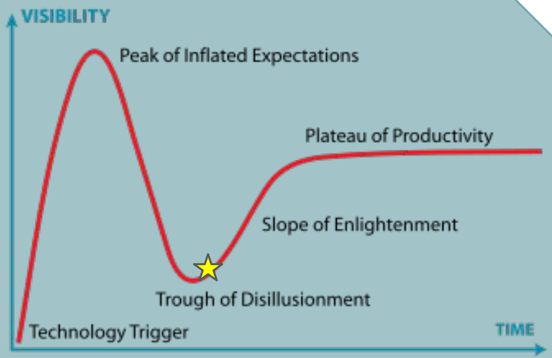


**NOT two use cases for the same infrastructure**


## Generative Pretrained Transformers

Major revolution in **Natural Language Processing** in 2023.

Requires:
- ❑ Several of GPUs to develop and operate
- ❑ Humongeous datasets (*e.g.* the whole Internet)

# Infrastructure for **Deep Neural Networks**

AI_INFN
Artificial Intelligence technologies for INFN research

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

terabit

# Application of **Deep Neural Networks** to INFN Research

The INFN Community has leadership roles in the development of the three "*components of Deep Learning*".

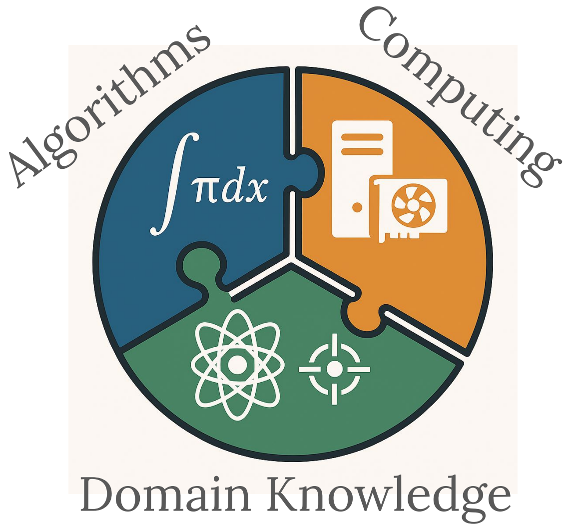


Challenge:

*keep the **communities connected** and the **solutions scalable**.*

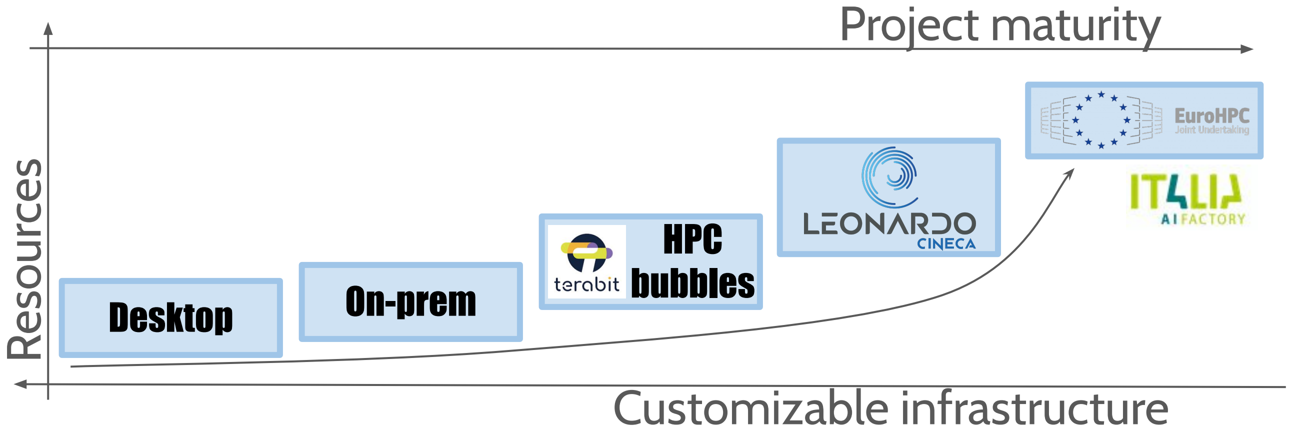

# A distributed infrastructure

INFN is a “distributed institute” by nature.

- Most INFN units have local computing resources, often including GPUs.
- Thanks to the TeRABIT project and Data Cloud integration,
  - Fast unit-to-unit network is provided
  - Bleeding-edge HPC resources are installed (HPC Bubbles)
    - 160 CPU nodes
    - 61 GPU nodes (244 nVidia H100)
    - 10 FPGA nodes (24 Xilinx + 16 Altera FPGAs)
    - 118 High Performance stroage nodes (45 PB)
- Provisioning models:

See Claudio’s talk

Rich matrix of options!

«Heterogeneity» is the keyword

## WP1 Infrastructure and Resource Provisioning

**Lots of resources coming from ICSC and TeRABIT?**
→ Less "*pressure for being in production*" on our farm
→ Opportunity for contributing to the provisioning model

Focus shifts towards R&D on the provisioning model, with a systemic view to ease ML workloads.

**Needs for an updated and well maintained farm.**

## WP2 Open Science and Advanced Education

**What will be added-value in our *hackathons*?**
→ ML_INFN has attracted a community of world-leading experts in the application of ML to research in physics
→ We can be more ambitious in the target of our *hackathons,* letting experts to discuss their code

**Focus shift towards *Advanced Hackathon Workshops.***

**AI_INFN**
Artificial Intelligence technologies for INFN research
ai.cloud.infn.it

## WP3 Scientific use cases

**Applications to scientific research remains central.**

To develop the tools for making it easier to do machine learning for INFN researchers,
we need them to use the tools and provide feedback.

***How?* User support and community engagement**

## WP4 ML on FPGA and Quantum Computers

**New hardware will change the landscape of computing.**

Deploying ML algorithms of **FPGAs** enables fixed-latency, low-energy inference.
**Quantum Computing** will enable extremely fast computations of specialized algorithms.

# The AI_INFN Platform on INFN Cloud

[hub.ai.cloud.infn.it](https://hub.ai.cloud.infn.it)

```
Context: default          <c>      Cordon    <u> Uncordon
Cluster: default          <ctrl-d> Delete    <y> YAML
User:    default          <d>      Describe
K9s Rev: v0.40.8 ⚡ v0.40.10 <r>    Drain
K8s Rev: v1.31.6+rke2r1   <e>      Edit
CPU:     0%               <?>      Help
MEM:     0%
```

v1/nodes(all)[13]

| NAME↑ | STATUS | ROLE | TAINTS | VERSION | PODS | CPU | MEM | %CPU | %MEM | CPU/A | MEM/A | AGE |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| h5-cpu | Ready | <none> | 0 | v1.31.6+rke2r1 | 14 | 169 | 19064 | 2 | 29 | 8000 | 64300 | 18d |
| h5-gpu-a30 | Ready | <none> | 1 | v1.31.6+rke2r1 | 18 | 149 | 8083 | 1 | 12 | 8000 | 64300 | 14d |
| h5-gpu-a100-1 | Ready | <none> | 0 | v1.31.6+rke2r1 | 27 | 1581 | 62048 | 5 | 36 | 28000 | 169031 | 18d |
| h5-gpu-a100-2 | Ready | <none> | 0 | v1.31.6+rke2r1 | 32 | 618 | 6982 | 2 | 4 | 28000 | 169031 | 18d |
| h5-gpu-a100-3 | Ready | <none> | 1 | v1.31.6+rke2r1 | 21 | 427 | 19398 | 1 | 11 | 28000 | 169031 | 18d |
| h5-gpu-rtx-1 | Ready | <none> | 0 | v1.31.6+rke2r1 | 19 | 234 | 19127 | 1 | 19 | 16000 | 100586 | 18d |
| h5-gpu-rtx-2 | Ready | <none> | 0 | v1.31.6+rke2r1 | 20 | 155 | 14887 | 0 | 14 | 16000 | 100586 | 18d |
| h5-gpu-rtx-3 | Ready | <none> | 0 | v1.31.6+rke2r1 | 16 | 83 | 4359 | 1 | 8 | 8000 | 50187 | 18d |
| h5-gpu-t4 | Ready | <none> | 0 | v1.31.6+rke2r1 | 18 | 146 | 20748 | 0 | 20 | 16000 | 100586 | 18d |
| h5-master-1 | Ready | control-plane,etcd,master | 0 | v1.31.6+rke2r1 | 24 | 326 | 4620 | 4 | 7 | 8000 | 64300 | 18d |
| h5-master-2 | Ready | control-plane,etcd,jhub,master | 0 | v1.31.6+rke2r1 | 18 | 208 | 4104 | 2 | 6 | 8000 | 64300 | 18d |
| h5-master-3 | Ready | control-plane,db,etcd,master | 0 | v1.31.6+rke2r1 | 18 | 155 | 3700 | 1 | 5 | 8000 | 64300 | 18d |
| interlink | Ready | agent | 1 | 0.3.7 | 0 | 0 | 0 | 0 | 0 | 10000000 | 2199023255552 | 17d |

## Welcome to the AI_INFN Platform!

You are logged as: boccali

You are member of the following projects: cms

The platform is currently under active development, and not compliant with GDPR rules, yet.
**Do not upload personal data. Uploaded data might be lost.**

You might want to read the User's Guide and contribute to our Knowledge Base.

Docker image: harbor.cloud.infn.it/datacloud-jupyter/jlab-ai-infn:2.0.0-ai1.3

Number of cores:
- (●) 1 ( ) 2 ( ) 3 ( ) 4 ( ) 8 cores

Memory size:
- (●) 2 ( ) 4 ( ) 8 GB

Hardware accelerator:
- (●) None
- ( ) nVidia A100 1g MIG partition (10 GB) *13/14 available*
- ( ) nVidia T4 (16 GB) *0/2 available*
- ( ) nVidia RTX5000 (16 GB) *0/5 available*
- ( ) AMD-Xilinx U55c FPGA *1/1 available*

Start

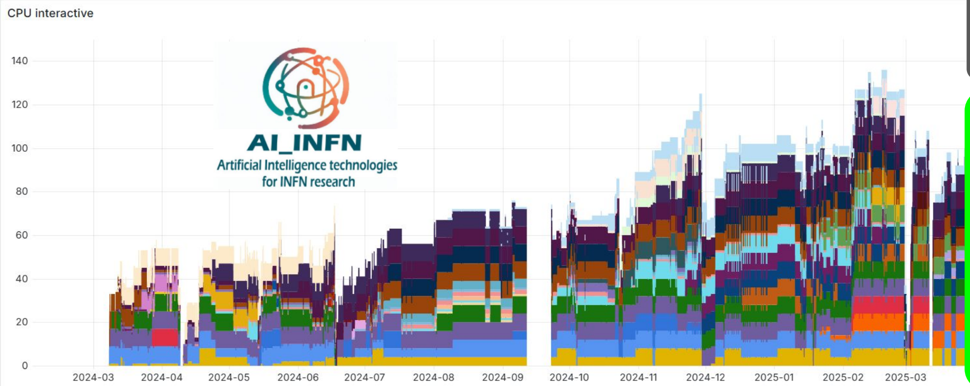


For more info:
[cloud-support@infn.it](mailto:cloud-support@infn.it)

Platform selected for: Fondazione ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

- ➢ **ICSC Spoke 0** Integration PoC
- ➢ **ICSC Spoke 2** – Flagship Project on Flash Simulations
- ➢ **ICSC Spoke 2** – Innovation Grant on Physics Informed ML

# Scaling out to the distributed INFN infrastructure

**Workload Offloading:** extending the container orchestration layer (K8s) to support the execution of a user payload to a remote resource instead of a physical node of the cluster

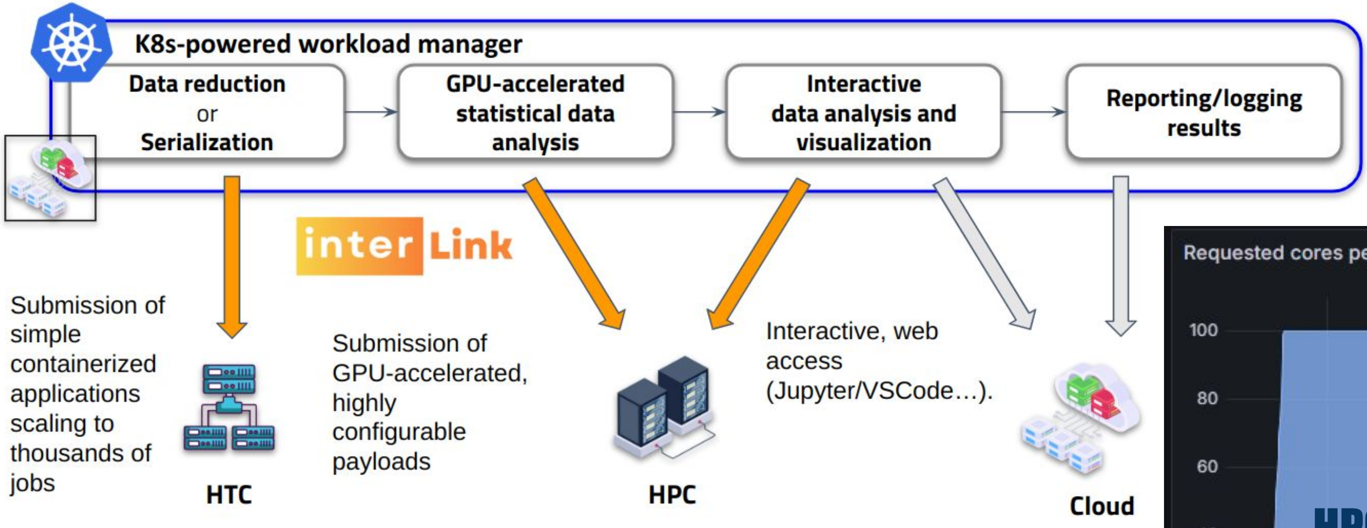


... more in Giulio Bianchini's talk

✅CINECA Leonardo
✅TeRABIT HPC Bubble PD
Poster by Antonino T.

✅INFN-CNAF Tier-1
✅ReCaS-Bari Tier-2

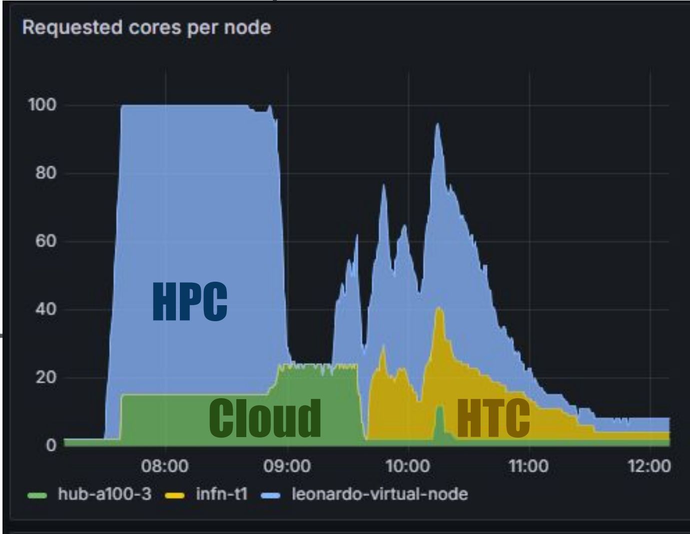

# What come next? MLOps

Workflow management:

- Integration with PSI ELOG for logging and tracking
- DAG management with Snakemake
- Tools beyond Jupyter: *Kubeflow* See Mauro's talk
- Integration with baltig.infn.it for **CI/CD** (planned)

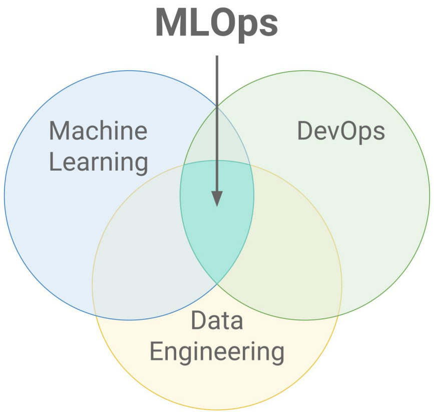


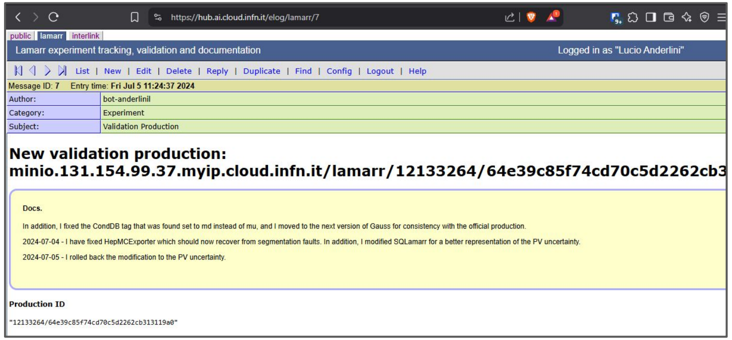


snakemake

Kubeflow

GitLab

# Storage & Data Management

## User's homes (software & configuration)

Developing ML requires **interactive tinkering**.

Switching from a GPU model → switch server.

An **responsive network** file system, tuned for intense **metadata activity** is needed:

→ *prototyped **bcached Cinder volume via NFSv4***

Talk by Rosa P.

## Data Management for Training & Validation

Training data and metadata are often the most precious asset.

**Security + privacy + throughput**

Several tools & backends explored: Ceph+RadosGW (CNAF), Invenio RDM (GLOS), MinIO on K8s (Firenze), xNATS (Pisa), OwnCloud (CH-Net)

## Artifact management and distribution

Trained models should be versioned, and distributed.

We prototyped a solution based on **ORAS** (OCI Registry As Storage) for using **INFN Cloud harbor** for artifacts (e.g. trained models).

Automatic upload of artifacts to **cvmfs** via unpacked.

Poster by Alessio F.

credits to *Marco Verlato*

## Multi-site distributed file system

POSIX access to data **decouples the AuthN/Z from the application layer**.

**Apptainer** leverages modern features of the **Linux kernel** to enable mounting custom **fuse volumes**, without admin privileges.

→ Successful integration tests with InterLink.

# Medical Data Platform

Poster by Antonino F.

When the artifacts are **intermediate data** (*e.g.* pretraining output), dedicated solutions are necessary.

- Development of a dedicated data platform for secure medical data management

It will be based on an open-source imaging IT platform developed by Washington University for neuroimaging data analysis: XNAT
Thanks to its **extensibility**, XNAT can be used to support a wide range of imaging-based projects.

***Extensible Neuroimaging Archive Toolkit*** *https://www.xnat.org/*

Data modeling and simulations

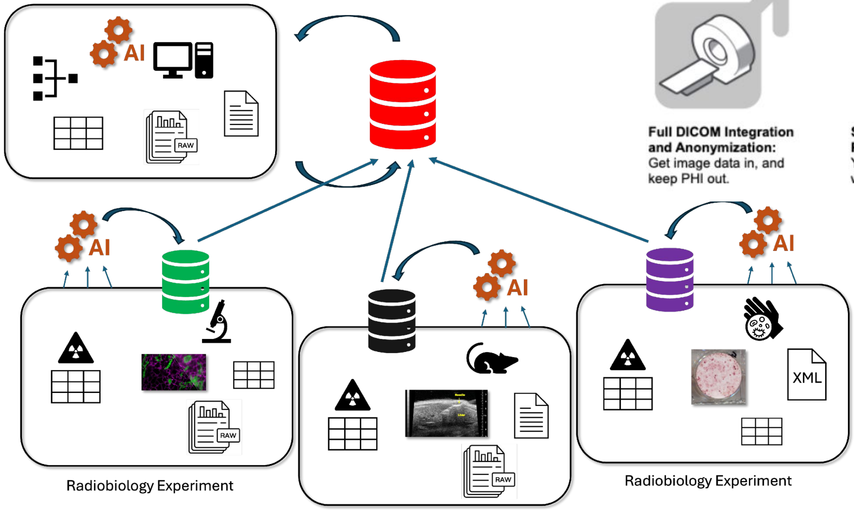


Radiobiology Experiment

Radiobiology Experiment

Radiobiology Experiment

**Full DICOM Integration and Anonymization:** Get image data in, and keep PHI out.

**Secure Access & Permission Control:** You decide who does what with your data.

**Integrated Search & Reporting:** Report on your image and clinical data together.

**Pipeline Processing:** Use the power of high-performance computing on your data.

**Modular Extensibility:** Expand the capabilities of your XNAT to meet your needs.

**Developer Community:** Benefit from an active and engaged set of XNAT power users.

Integrated with INFN AuthN/Z.

Integrated with pipelines running in SLURM in HPC@Pisa.

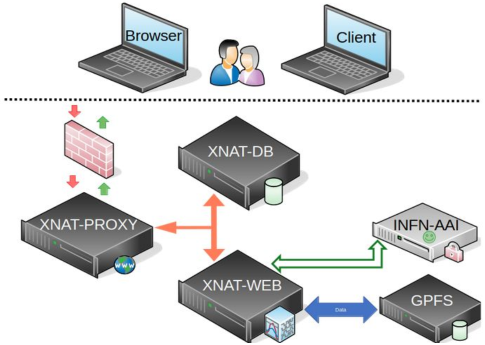

# Infrastructure for **GPT**

# Tens or hundreds of GPUs

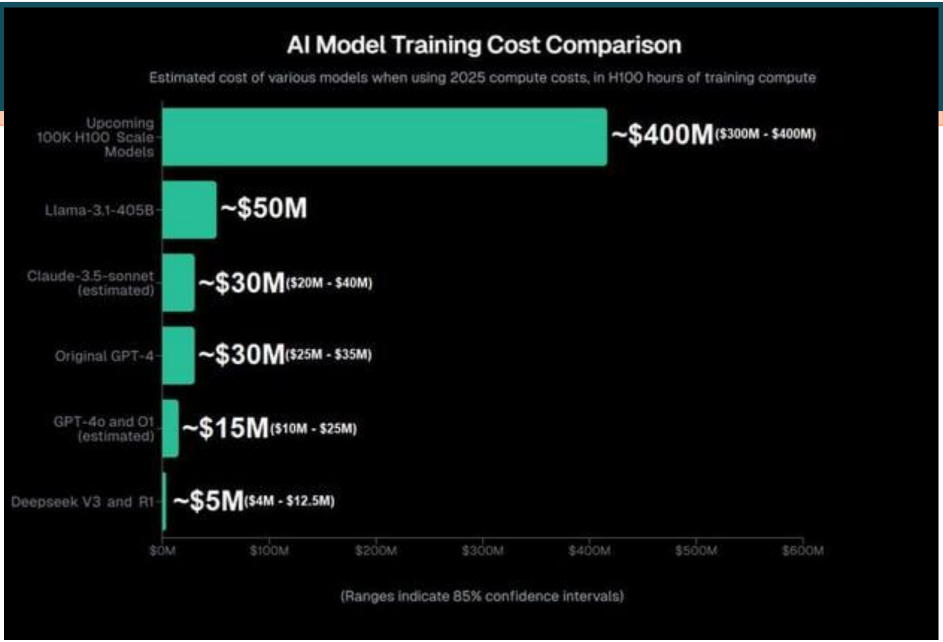


Transformers (the T in GPT) are made special by *attention layers*, enabling extraordinary scalability.

Large Language Models enabled scaling to thousands of GPUs processing "*a large fraction of the Internet*".

Training a model costs as a HEP experiment.

→ **Hardware and Software specialization**


NVIDIA is investing in integrated Arm + GPU chips providing extremely fast communication between CPU and GPU operations.

Architecture specialied for training and inference of DNNs and Transformers.

**Pioneering studies at CNAF and in Ferrara.**

Talk by Alessandro Pascolini

Poster by Enrico Calore

# IT4LIA AI Factory

## Major investment from EU

Italy (and Bologna) selected for one of the first-round EuroHPC factories: O(430 M€)

*Target*: Large Language Models and Transformers.

CINECA leads the project, the installation site: next to INFN Tier-1.

## Project evolution

**Kick-off meeting**: May 20-21, 2025 (last week)

The project is currently defining the *procurement strategy* and *collecting requirements* from the communities of the future users.

**INFN involved in the Data Management WP**, *in continuity with the CINECA-INFN Integration PoC.*


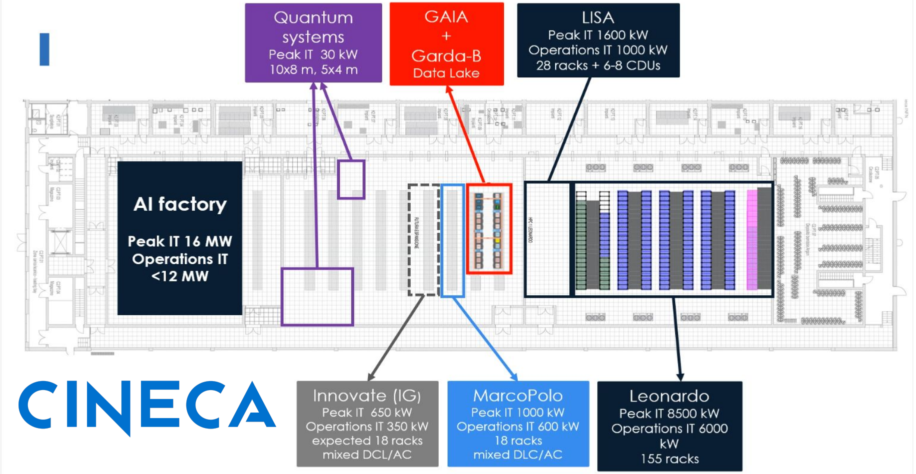

# Conclusion

# Conclusion

**Deep Learning is beneficial to INFN research already**, as evident from the agenda of this workshop!

**INFN is heterogeneous and distributed by nature**: need something for you ML? INFN has it somewhere.

**AI_INFN and DataCloud are tracking and developing tools to benefit from such diversity.**

**Offloading** *for computing power*, **Data Management** *for distributed access*, **MLOps** *for collaborative work*

Solutions explored with **ICSC Spoke 0 Integration PoC inspire future challenges** (e.g. the AI factory)

**Large Language Models are defining the top-tier hardware solutions.** Are we ready with our software?

# Backup

# M2.1 - Organization of the first Advanced Hackathon Worksho

**Proposed status**
**100%**

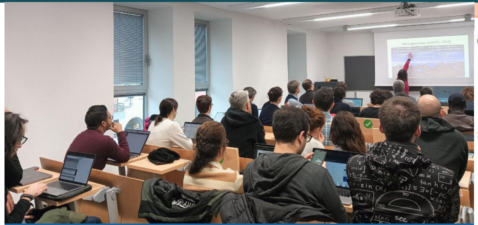

## 1st AI-INFN Advanced Hackathon

26–28 Nov 2024
University of Padua, Complesso Paolotti
Europe/Rome timezone

- Overview
- Timetable
- Registration
- Experts and tutors
- Access to cloud resources
- Groups and servers
- Streaming

Contact
ml-infn-hackathons@list...

Welcome to the First edition of the Advanced Artificial Intelligence @ INFN (AI_INFN) hackathon, dedicated to INFN affiliates. This edition is hosted at INFN Sezione di Padova.

Notably, it is the third Hackathon to happen in Person, so please apply only if you are planning to come to Padua. The logistics allow for ~ 20 participants.

agenda.infn.it/event/43129/

### Day 1

- 09:00 - 09:15 Welcome — Classroom P4C, University of Padua, Complesso Paolotti Via Paolotti
- 09:15 - 09:50 AI in Europe and WLCG vision — Tommaso Boccali — Classroom P4C, University of Padua, Complesso Paolotti Via Paolotti
- 09:50 - 10:25 Cloud Veneto and INFN Cloud — Marco Verlato — Classroom P4C, University of Padua, Complesso Paolotti Via Paolotti
- 10:25 - 10:45 On the physics case of the LHCf exercise — Elena Gensini — Classroom P4C, University of Padua, Complesso Paolotti Via Paolotti
- 10:45 - 11:00 Hands-on: accessing the Hackathon resources — Lucio Anderlini — Classroom P4C, University of Padua, Complesso Paolotti Via Paolotti
- 11:00 - 11:20 Coffee break — Lucky bar - The Brick
- 11:20 - 12:00 Processing data from the LHCf detector — Mr Andrea Paccagnella et al. — Classroom P4C, University of Padua, Complesso Paolotti Via Paolotti
- 12:00 - 13:00 Generative models to unfold detector effects — Fabio Rossi et al. — Classroom P4C, University of Padua, Complesso Paolotti Via Paolotti
- 14:30 - 16:00 Use of a multidimensional CNN for particle identification in the LHCf experiment — Mr Andrea Paccagnella et al. — Classroom P4C, University of Padua, Complesso Paolotti Via Paolotti
- 14:30 - 16:00 Generative Adversarial Networks as a tool to unfold detector effects — Tommaso Vittorini et al. — Classroom P2B, University of Padua, Complesso Paolotti
- 16:00 - 16:20 Coffee break — Lucky bar - The Brick
- 16:20 - 18:00 Generative Adversarial Networks as a tool to unfold detector effects — Tommaso Vittorini et al. — Classroom P2B, University of Padua, Complesso Paolotti
- 16:20 - 18:00 Use of a multidimensional CNN for particle identification in the LHCf experiment — Mr Andrea Paccagnella et al. — Classroom P4C, University of Padua, Complesso Paolotti Via Paolotti

### Day 2

- 09:00 - 10:35 Introduction to Quantum Machine Learning — Laura Cappelli et al. — Classroom P4C, University of Padua, Complesso Paolotti
- Quantum-inspired tensor-network machine learning: finding optimal hyperparameters, libraries, and hardware — Daniel Jaschke
- 11:00 - 11:20 Coffee — Lucky Bar - The Brick
- 11:20 - 12:20 Introduction to Magnetic Resonance Imaging — Francesca Brero — Classroom P4C, University of Padua, Complesso Paolotti
- 12:20 - 13:00 Introduction to ASD/TD classification on structural and functional MRI — Francesca Lizzi — Classroom P4C, University of Padua, Complesso Paolotti
- 14:30 - 16:00 Autism Spectrum Disorders (ASD) diagnosis using structural and functional Magnetic Resonance Imaging and Radiomics — Francesca Lizzi et al. — Classroom P4C, University of Padua, Complesso Paolotti
- 14:30 - 16:00 Quantum Machine Learning applications: classification, anomaly detection and QUBO problems — Laura Cappelli et al. — Classroom P2B, University of Padua, Complesso Paolotti
- 16:00 - 16:20 Coffee break — Lucky Bar - The Brick
- 16:20 - 18:00 Quantum Machine Learning applications: classification, anomaly detection and QUBO problems — Laura Cappelli et al. — Classroom P2B, University of Padua, Complesso Paolotti
- 16:20 - 18:00 Autism Spectrum Disorders (ASD) diagnosis using structural and functional Magnetic Resonance Imaging and Radiomics — Francesca Lizzi et al. — Classroom P4C, University of Padua, Complesso Paolotti

### Day 3

- 09:00 - 09:40 Machine Learning for Future Colliders — Dr Riccardo Torre — Classroom P2B, University of Padua, Complesso Paolotti
- 09:40 - 10:20 Applications of machine learning in the event reconstruction of Imaging Atmospheric Cherenkov Telescopes — Ilaria Viale — Classroom P2B, University of Padua, Complesso Paolotti
- 10:20 - 10:40 What's happening in EU? A gentle introduction to the AI Act — Francesca Lizzi — Classroom P2B, University of Padua, Complesso Paolotti
- 10:40 - 11:00 An introduction to MLOps — Luca Clissa — Classroom P2B, University of Padua, Complesso Paolotti
- 11:00 - 11:20 Coffee — Lucky Bar - The Brick
- 11:20 - 12:05 Experiment Tracking: Hands on — Luca Clissa — Classroom P2B, University of Padua, Complesso Paolotti
- 12:05 - 13:00 Advanced features of the AI_INFN Platform and presentation of the afternoon activity — Lucio Anderlini — Classroom P2B, University of Padua, Complesso Paolotti
- 14:30 - 15:30 Customizing the software environment with conda and apptainer — Lucio Anderlini — Classroom P2B, University of Padua, Complesso Paolotti
- 15:30 - 15:50 Final remarks and closure — Francesca Lizzi — Classroom P2B, University of Padua, Complesso Paolotti

# M2.2 - Lectures available through an e-learning platform (e.g. moodle)

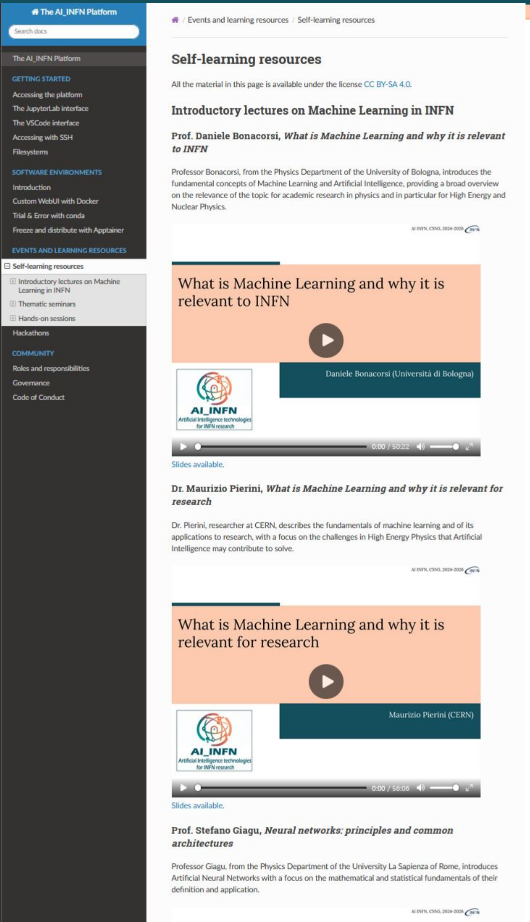


We collected video-lectures from previous hackathons and events in the documentation page, together with self-learning resources.

In total they cover 6 hours of introductory topics in Artificial Intelligence and its applications within INFN research.

# M3.2 - Identified a technological solution for user-self-support
# M3.3 - Platform for user-self-support configured and operational

We made a survey among the early adopters of the AI_INFN Platform.

It was decided not to diverge from the underlying INFN Cloud infrastructure.
We agreed with the cloud-support@infn.it team a policy for user support.

In addition, the survey highlighted:

- the need for documentation: set up con the INFN GitLab instance
https://ai-infn.baltig-pages.infn.it/wp-1/docs/

- the need for an *in-person* workshop other than the hackathon:
  Users' forum organized at CNAF https://agenda.infn.it/event/40489

# M4.2 - Quantum Simulators interfaced to the platform

*Pennylane* e QUBO sono stati integrati nella piattaforma e un esercizio interattivo che li utilizza è stato proposto all'hackathon di novembre.

[link esercizio]

[link slide]

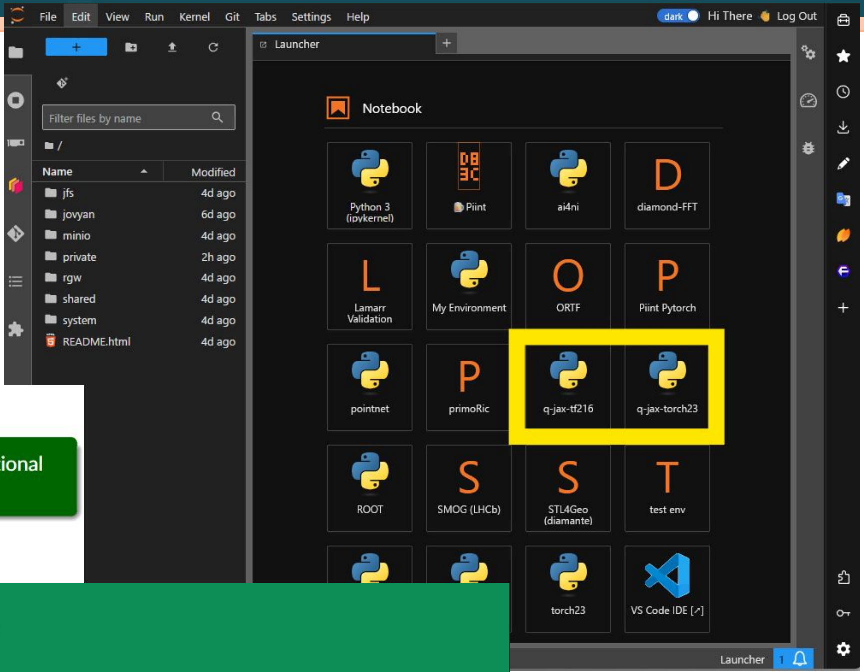


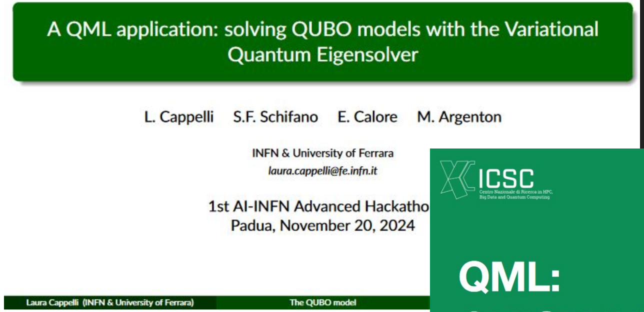


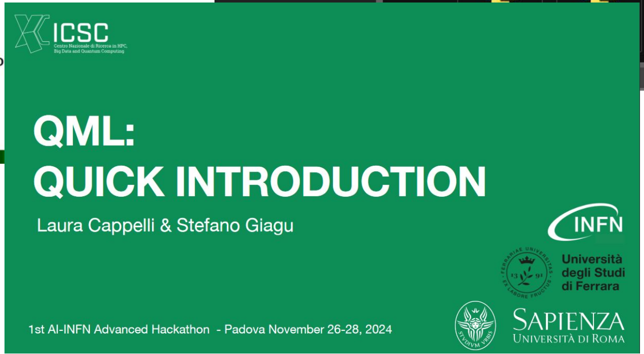

# M1.3 Batch System for opportunistic usage of the cluster resources

Resources left free from interactive usage, can be used for batch workloads.

Symmetric submission mechanism for local and offloaded payloads.

We have three alpha-testers of the currently-evolving setup.

Missing item for entering BETA phase: documentation.


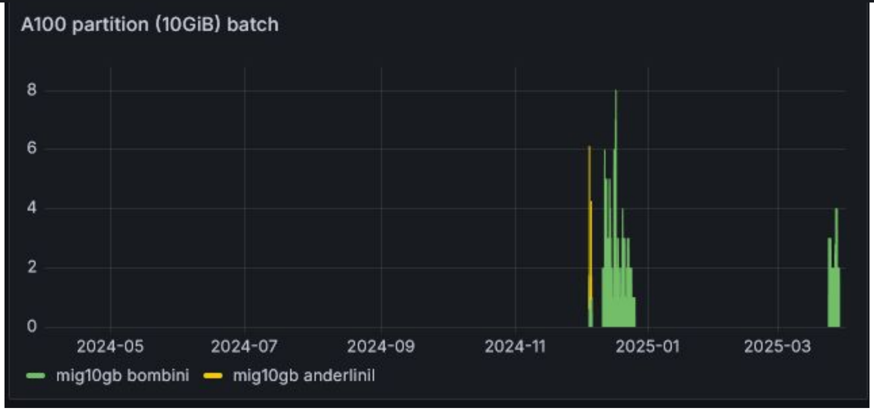

# M1.2 Monitoring and accounting implemented and documented

Dedicated monitoring and accounting dashboard based on:

- Prometheus
- Postgres
- Grafana

Plus custom Python glue here and there.


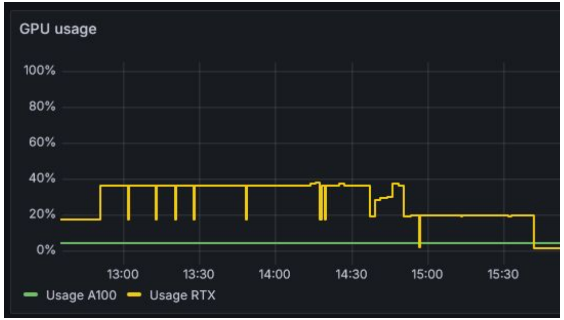



Implementation is converging: *we are starting to document*

# M2.3 Advanced Hackathon Workshop intermedio

As of today, we have:

- a candidate host (to be announced soon);
- the promise for 1 out of 4 new exercises.

# M4.1 Report sull'utilizzo di FPGA in Cloud come acceleratori rilasciato

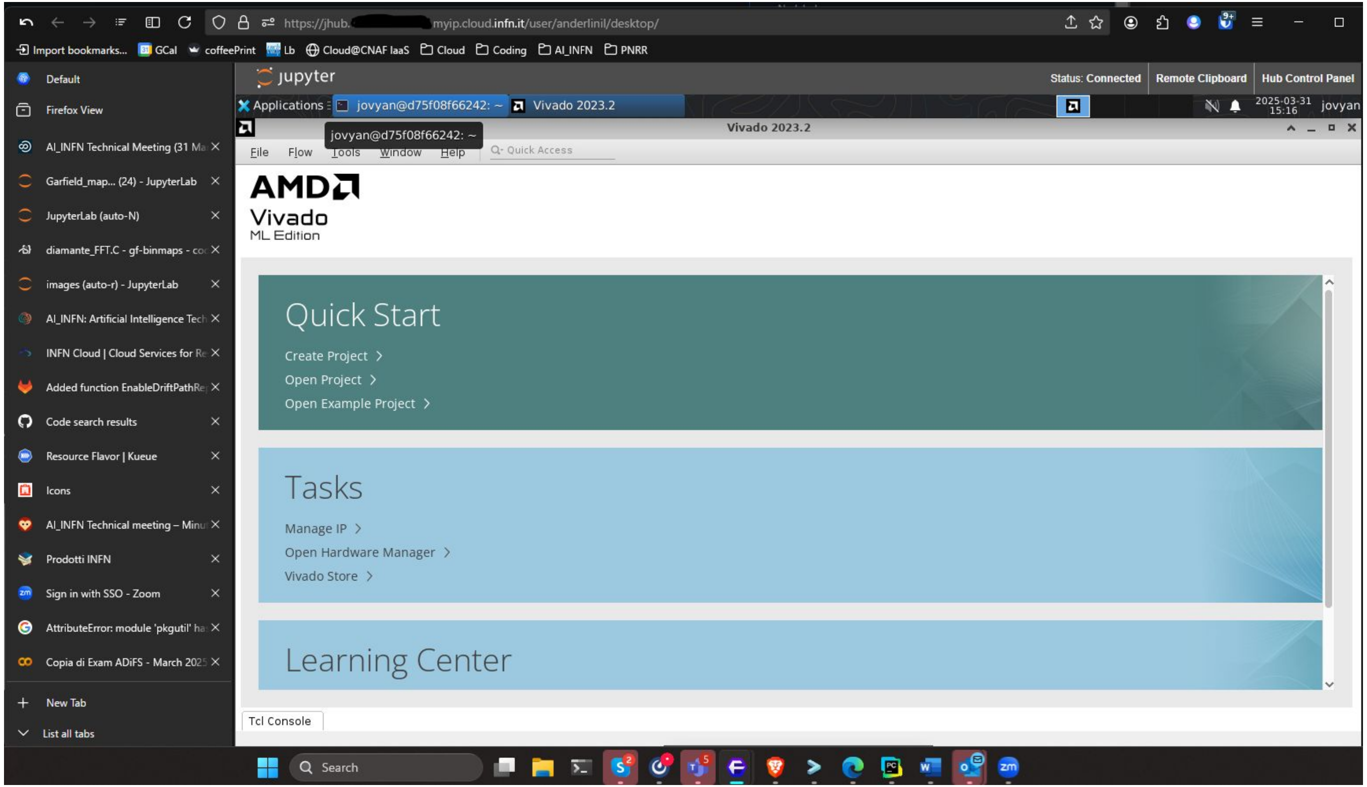


Procedura di acquisto e installazione delle V70 completata.

Purtroppo le V70 non sono utilizzabili.

Trattativa con AMD in corso.

Intanto… prototipo per il provisioning in Cloud di due FGPA Xilinx U55C installate a Perugia.

interLink

# M4.2 Pipeline di compressione DNN per FPGA sviluppate e documentate

Abbiamo una pipeline pronta che include:

- ❑ weight pruning;
- ❑ weight and activation quantization;
- ❑ model compression by distillation (sia per classificatione che regressione);
- ❑ input data compression con auto-encoder.

Resta da portarla “in cloud” e documentarla.

# Challenges

Re-allocating GPU accelerators to different clusters requires features not available with the K8s as deployed from the INFN Cloud PaaS.
We are using RKE 2, with some risk of **diverging developments**.

Mitigation: the deployment of the platform is being documented for both Kubernetes distributions, while requesting the missing feature.

Offloading can expand the computing power of the platform *ad libitum*, with distributed storage performance quickly becoming a bottleneck.

We have investigated several Open Source options, none is competitive with commercial solutions (which are also problematic).

Dedicated SSD storage may help, and will probably drive 2026 requests.

# XNAT https://www.xnat.org/

An open-source informatics platform originally developed for neuroimaging data (Washington University).

Thanks to its extensibility and modularity, XNAT can be used to support a wide variety of projects.

## Original Data Model

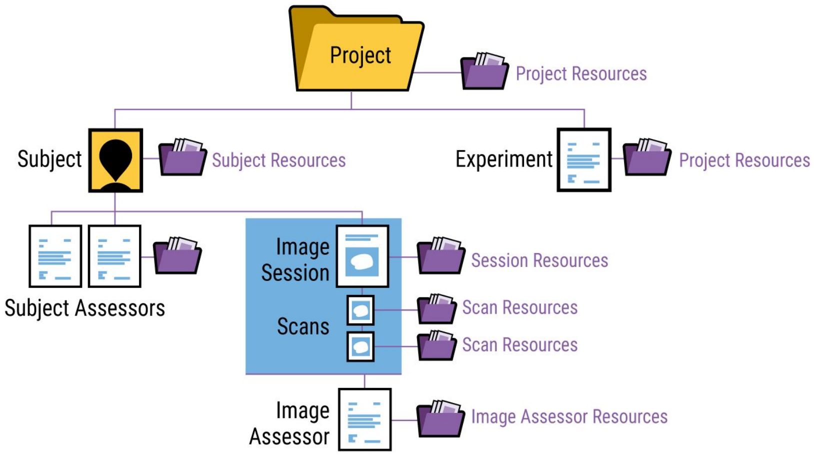


## Customized Data Model

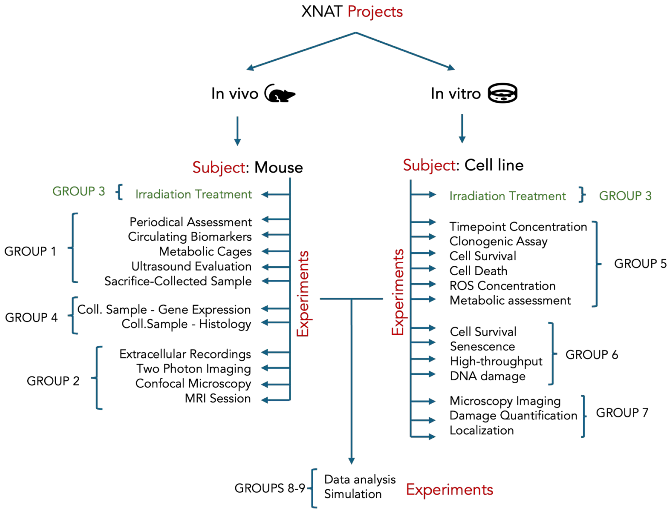

# PERCHÉ UNA PIATTAFORMA DATI?

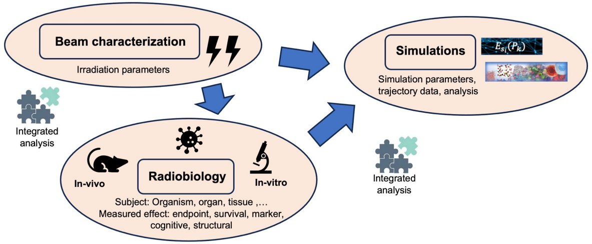


## XNAT-based data platform

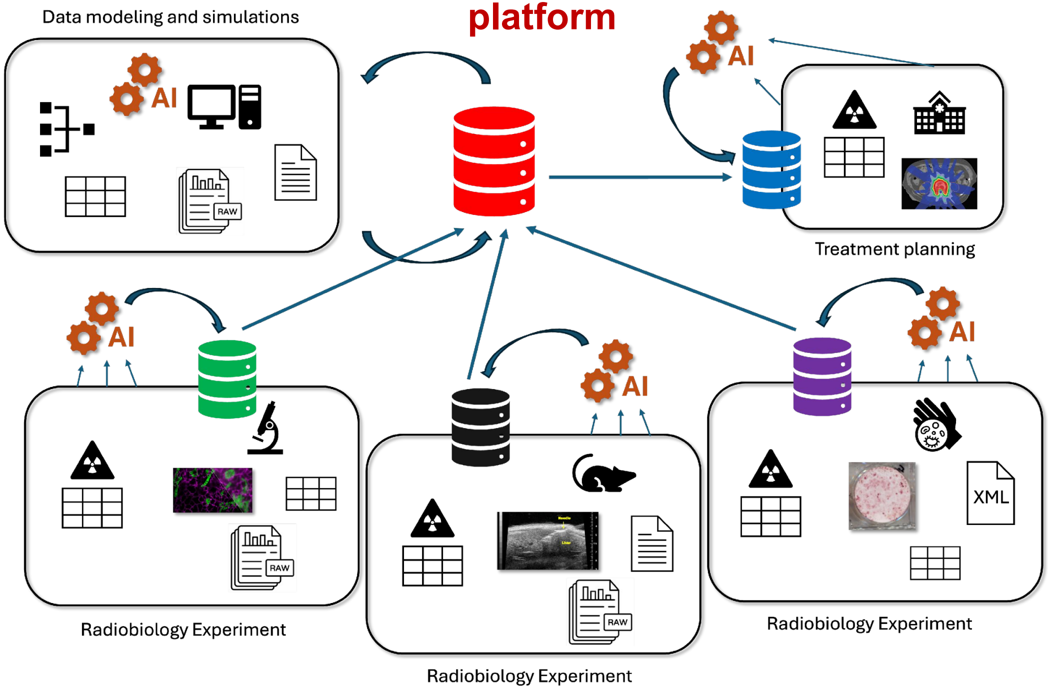


## Multimodal data and heterogeneous parameters to be linked

- Microscopy images, slide photographs, MRI or ultrasound images of small animals, etc.
- Raw files, XML, Excel, CSV, etc.
- Tabular and textual data.
- Heterogeneous parameters: strings, categorical variables, and continuous numerical variables across different ranges.

INFN – CSN5

# Artificial Intelligence in Medicine: focus on Multi-Input Analysis

## The FAIR Guiding Principles for scientific data management

Scientific data should be: **F**indable, **A**ccessible, **I**nteroperable, and **R**eusable

[Wilkinson, M., Dumontier, M., Aalbersberg, I. *et al.* The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data* **3**, 160018 (2016). https://doi.org/10.1038/sdata.2016.18]

XNAT Data Platform to:

- Collect multimodal medical imaging data
- Share original and pre-processed data among the groups
- Integrate analysis pipelines within the platform to preserve new knowledge (e.g. pre-processing step)

☐ Development of a dedicated data platform for secure medical data management

It will be based on an open-source imaging IT platform developed by Washington University for neuroimaging data analysis: XNAT

Thanks to its **extensibility**, XNAT can be used to support a wide range of imaging-based projects.

XNAT open source imaging informatics

***Extensible Neuroimaging Archive Toolkit*** *https://www.xnat.org/*

MIRO
Minibeam Radiotherapy

# INFRASTRUCTURE

## XNAT installed on computing center machines @INFN Pisa.

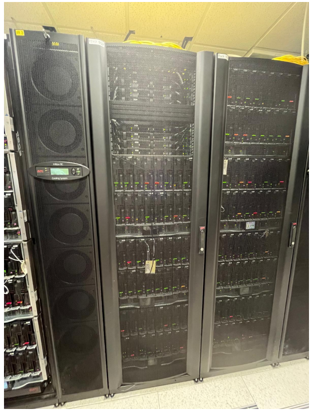

Analysis pipeline processed within the same server hosting the database

Limitations in terms of resource scalability

Interfacing XNAT with HPC Resources

- Architecture setting
- User Access management
- Technical support