Enabling Italian Astrophysics at Scale

Results from the National Centre for HPC, Big Data and Quantum Computing and future perspectives through the Deep Tech Living Lab

ICSC Spokes legacy • DT-Lab vision • 3-year outlook

Agenda

From project results to a three-year strategic perspective

Why astrophysics is an extreme laboratory for digital technologies

What the National Centre and Astro SPOKEs delivered

The strategic legacy for INAF and Italian astrophysics

Why DT-Lab is the natural next step

The next three years: structure, roadmap, and sustainability

Past
Capabilities

Transition
Ecosystem

Future
Living Lab

Astrophysics as an extreme laboratory

Large scientific infrastructures continuously push the limits of computation and data science.

Extreme data volumes

SKA-scale streams, survey archives, simulation ensembles, and heterogeneous data products.

Complex workflows

Pipelines combine HPC, AI/ML, uncertainty quantification, visualization, and distributed services.

Interoperability pressure

Science requires FAIR data, EOSC alignment, HPC, archive services, and cross-facility reuse.

SKA, CTA, Euclid, LSST, and LISA are not only science projects: they are also technology stress tests.

Astrophysics as an extreme laboratory

Large scientific infrastructures continuously push the limits of computation and data science.

Facilities are not only ”telescopes”

New exascale capable laboratories allow to increase dramatically the dynamical range (from cosmological scales down into galaxies)

Crucial for scientific exploitation of a variety of observational data !!!!!

New experiments require new exascale capable laboratories!

Are Astrophysical (HPC) codes ready for that?

Not ONLY Simulations, but also data analysis
Not ONLY HPC, bit also AI and QC

Why AI, HPC, and Quantum Computing matter for the future of astrophysics?

Because the next generation of astrophysical science is defined by scale, complexity, and speed

HPC

Exascale simulations and large survey processing require scalable computing

Essential for multi-scale, high-resolution, and data-intensive astrophysical workflows

The foundation for next-generation numerical modelling and large collaborative infrastructures

AI

Increasingly necessary for detection, classification, anomaly finding, surrogate modelling, and accelerated inference

Critical for extracting science from massive and heterogeneous datasets

A key enabler for more adaptive, automated, and efficient scientific pipelines (**HPC-AI**)

QC

Still exploratory, but strategically relevant for future hybrid workflows

Promising for optimisation, sampling, and new algorithmic approaches

Important to develop competences now, as **HPC–AI–QC** convergence becomes part of the international landscape

Astrophysics is becoming one of the domains where HPC, AI, and Quantum Computing evolve together.

AI in Europe and Italy: opportunity, urgency, and a moving target

AI is no longer only a software topic: it is becoming a strategic infrastructure layer.

Europe is moving fast

The EU is building an integrated AI strategy around infrastructure, data, talent, adoption, and trustworthy AI

AI Factories have become a central instrument of this strategy

In April 2026 the Commission reported 19 AI Factories and 13 AI Factory Antennas already operational or deployed across Europe

This means that AI is being treated as a continental capability, not only as a research topic

Italy is inside this transition

Italy participates directly through IT4LIA, the national AI Factory hosted by CINECA at the Bologna Tecnopolo

IT4LIA builds on Leonardo and on a new AI-optimised system now being deployed

The Italian AI Strategy 2024–2026 also frames AI as a national priority across research, public administration, enterprise, and training

This gives Italy a concrete position inside the new European AI infrastructure

Why this matters for astrophysics

Our field is naturally aligned with this agenda because it **combines extreme data volumes, complex simulations, advanced inference,** and strong workflow requirements

Astrophysics is therefore not just a user of AI, but a domain that can help **shape AI for science**

The opportunity is real, but it is tied to access, competences, software engineering, and long-term infrastructure

The opportunity is clear: Europe and Italy are building AI ecosystems at scale, and astrophysics is well placed to be part of them.

Opportunities and Challenges

The challenge is not only to adopt AI, but to guide its scientific use at scale.

Why AI is even more demanding than classic HPC

AI requires not only compute, but also very large storage, curated datasets, continuous data preparation, software stacks, engineering support, model serving, and user assistance

In this sense, AI often implies an infrastructure ecosystem broader than traditional HPC access alone

The real challenge is therefore organisational and infrastructural as much as algorithmic.

What INAF and astrophysics need

A guided AI strategy rooted in scientific priorities

Long-term competences that remain inside the community

Integration of AI with HPC, data infrastructures, FAIR services, and advanced workflows

A laboratory environment able to connect frontier research, infrastructure, and innovation

The key issue is not whether AI matters, but whether we can build the resources, competences, and leadership needed to use it well: from Data sovereignty to AI SOVEREIGNTY.

Astrophysics as an extreme laboratory

Sustained vs Peak performance

Algorithms : memory intensive vs computation-intensive

We can engineer far more floating point capability onto a chip than can reasonably be used by an application. Today.

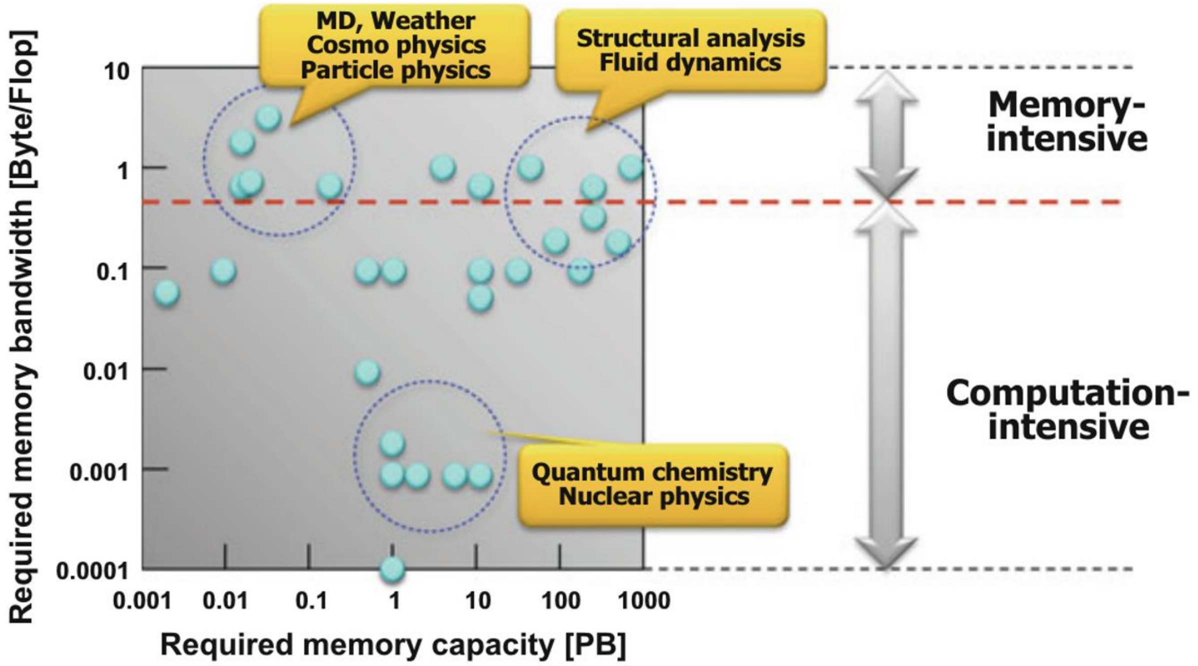

Credits: Hiroaki Kobayashi 2014

Astrophysical workflows are also challenging for any platform

Astrophysics as an extreme laboratory

The Ecological Impact of Computing in Astrophysics

[Simon Portegies Zwart](#)

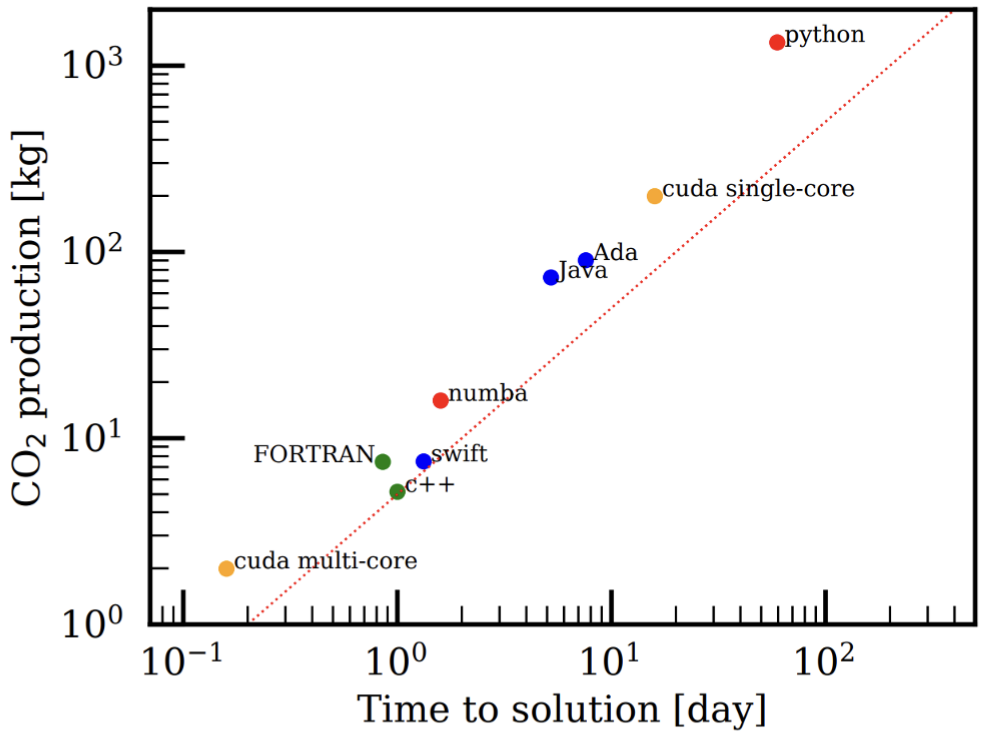

Figure 3: **Here we used the direct N-body code from [23] to measure execution speed and the relative energy efficiency for each programming language from table 3 of [22]. The dotted red curve gives a linear relation between the time-to-solution and carbon footprint ($\sim$ 5 kg CO_2/day). The calculations were performed on a 2.7GHz Intel Xeon E-2176M CPU and NVIDIA Tesla P100 GPU.**

★Gen.IA: iniziativa bottom-up nata in CSN5, supportata da USC-C

Architectural and programming choices matter!

The National Centre and Astro Community role

A national programme that connected infrastructures, methods, and communities.

ICSC mission

Strengthen national capacity in HPC, Big Data, AI, and QC

Connect research, infrastructures, and innovation

Build strategic assets through the PNRR framework

INAF contribution

Participation in multiple spokes:

Spoke 1: Future HPC and Big data

Spoke 2: Fundamental research and Space Economy

Spoke 10: Quantum Computing

Coordination of Spoke 3:
Astrophysics and Cosmos Observations

Astrophysics as a flagship use case for deep technologies

Why this mattered

It created long-term capacity, not just short-term project support

It connected infrastructures, software development, data services, and scientific communities

It produced reusable assets: codes, platforms, archive services, workflows, and expertise

It strengthened INAF's ability to support future large-scale astrophysical infrastructures

What Spoke 3 was designed to build

An integrated programme across computation, data, services, and community.

HPC code enabling
Refactoring and GPU acceleration of scientific codes toward the exascale era.

Algorithms and methodologies
Innovative numerical methods for scaling simulations and data analysis on emerging architectures.

Big data, AI, visualization
Machine learning, Bayesian inference, remote and in-situ visualization for large astrophysical datasets.

FAIR data and archives
Data models, interoperability, distributed archives, storage optimization, and EOSC alignment.

Services, access, training
Scientific hub, cloud/HPC integration, collaborative platforms, dissemination, and community growth.

Selected technical outcomes

The SPOKE 3 flagship program.

EAGER: Evolution of gAlaxies and Galaxy clustErs in high-Resolution cosmological simulations

SLOTH: Shedding Light On dark matter wiTH cosmological simulations

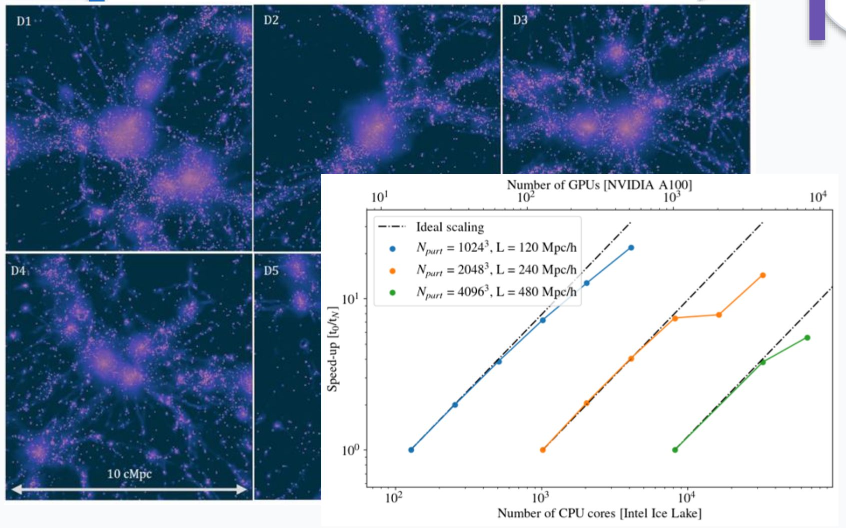

EuMocks: Simulating thousands of Euclid skies

Multi-tracer inference of the first billion years

ISTEDDAS: Revealing the populations of compact remnants in dense stellar systems

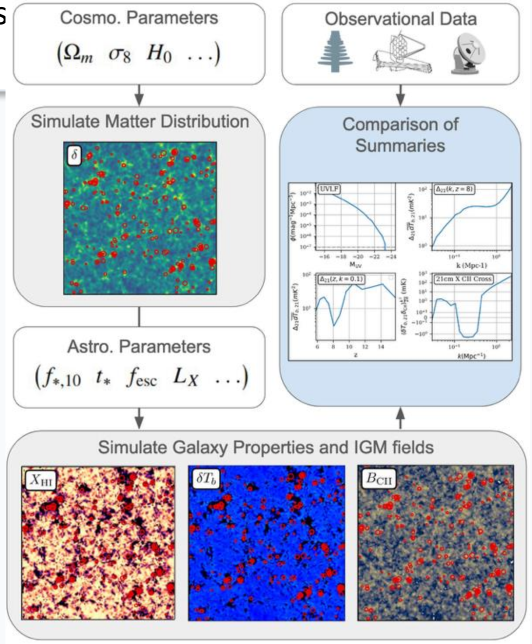

OPAL: simulating the Origins of Planets for ArieL

VisIVO: a framework both for batch and interactive visualization services

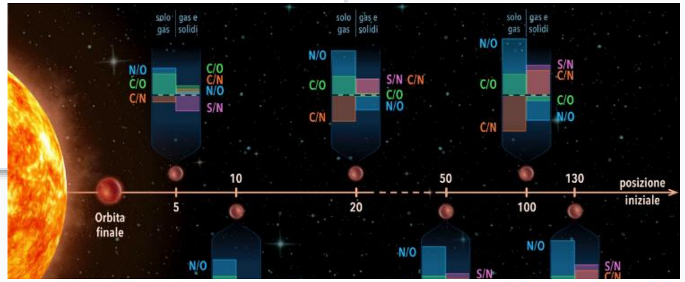

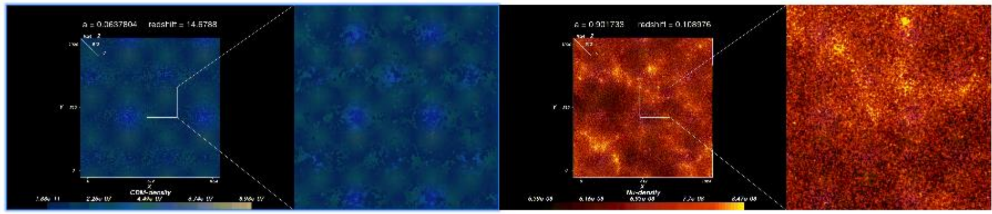

Results in brief (2022–2026)

A high completion rate, a broad technical portfolio, and a visible community effect.

97%
completion of planned activities

36
cascade funding projects

8
innovation grants

>30
cross-sector transfer cases and synergies

150 TB
data-lake space explicitly provisioned

>200
papers and tech reports, the Astronomy & Computing special issue

From astrophysics to wider impact

One of the strongest results was proof of transferability beyond the scientific domain.

The Innovation Grants: when industry involves science

HaMMon / HMMA
Environmental monitoring, seasonal forecasting, AI-supported visualization

HAEQ
Seismic risk assessment and integrated data/AI methods

ATS / TS
Financial time-series analytics and anomaly detection

IDL / IGUC
Interoperable data lakes and data-platform services

Cascade Fundings: **33** projects when science involves industry

Astrophysics has already functioned as a driver for environmental, financial, risk-assessment, and data-infrastructure applications.

The strategic legacy for INAF and A&C community

The project leaves assets that should now be consolidated rather than dispersed.

Competence and people

A larger community of developers, scientists, and users trained around HPC, AI, FAIR data, and advanced workflows.

Methods and platforms

Reusable codes, pipelines, visualization tools, archive services, and collaborative platforms already exist.

National positioning

Astrophysics is better placed to connect frontier science cases with infrastructures, services, and technology transfer.

The key point is that the Centre created assets that remain strategically relevant beyond the project itself.

How can we proceed now?

A strategic choice for continuity, consolidation, and expansion.

The key question now is not what was achieved, but how to turn this legacy into a lasting national capability.

For ICSC

Give continuity to the scientific, technological, and organisational assets generated within the National Centre.

For INAF and Italian astrophysics

Consolidate competences, software, data services, and methods into a laboratory able to support major projects over time.

For research and innovation

Strengthen the pathway from frontier science to industry and society, while keeping basic research as the main driver.

HPC Laboratory @ CINECA

From project legacy to a long-term innovation environment.

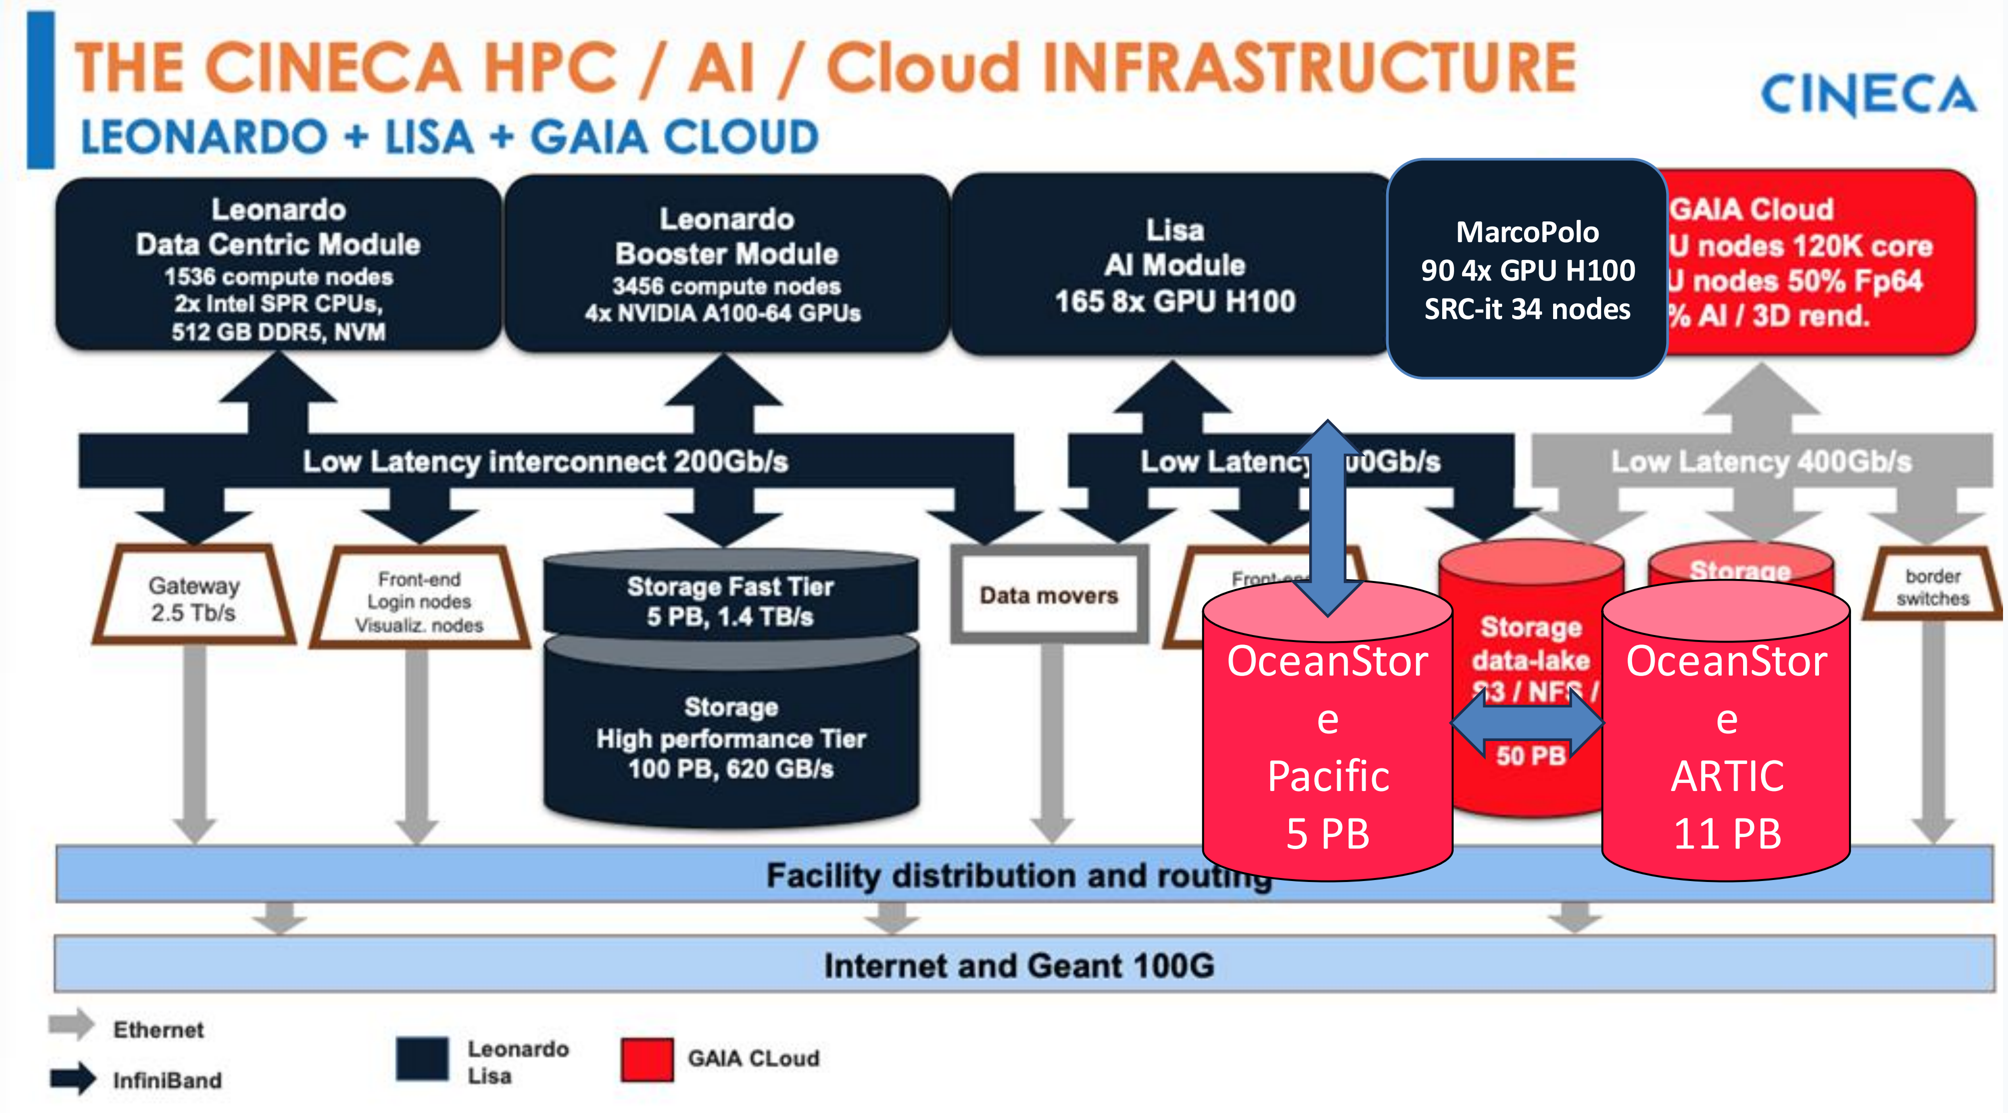

MARCO POLO Super Computer @ CINECA

CPUs

Rack: 4
Nodes: 228 (21 exclusive nodes + 55 nodes)
Model: Lenovo ThinkSystem SD665 V3
Processor: AMD EPYC 9745 (Turin-Zen5c)
Core x node: 128
Total number of cores: 29184
Frequency: 2.3 GHz
Power: 400W
Memory x node 1: 768 GB RAM DDR5 (3GB per core)
Memory x node 2: 26 nodes wirh 1.152 TB DDR5.
Interlink

- 2 × ConnectX-7 NDR200 dual-port x node
- 2 × 100 Gb Ethernet per nodo

HPL: 3,9 PFLOPS
Consumtion: 256 kW (including switch IB core)

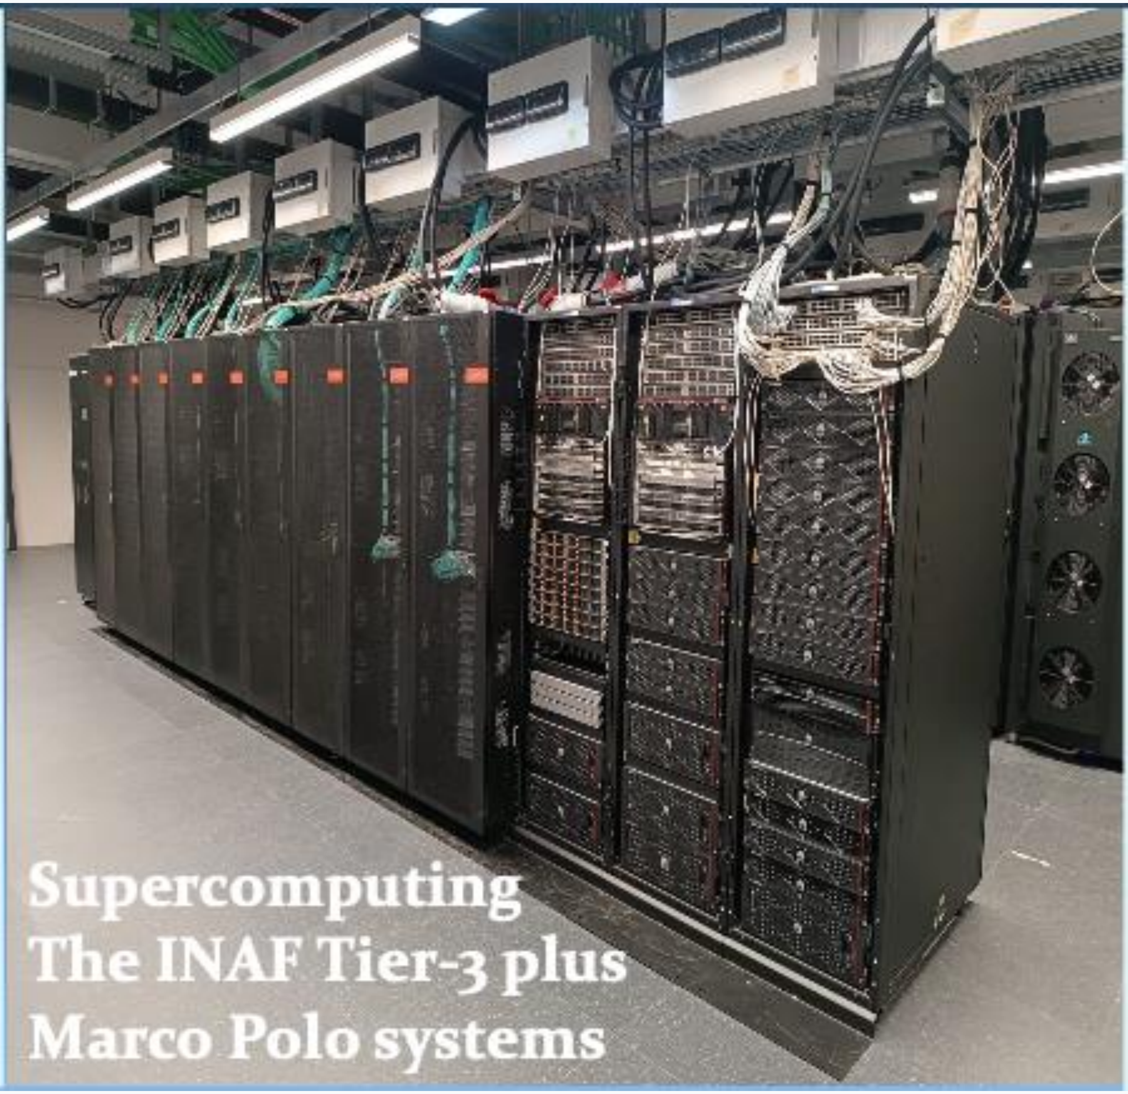

GPUs

Rack: 5
Nodes: 90 (4 exclusive nodes + 30 nodes)
Model : Lenovo ThinkSystem SD650-N V3
Processor : Intel Emerald Rapids 8592+
Cores x node: 64 x 2 sockets
Total number of cores: 218
Frequency: 1.9 GHz
Accelerator: Nvidia H100 SXM5

- 4 GPU per nodo
- Memoria HBM3 per GPU: 80GB
- NVLink 4.0
- Consumption GPU:
  - *HPL*: 700W
  - *Normal operationsi*: 600W

Memory x node: 512GB DDR5 5200 MHz
Interlink:

- 4 × NDR200, each directly linked to GPU
- 2 × 25GbE + 1 × RJ45 GbE x node

Local Storage: 960GB NVMe x node (scratch)
HPL: 15,24 PFLOPS
Consumption: 421 kW (GPU at 700W), 3.317W x node

HPC for INAF and Astrophysics Community?

Because the next generation of astrophysical science is defined by scale, complexity, and speed

ISTITUTO NAZIONALE DI ASTROFISICA INAF

Pleiadi

cluster CPU e GPU distribuiti nelle sedi INAF (OACT,OATS,IRA,OAPA)

fino a ~500k core-hours per progetto

Cineca

accesso ai sistemi di supercalcolo del cineca , oggi Leonardo Booster in futuro Marco Polo

fino a ~125k standard hours per progetto

IA2 Data preservation

a supporto del calcolo

archiviazione e preservazione dati scientifici

tipicamente fino a 20 TB per progetto

La comunità INAF può accedere a risorse di calcolo e archiviazione tramite call competitive periodiche.

HPC for INAF and Astrophysics Community?

Because the next generation of astrophysical science is defined by scale, complexity, and speed

PLEIADI è l'infrastruttura di calcolo distribuita di **INAF**, progettata per supportare simulazioni numeriche, analisi dati e workflow scientifici su larga scala.

Pleiadi

Cluster basati su:

- Intel Xeon E5-2697 v4 (Broadwell): 36 cores per node (2 × 18 cores)
- 128–256 GB RAM per node
- interconnessione Intel Omni-Path 100 Gb/s

Risorse ospitate da OACT (56 nodi) + OATS (60 nodi) + IRA (48 nodi)

Pleiadi-GPU

Cluster basati su:

- IBM POWER9 CPUs 36 cores per node (2 × 18 cores)
- 256GB RAM per node
- NVIDIA Tesla V100 GPUs (16GB)
- interconnessione InfiniBand 100 Gb/s

Risorse ospitate da OACT (10 nodi) + OATS (4 nodi) + OAPA (4 nodi)

A three-year roadmap for the DT-Lab

From setup to consolidation of a long-term deep-tech ecosystem.

Year 1
Build the foundation

Year 2
Connect and accelerate

Year 3
Consolidate and project forward

1

2

3

Focus

Governance and operational units
Core activities in HPC, AI, data, and transfer
Recruitment and first training actions
Integration of Spoke 3 legacy assets

Focus

Cross-domain demonstrators
Stronger integration across operational units
Expanded collaboration with industry and stakeholders
Scalable workflows and interoperable services

Focus

Technology maturation toward higher TRLs
Stable service and innovation environment
Long-term national and European positioning
Sustainability beyond the initial project

Implementation and governance

A distributed living lab with shared coordination and industrial participation.

Governance model

Management Board: strategy and priorities
Coordinator: operational leadership and external representation
Technical Board: scientific and technical guidance
OU coordinators: execution and cross-OU integration

Consortium (tentative)

INAF (leader)
INFN
University of Trieste
University of Tor Vergata
E4 Computer Engineering
IFAB
Koexai
ExactLab

Physical footprint

Distributed activity model
Catania site from year 1
Bologna Tecnopolo site from year 2
Access to INAF, INFN, Leonardo, AI Factory, EuroHPC resources

The five Work Packages

A structure designed to connect technological development, transfer, and human capital.

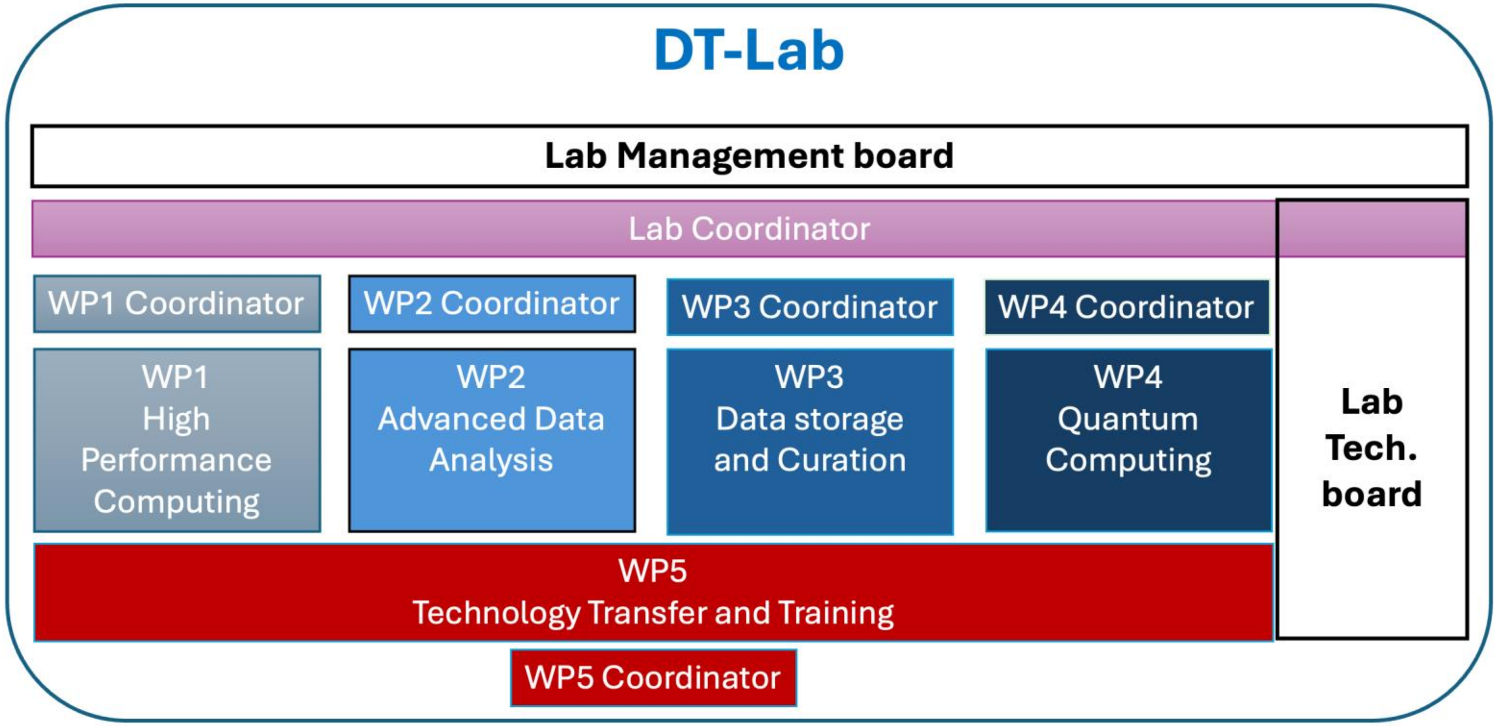

Why this can last beyond the initial project

The proposal is conceived as a long-term deep-tech infrastructure, not a short-lived initiative with a direct IMPACT on main scientific projects.

Funding logic

National programmes and follow-on initiatives
European programmes: EuroHPC, EOSC, Horizon Europe
Industrial collaborations and service-oriented activities

Operational logic

Use of existing INAF and national infrastructures
Open-source software as the default model
Training, PhDs, and technology transfer as structural pillars

Expected outcome

A stable and internationally visible hub
Higher TRLs for selected technologies
A stronger bridge between science, innovation, and society

Take-home messages

From project-based results to a sustainable national deep-tech ecosystem.

Past

Spoke[1,2,3,10] proved that Italy can develop integrated capabilities in HPC, AI, FAIR data, scientific software, and services for **astrophysics**.

Present

The assets already exist: people, codes, archives, platforms, partnerships, and a first track record of technology transfer.

Future

DT-Lab offers a credible 3-year path to consolidate these assets into a long-term innovation environment for INAF and Italian astrophysics.

The opportunity now is not only to report what was done, but to decide how to carry it forward.

Thanks for your attentions

Questions?