AI_INFN
Artificial Intelligence technologies for INFN research

INFN
Istituto Nazionale di Fisica Nucleare

# AI_INFN & CREST: towards a Geo-Distributed Platform for Machine Learning across the INFN Federation

Gioacchino Vino (INFN Bari) — speaker
Lucio Anderlini (INFN Firenze)
Francesco Debiase (INFN Bari)
Rosa Petrini (INFN Firenze)

WMLQ 2026 — Warsaw 2026/09/07-11

# Outline

1. **The platform** — what runs today: architecture, storage, workflows
2. **A second site** — what geo-distribution added, and cross-site offloading
3. **What comes next** — CREST: funded infrastructure, and the constraints attached to it

# INFN and its computing ecosystem

- **INFN**: the Italian national institute for nuclear and subnuclear physics research — ~5,000 people, 20+ sites across Italy
- Operates one of the largest **research computing infrastructures** in Italy
- Tier-1 centre at **CNAF** (Bologna) and several Tier-2 centres for LHC computing
- **INFN Cloud / DataCloud**: federated cloud platform across INFN sites
- Partnership with **CINECA** (Leonardo, EuroHPC pre-exascale system) within the **ICSC** national centre (Italian recovery plan, PNRR)
- This talk: how we build an AI platform on top of this federation

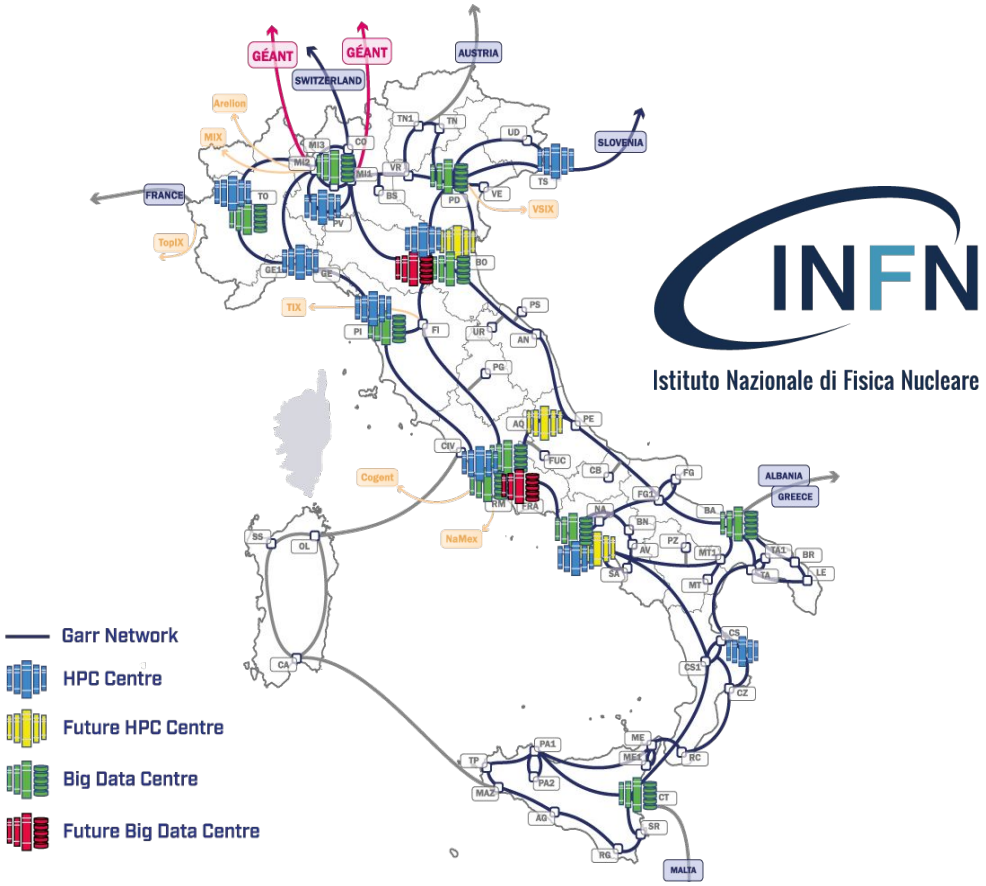


INFN
Istituto Nazionale di Fisica Nucleare

# AI_INFN at a glance

AI_INFN
Artificial Intelligence technologies for INFN research

- **Mission**: enabling AI for physics research across the INFN federation — from classical ML to LLMs — leveraging DataCloud and INFN Cloud solutions
- **Heterogeneous infrastructure**: GPU clusters, HPC, FPGA nodes, cloud, high-performance storage
- **Strategic axes**: scalability, open science, emerging paradigms (e.g. **quantum**)
- **Growing expert community**: collaborations, hackathons, training events
- **Medical physics on the platform** (work of INFN groups): breast-lesion classification on DCE-MRI · deep-learning denoising of low-dose chest CT · 3D dose reconstruction for radiotherapy


# The platform at CNAF

- Deployment at **Cloud@CNAF**, Kubernetes-based orchestration on heterogeneous resources
- **User access**: JupyterHub frontend; AuthN/Z via **INDIGO IAM** (user/admin roles); admins also via SSH
- **Objectives met**: improved hardware orchestration, scalability for AI workloads
- **PNRR contribution**: selected for the INFN-CINECA PoC and for the Flash Simulation flagship project

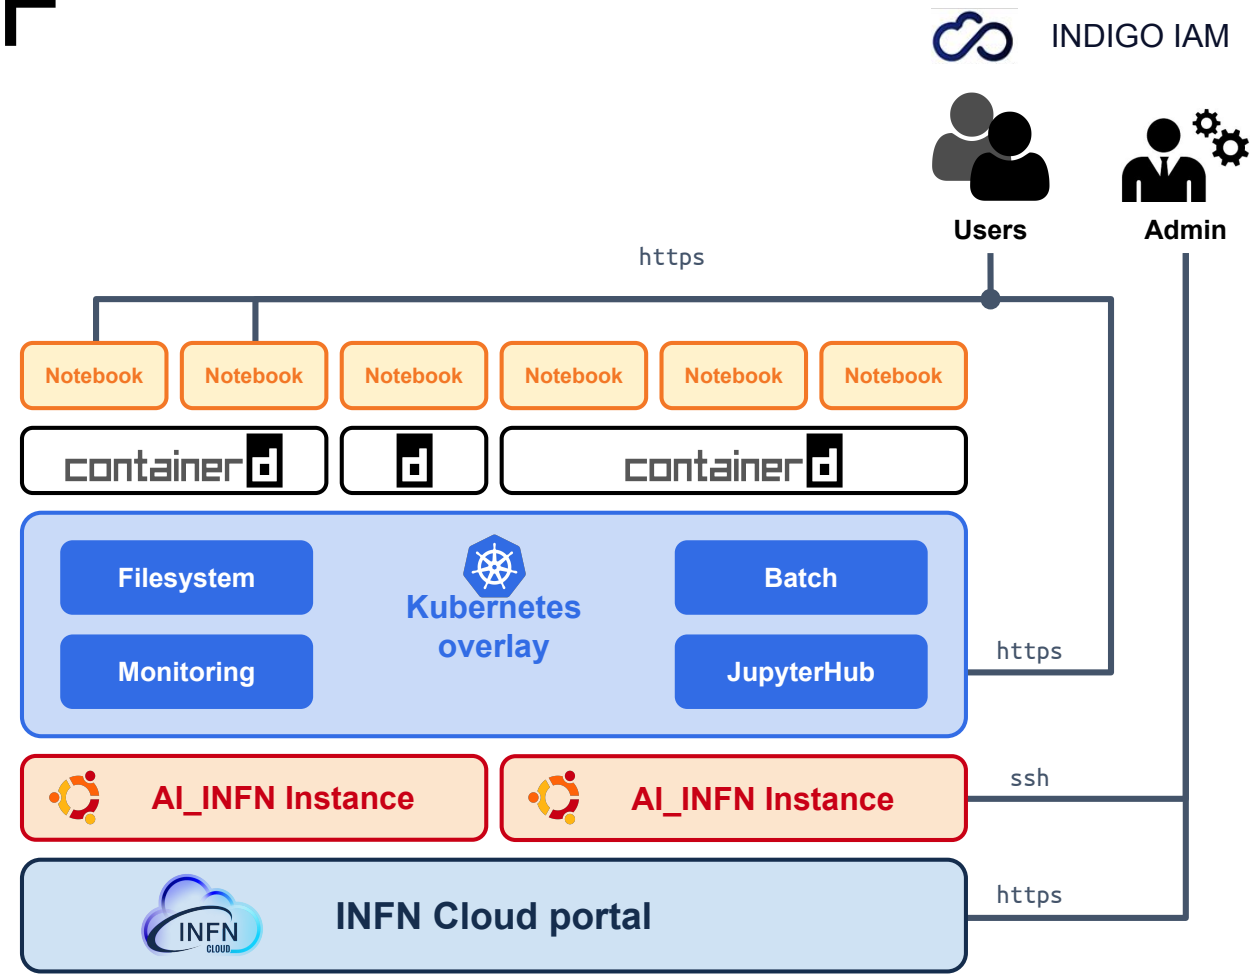

# The platform in numbers

**Objective**:

- Immediate, interactive, web-based access for development;
- Offload heavier tasks to the national infrastructure

**180**
virtual CPUs

**1.2**
TB of memory

**3**
NVIDIA A100
*up to 21 MIG partitions*

**~90**
registered users
*a couple of dozen active*

**small by design**
a notebook and a GPU slice in seconds — the heavy lifting is offloaded

also: 2 × T4 · 5 × RTX 5000
FPGA: AMD Xilinx U55C · Alveo V80 (commissioning)

35 peak concurrent · ~20 weekly · ~28 monthly

# Storage: the CNAF baseline

- NFS-based, **single VM**, single-tenant design
- User home directories only:
  O(10 GB/user), O(1 TB/tenancy)
- **Limits**: fragile backup solution (Borg); suboptimal metadata performance
- Adequate for early operations, not designed to scale across tenants
- The metadata cost, measured — creating a conda environment:
  - on local disk => ~7 s
  - on NFS => < 20 s

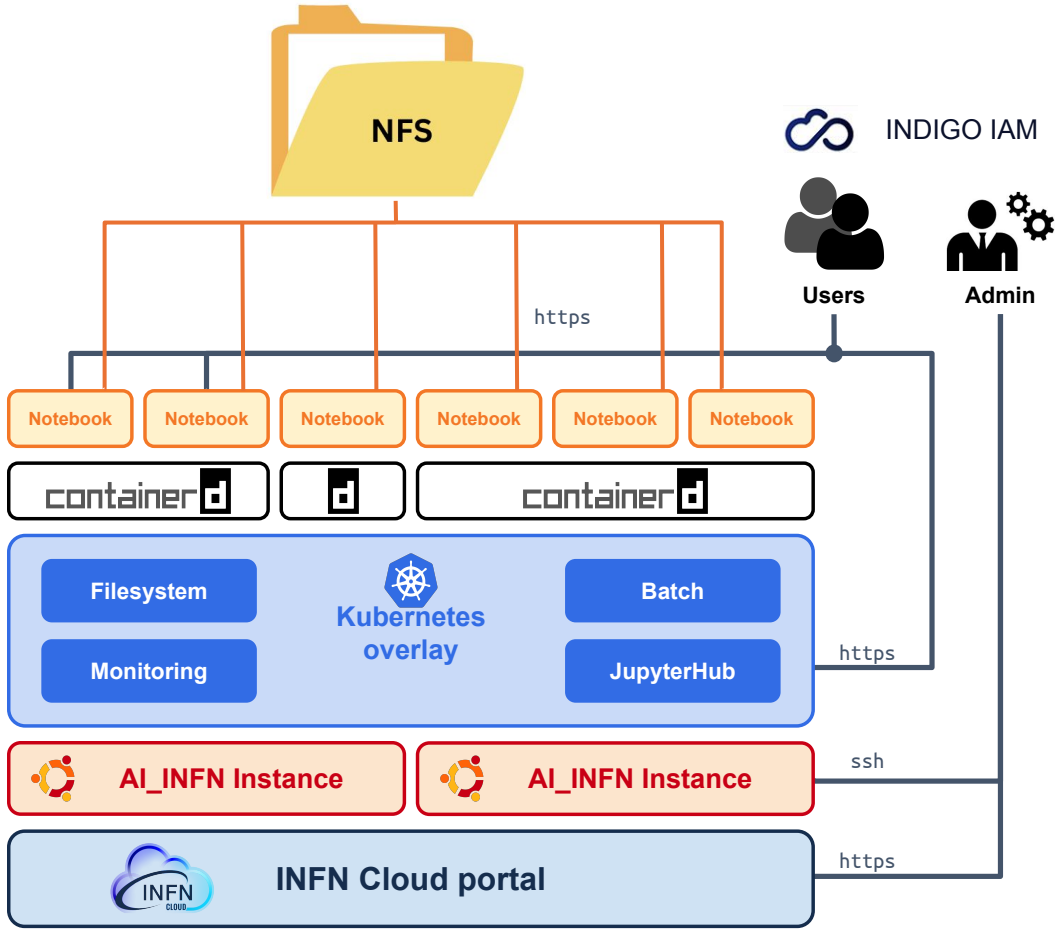

# Storage Evolution: Multi-Tenant NFS

- **Cloud-native**: NFS served by k8s-managed pods, **no VMs to maintain**
- Built-in **per-tenant filesystem segmentation** — no I/O contention
- Automated provisioning at user creation (JupyterHub <-> IAM)
- Designed to **scale across tenancies** and to be **replicable across sites**
- Filesystems can be mounted from **remote sites** via the **SFTP adapter**
- Successfully tested during the 2025 hackathon in Pavia. **Soon in production**

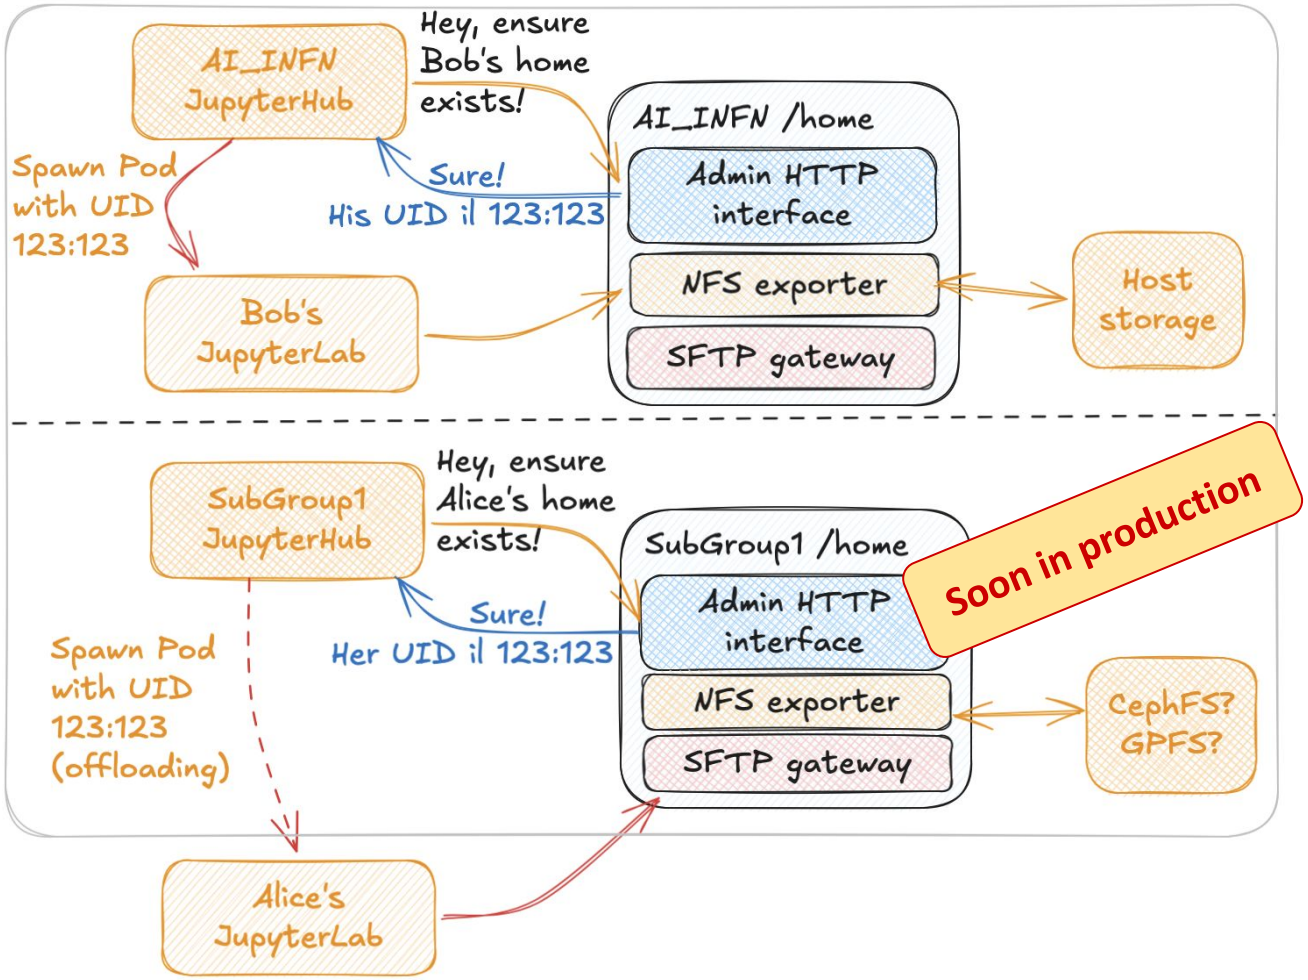

# POSIX access to DataCloud S3 user space

- Users' DataCloud **S3 space mounted POSIX-style** inside the platform
- Implemented with **rclone**, in the **JupyterLab instances**
- The mechanism is not tied to notebooks — it can be extended to other workloads
- **Mature tooling**: small latency cost on many small files, **paid back in reliability**
- Scope: **datasets, not home directories**

# Workflows: bridging batch and cloud

- First implementation: **Snakemake on Kubernetes**
- Same workflow runs **anywhere** — from local JupyterLab to CINECA
- Backend integration: **Kueue** for k8s queueing, **interLink** for offloading to HPC/HTC sites
- **Trade-off**: DAG-only model — no support for ephemeral services within the workflow lifecycle

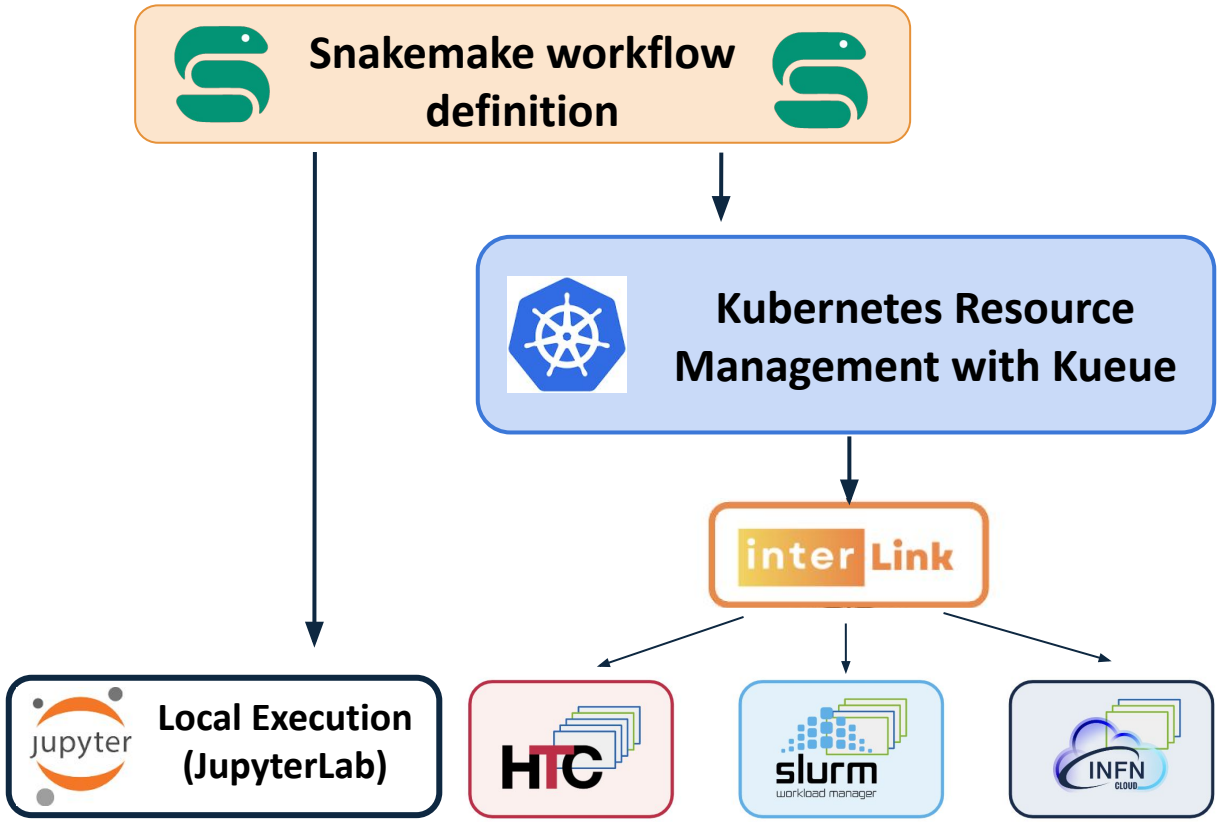

# Argo Workflows: the other trade-off

- Complex scientific workflows need **ephemeral services** — a database, a model registry, monitoring
- **Argo Workflows**: Kubernetes-native, services within the workflow
- **Two models, not a succession** — fine-grained permissions vs standard and maintainable
- **Web-based authentication** — cross-site submission with INFN IAM
- Operational — **first scientific projects being onboarded**

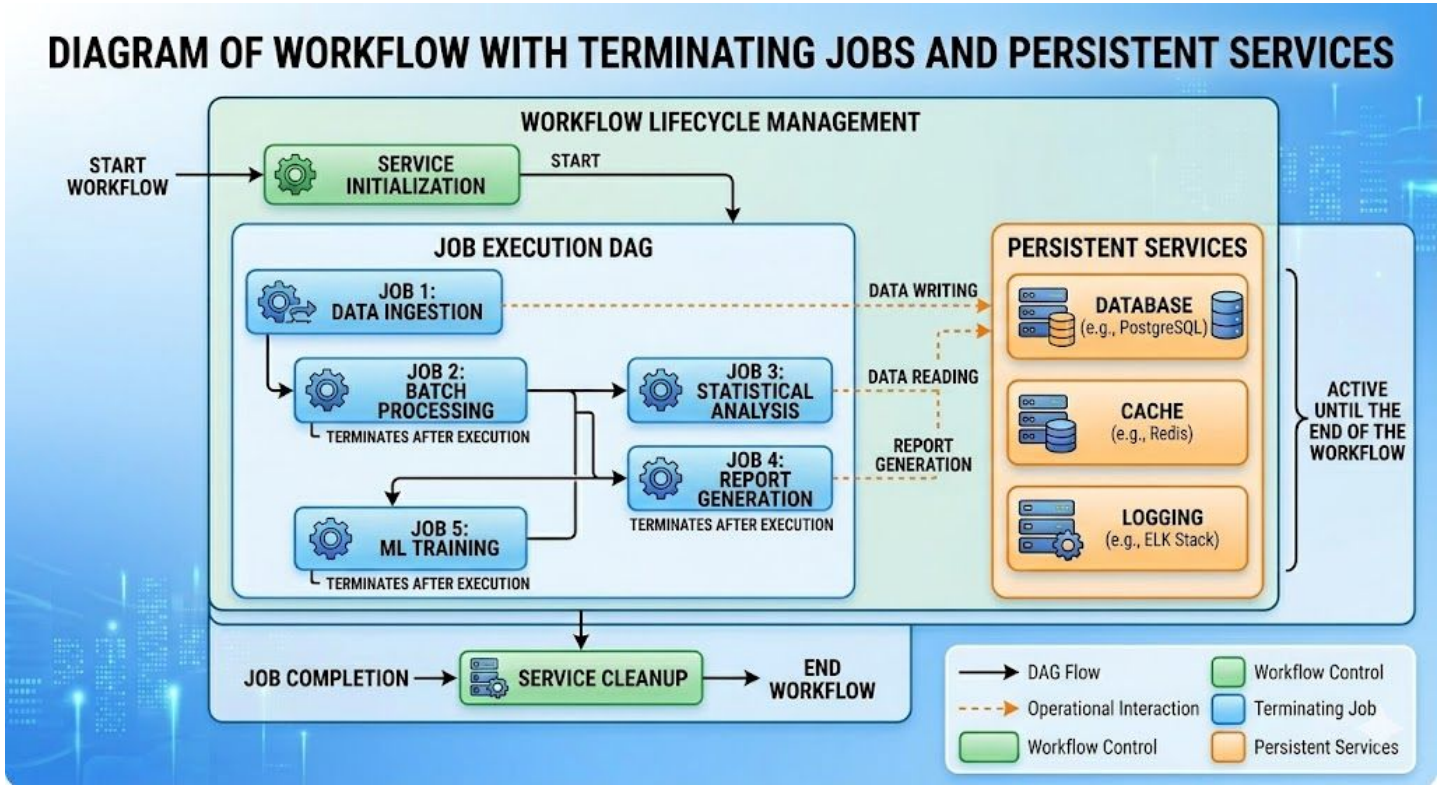

# Extending the cloud with HPC centres

- Federated resources are deeply **diverse**: HTC, HPC, amd64, arm64, GPUs, FPGAs …
- The platform offloads tasks to **national sites** via **interLink**: a plugin-based Virtual Kubelet provider
- **One submission interface**: you submit to **Kueue** or to **Argo**; interLink translates the job for **HTCondor** or for **Slurm** — no need to learn each site's batch system

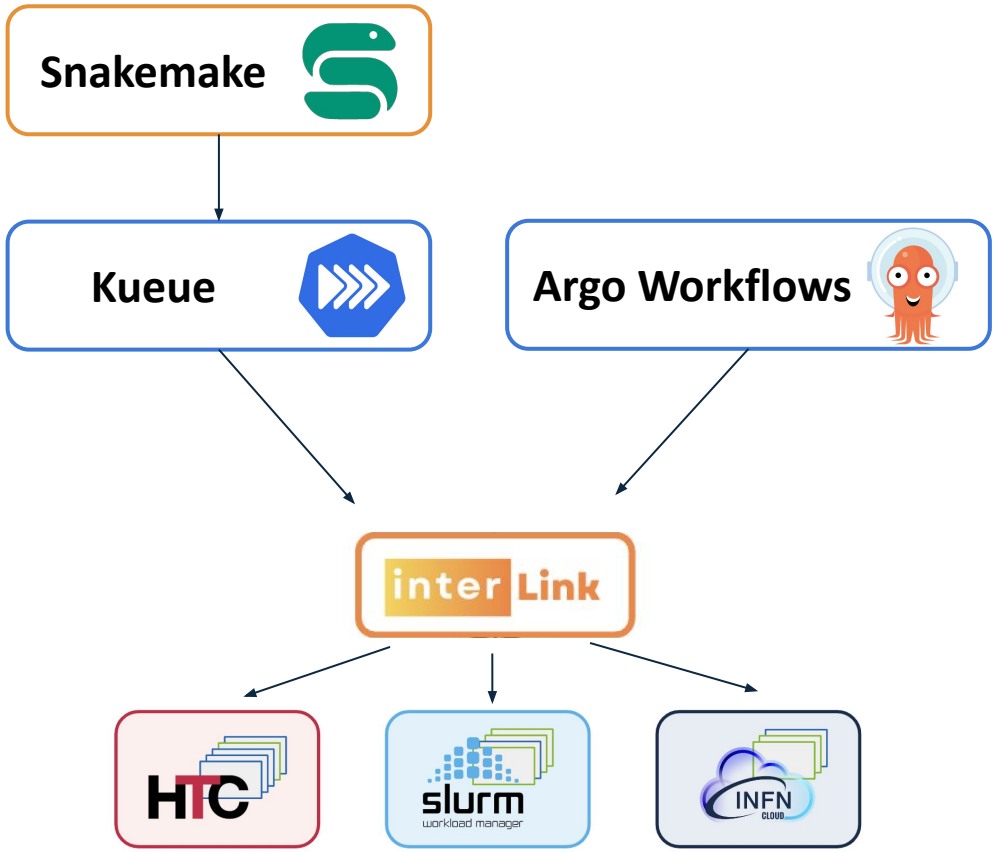

# Offloading in action: Tier-1 and Leonardo

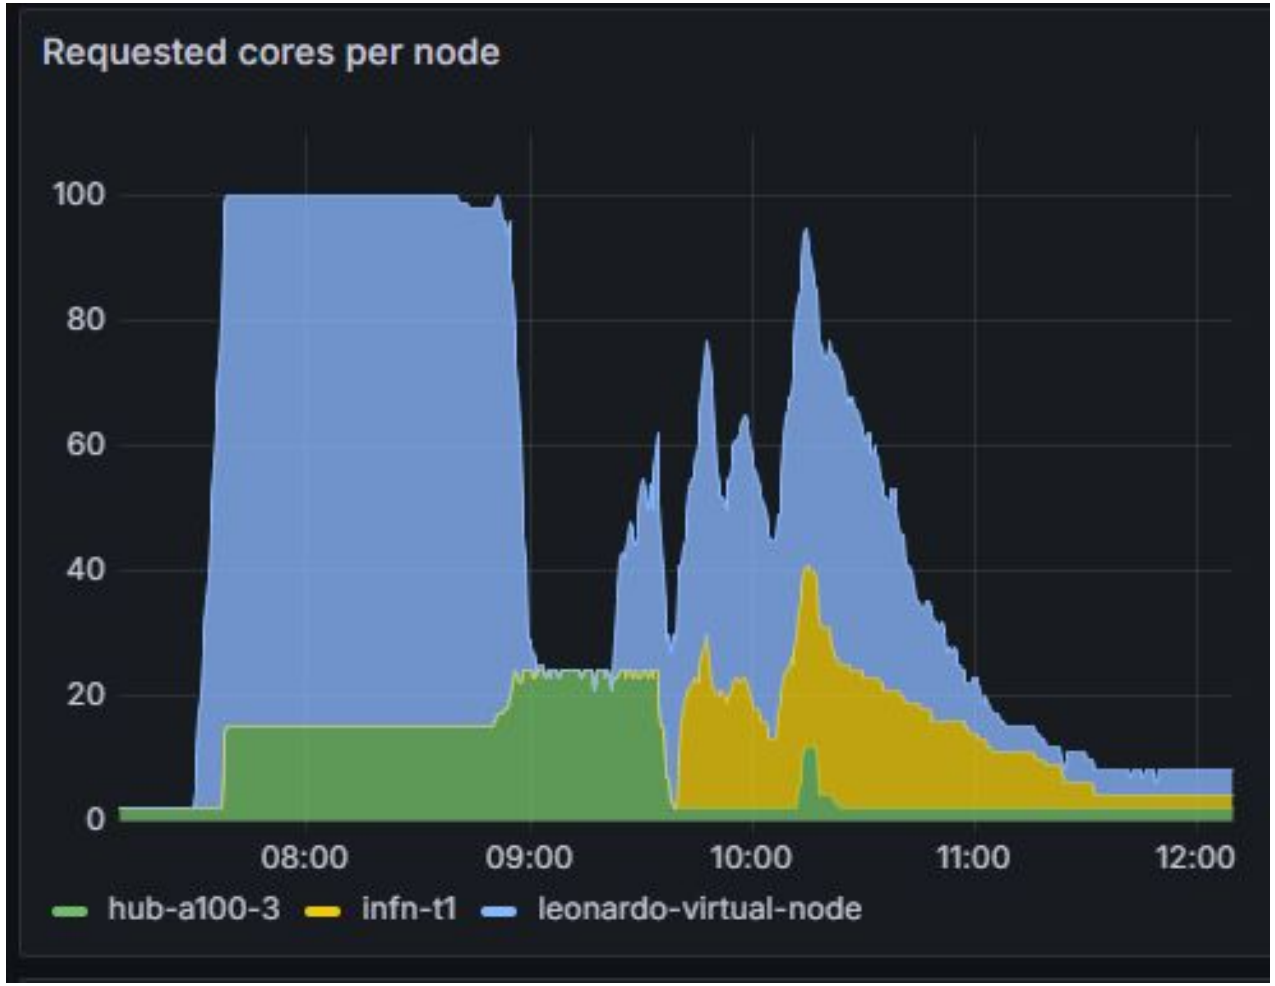

# Federating with ReCaS-Bari

- **Goal 1**: validate **replicability** and integration of the platform design across INFN data centres
- **Goal 2**: enable **cross-site workload offloading**
- Testbed sites: Cloud@CNAF and ReCaS-Bari (**INFN**)
- Offloading: Leonardo (**CINECA**) and HPC Bubble (**INFN Padova**)
- Validation refined small design details → a **more general, more replicable** design

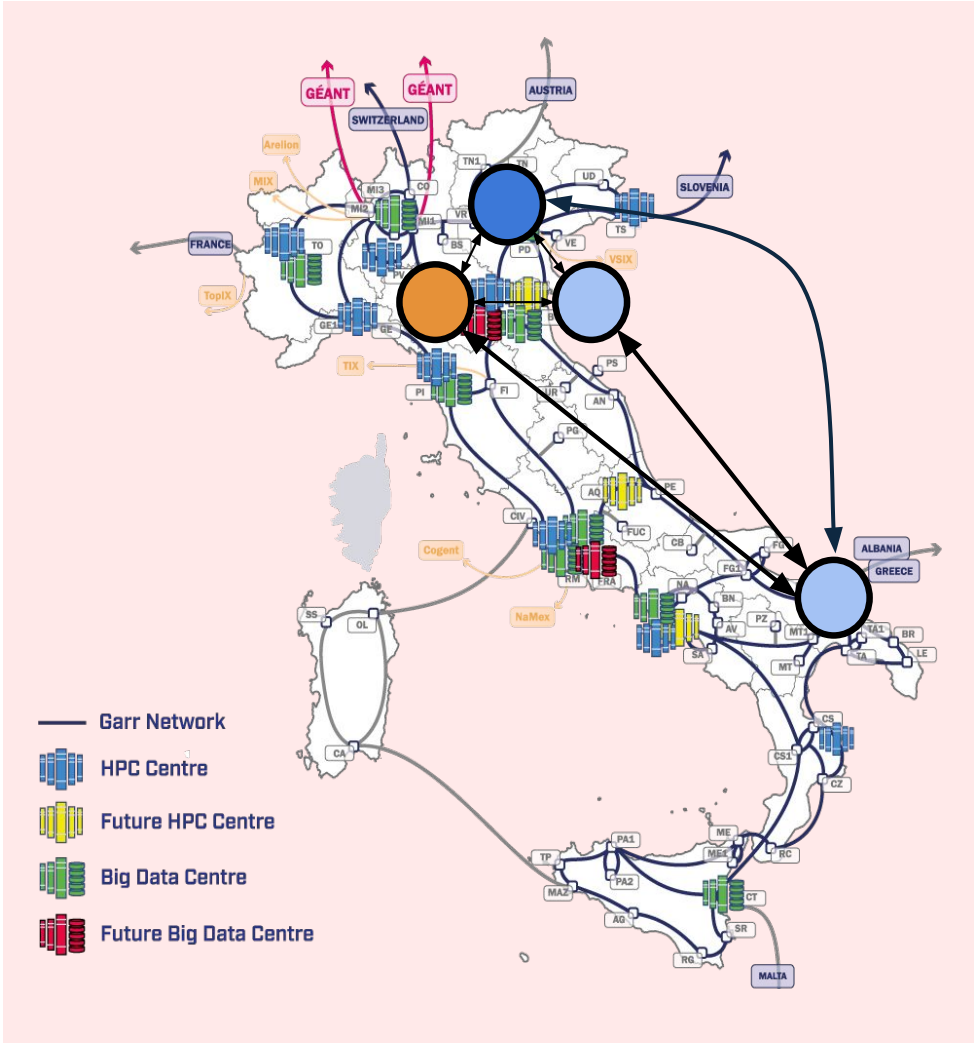

# ReCaS-Bari at a glance

- Data centre operated jointly by **INFN and University of Bari**
- Role within INFN: WLCG Tier-2, multi-experiment site, node of the INFN DataCloud federation
- Compute: HTCondor and OpenStack — **50k+ cores** · Storage: GPFS and Ceph — **23+ PB**
- Production **Kubernetes cluster**: 15+ nodes, **50+ NVIDIA GPUs**, 100+ TB local SSD, direct Ceph and GPFS access
- Established expertise in distributed computing → natural candidate for federation

# Home backup: restic in production

- At CNAF, Borg will be retired in the coming months, together with the storage service it runs on
- restic + backrest — **tested at CNAF** and **in production at ReCaS-Bari**
- backup is already **cross-site**

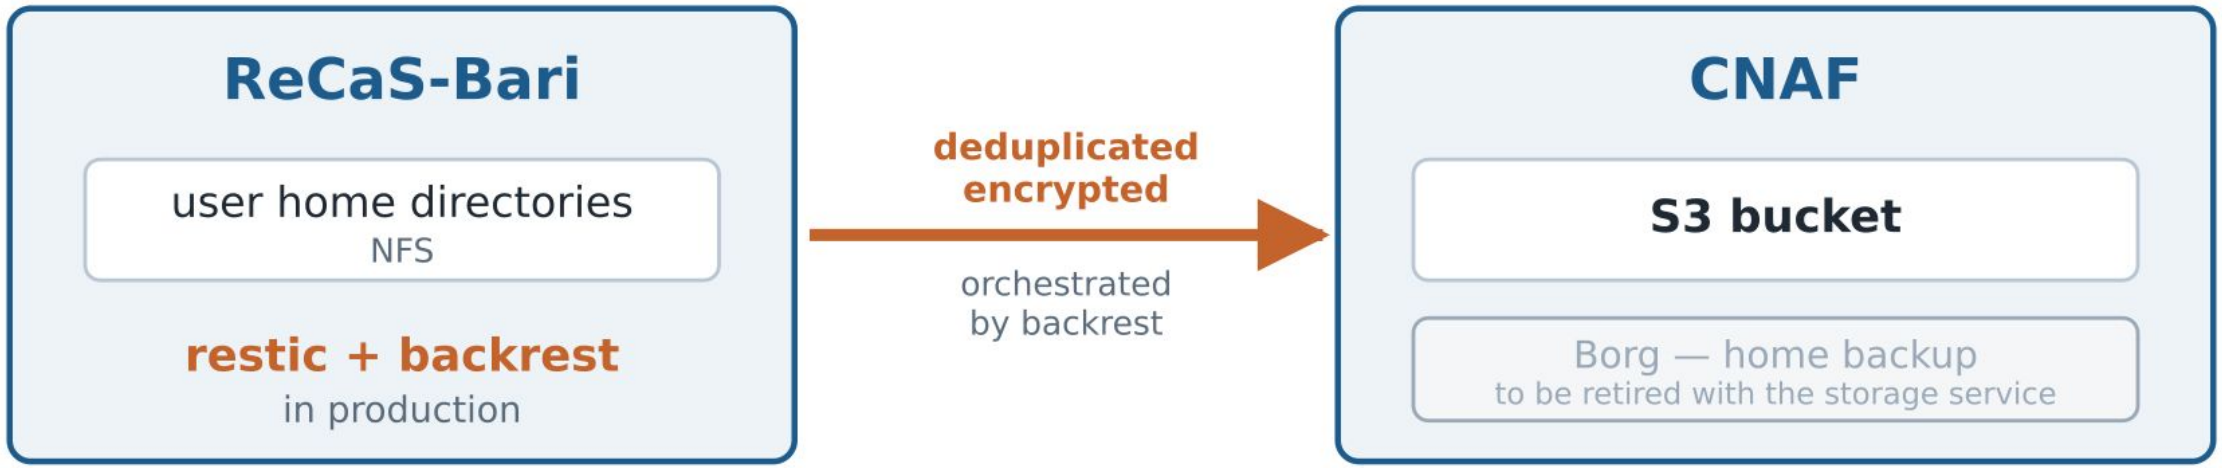

# Cross-site workload offloading

- **Vision**: a job submitted to AI_INFN — from anywhere — can consume resources at remote sites (Tier-1, ReCaS-Bari, Padova, CINECA)
- **Validated from CNAF**: Argo Workflows submitting pods via interLink — to the Tier-1, to Leonardo and to the HPC Bubble in Padova
- **At ReCaS-Bari**: platform and storage in place; the **Virtual Kubelet node is still being commissioned**
- Once that node is in place, Bari becomes a **second dispatch point**, not only a target

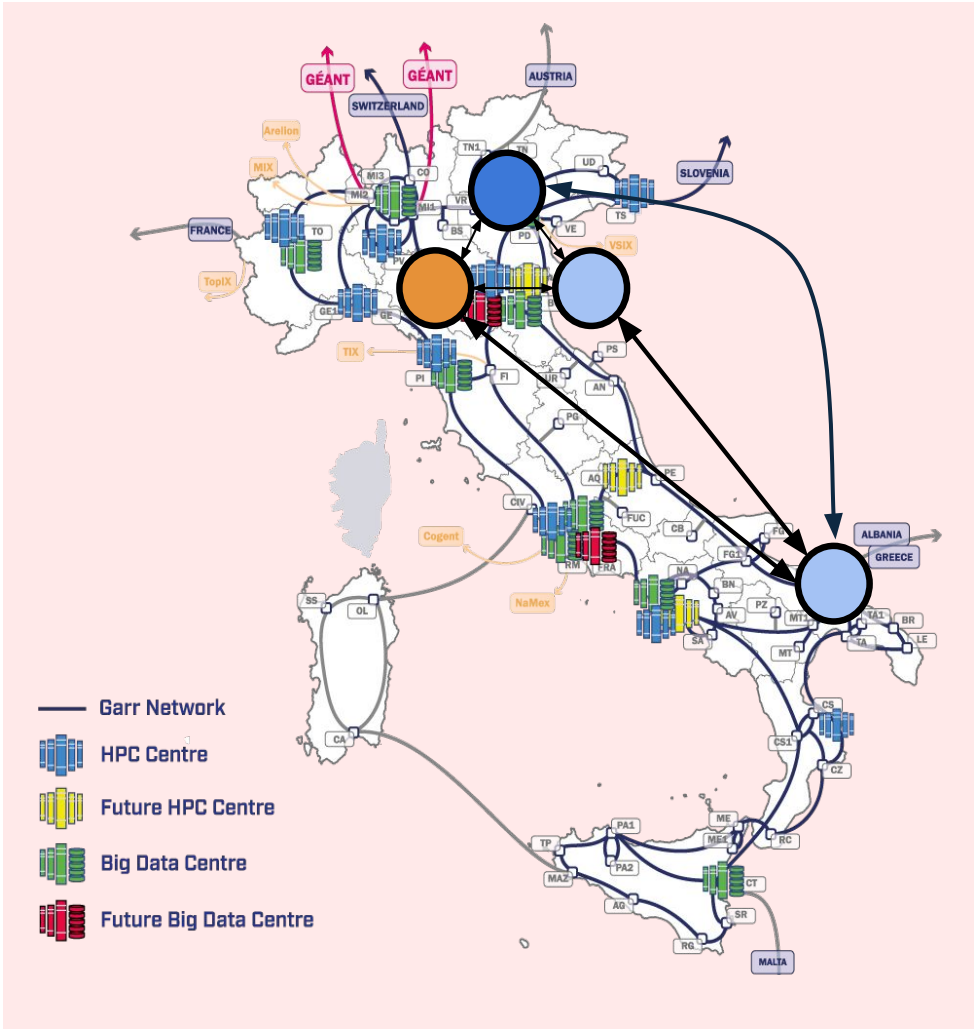

# Monitoring & user dashboards

- Observability stack: **Prometheus + Grafana**, Grafana integrated with INFN IAM
- **Per-user dashboards**: allocated/used resources, Kueue job status, JupyterLab pods
- Cluster health monitoring; alerting via Microsoft Teams


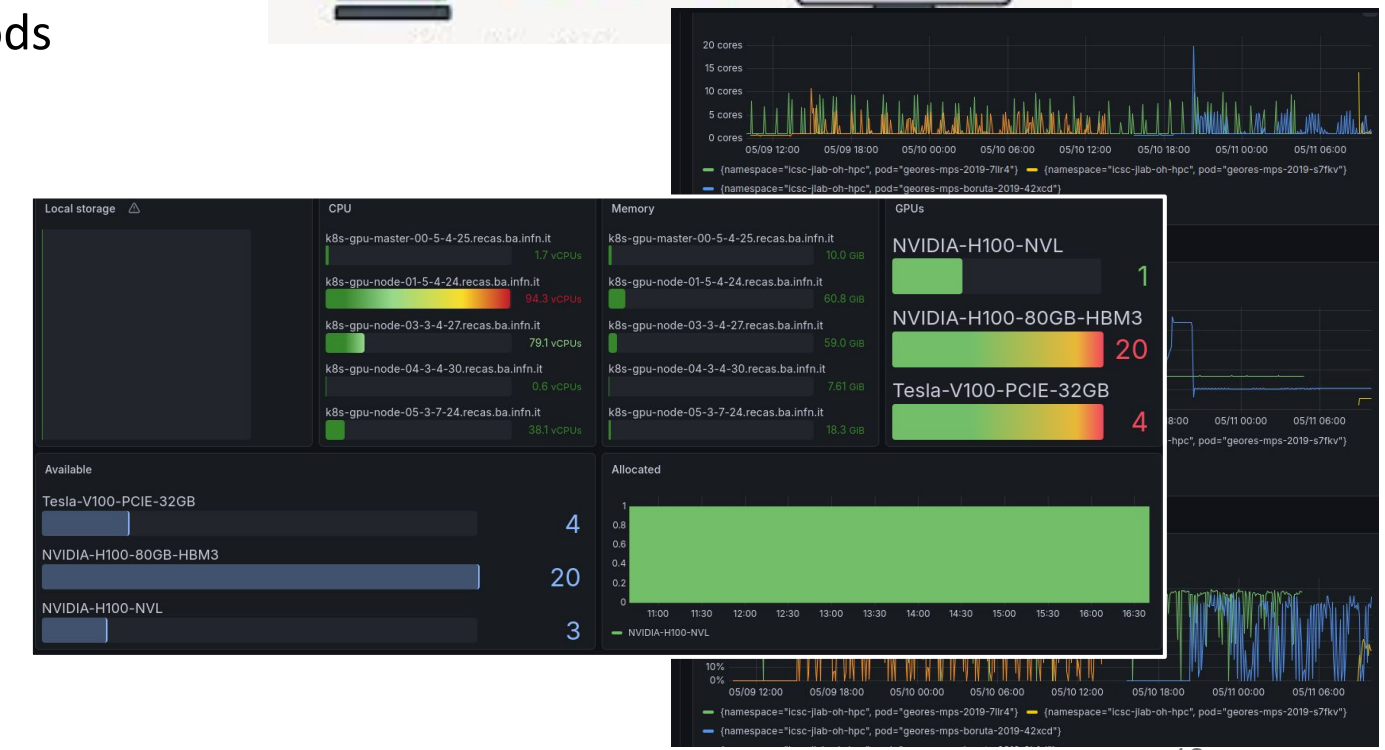

# Kubernetes accounting with AUDITOR

- Last missing piece: accounting of Kubernetes resource usage
- **AUDITOR** (ALU Freiburg): modular accounting toolbox
- Records **benchmark-normalized** with HEPScore23 — a core-hour means the same thing on different hardware
- Per-tenant **views** on shared infrastructure; the same records feed **APEL/EGI**
- Status: **in progress** — collector deployment at ReCaS-Bari

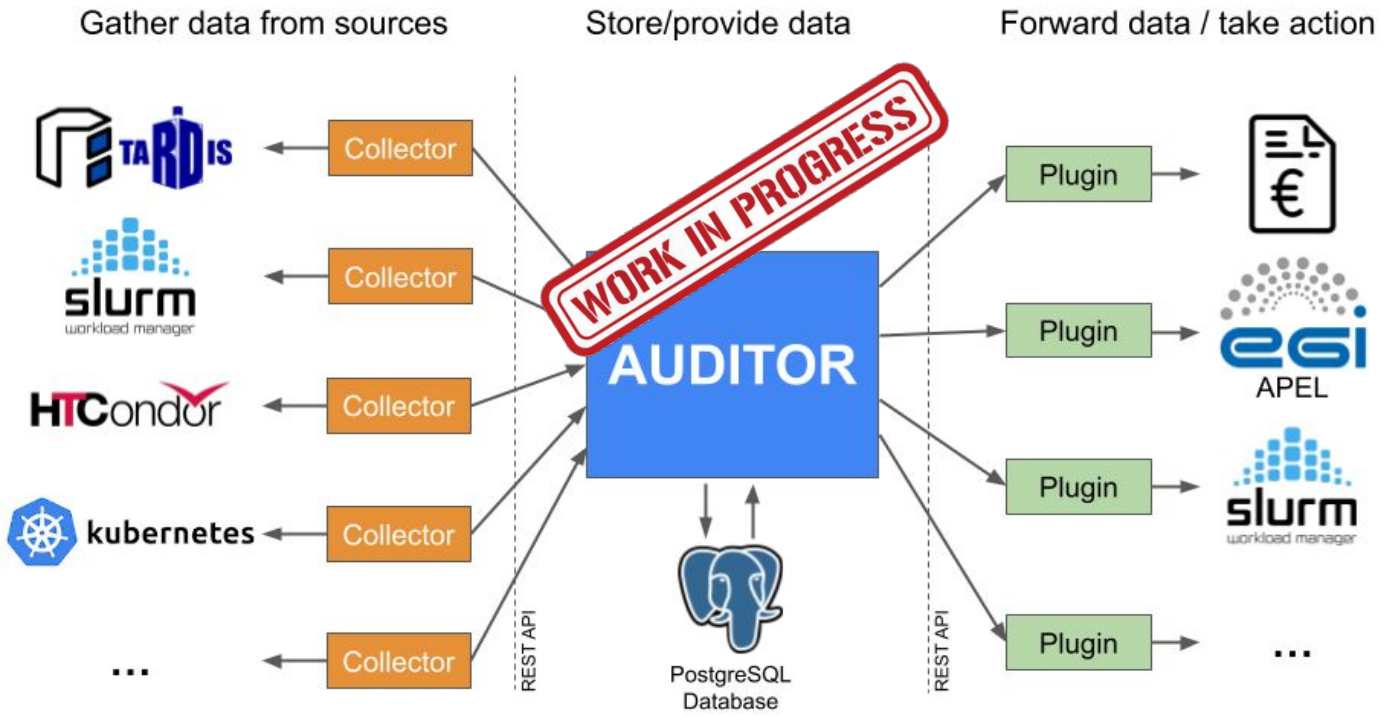

# Lessons learned

- **Replicability validated**: (almost) everything built at CNAF was reproduced at ReCaS-Bari, in production
- **The validation path has settled**: new components are exercised at Bari first, then propagate
- **Cross-site offloading** transparent at the user level
- **Delivered since spring**: S3 POSIX access · home backup at Bari · offloading validated from CNAF


- **Kubernetes Accounting**: Auditor
- **Next**: once all components are in place at ReCaS-Bari, begin **moving part of the workloads there**
- **The limit**: all of this runs on a **prototype** infrastructure — resources granted case by case or opportunistic

# What comes next: the CREST programme

- The infrastructure under the AI_INFN platform is **being upgraded right now**
- **CREST** — Computing Research infrastructure for Engineering, Science and Technology
- **15.5 M€** grant, Italian Ministry of University and Research, **ERDF National Programme for Research 2021-2027**
- **36 months**, started **June 2026** · 6 institutions · 8 work packages
- Five data centres: **CNAF, Bari, Naples, Catania, Cagliari** — including **both sites in this talk**

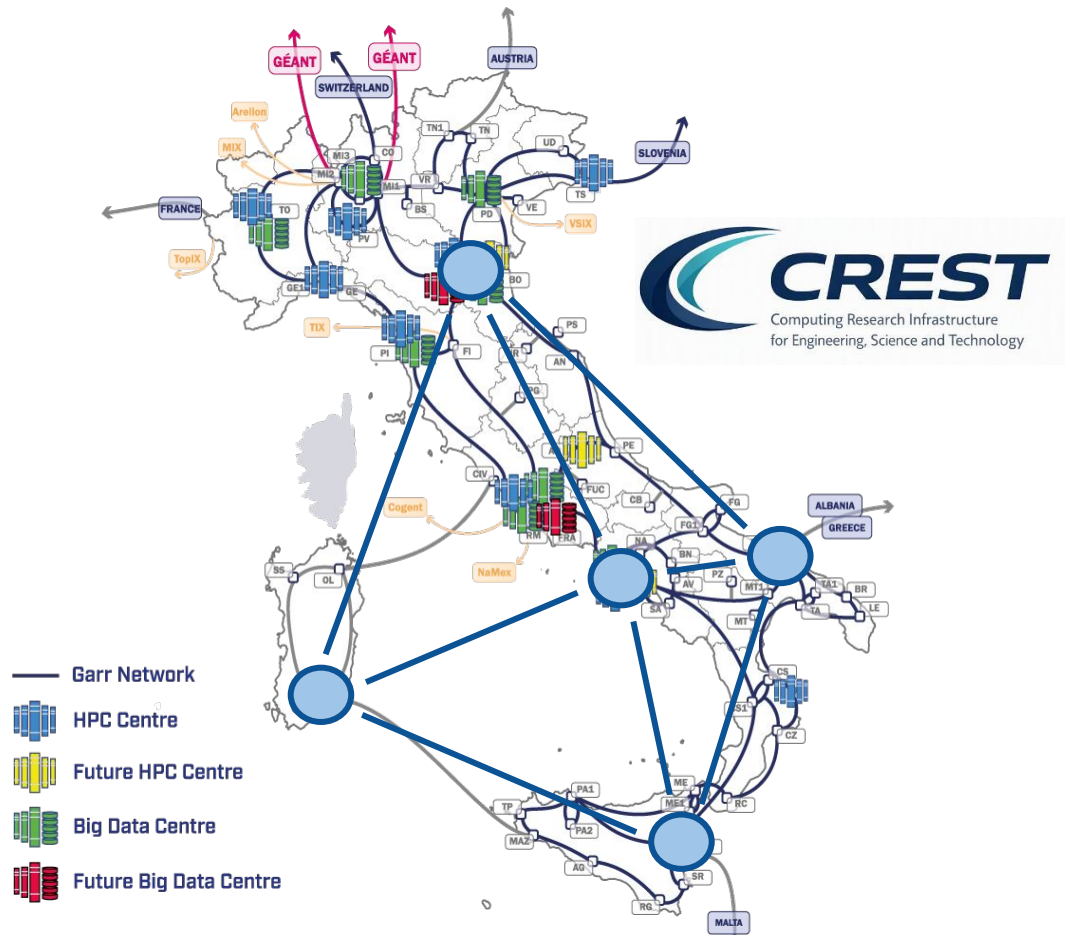


COESIONE ITALIA 21-27
RICERCA, INNOVAZIONE E COMPETITIVITÀ
per la transizione verde e digitale
Cofinanziato dall'Unione europea
Ministero delle Imprese e del Made in Italy
MUR
INFN
Istituto Nazionale di Fisica Nucleare

# What the programme builds

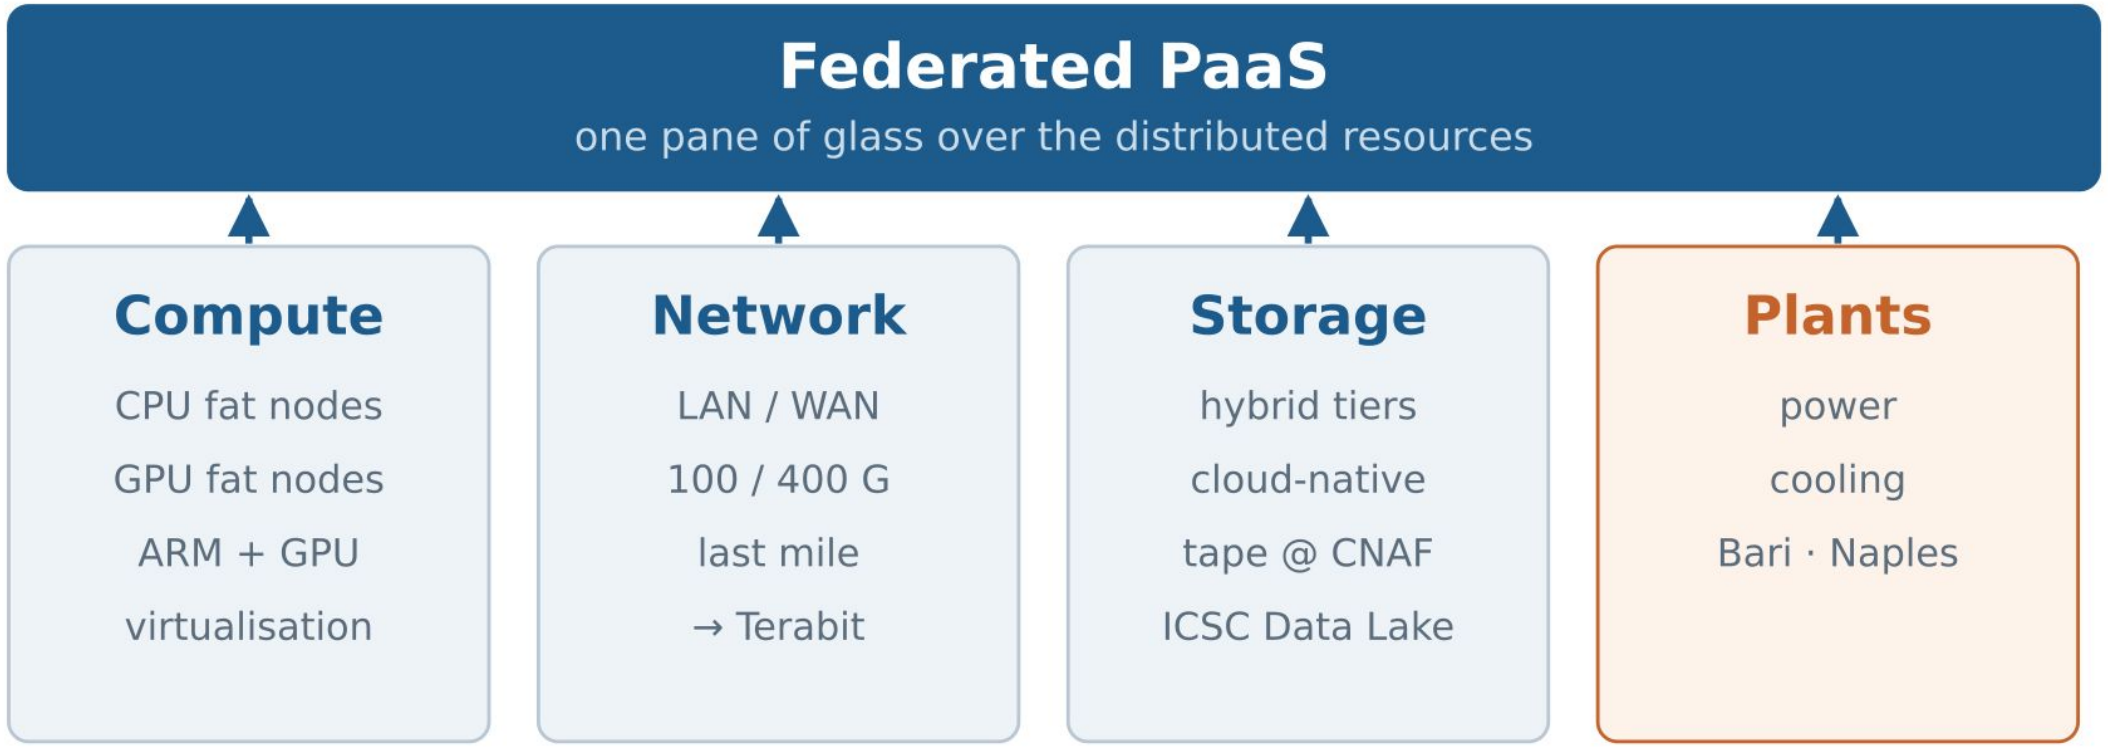


- Plants are on the critical path — air cooling alone is no longer sufficient
- Without the federation layer: five clusters, not one infrastructure

# The AI-relevant hardware

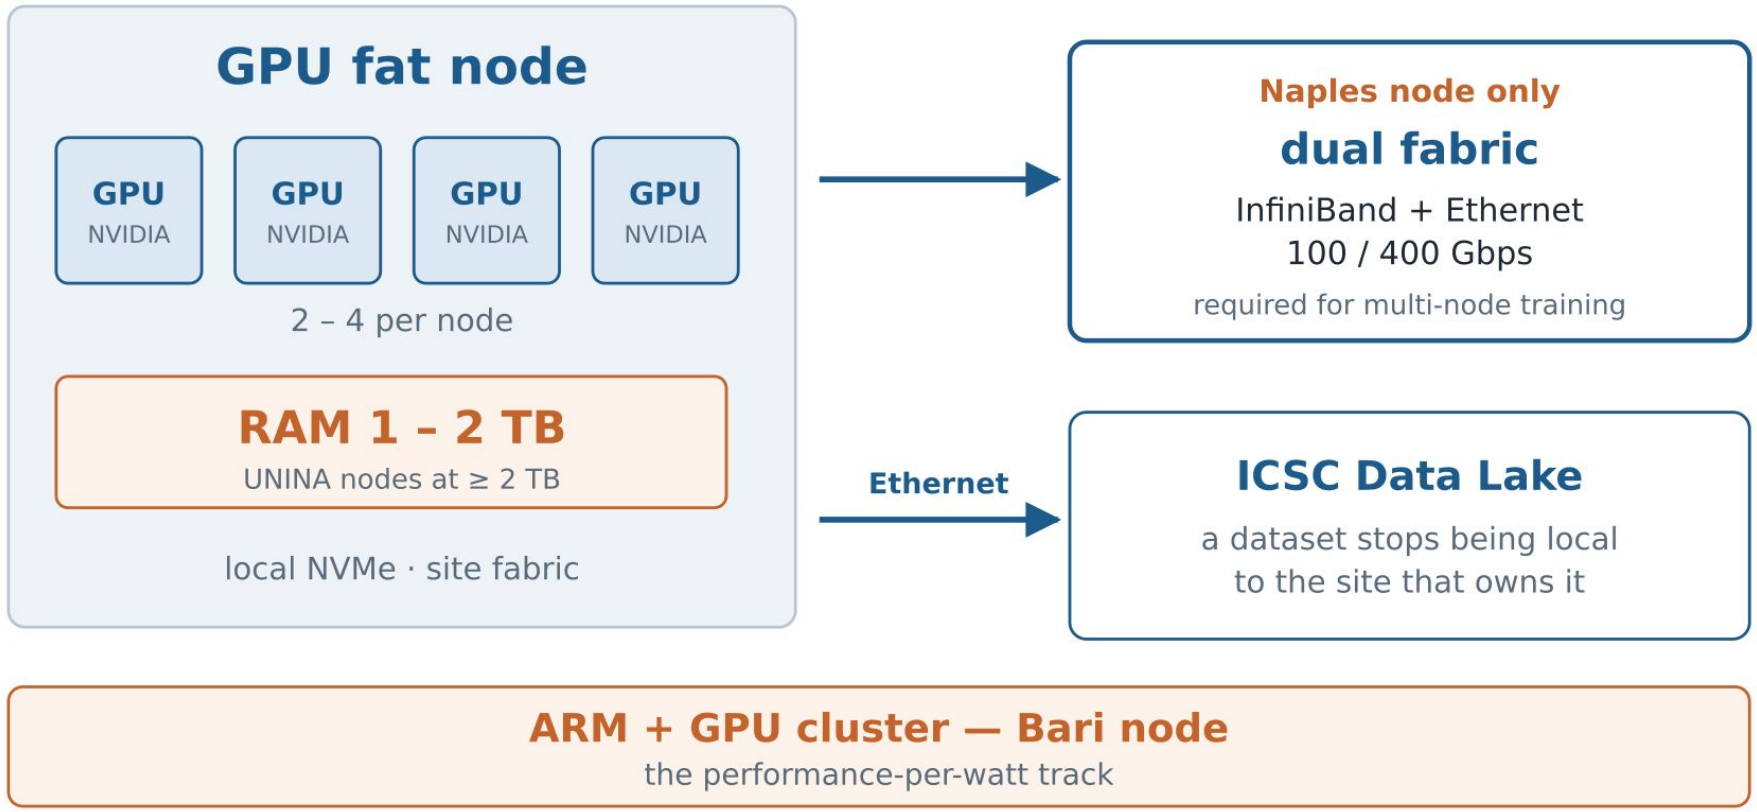

# The federation

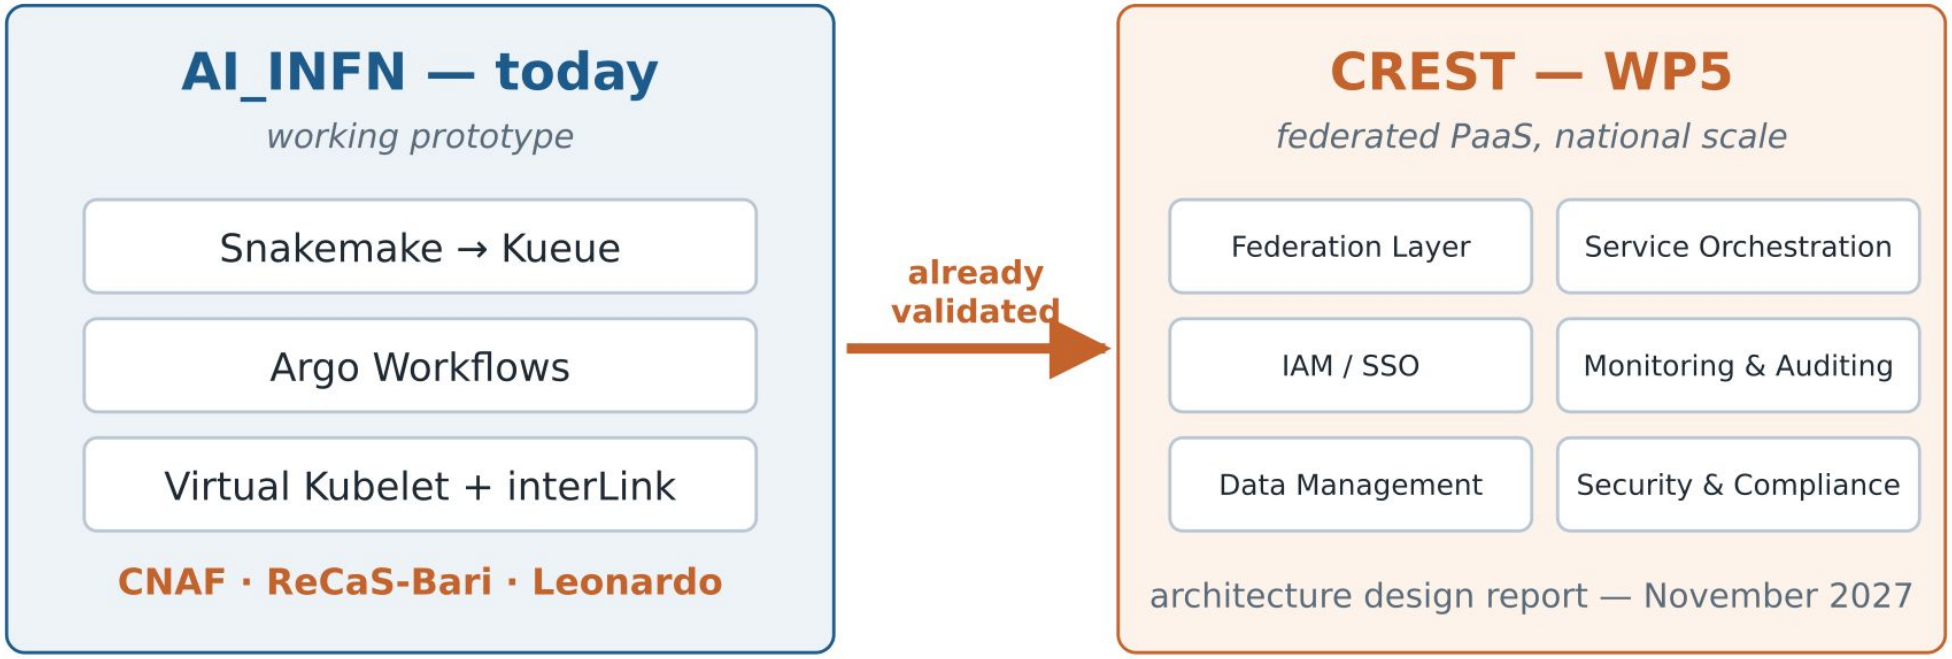


*the placement layer CREST has to build is the one we already prototyped*

- The open problems are: **identity across sites, workload placement and usage accounting**
- Same lineage underneath both: **DataCloud-INFN**. CREST must build a federated AI platform; AI_INFN already is one

# Sensitive data and compliance

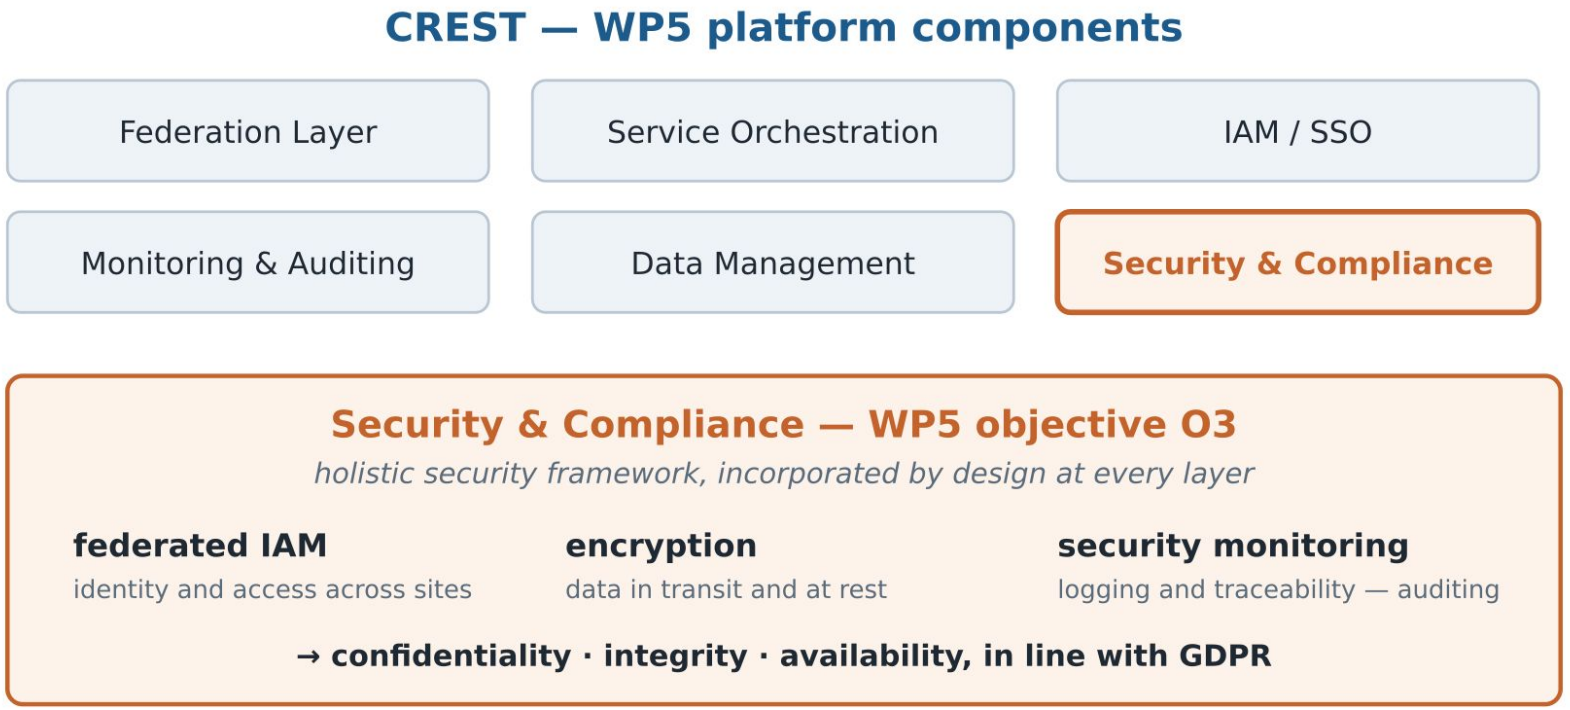


- A prerequisite for hosting **medical imaging or clinical data**
- Multi-tenancy is the enabling piece — the **per-tenant design** shown earlier
- **A design requirement today, not a certification**

# Energy monitoring

- Energy monitoring: a **contractual deliverable, deadline** (May 2028)
- The **ARM+GPU** track: capacity has to grow under a **fixed power ceiling**
- Environmental requirements must be assessed and stated **inside the tender documents**
- For us the real engineering problem will be FLOPS per watt, per euro, per square metre

# Award dates, not delivery dates

| **Contractual milestone** | **Deadline** |
|---|---|
| Federated architecture design report | November 2027 |
| Plant upgrades Bari / Naples awarded | November 2027 |
| CPU/GPU servers purchased | May 2028 |
| Energy-efficiency monitoring | May 2028 |
| Network systems purchased | September 2028 |

- These are **award and completion milestones under the grant agreement**. Between award and production there is delivery, installation and acceptance testing
- **No hardware delivery dates**: the procurement has not started

# The constraints that come with the funding

- The money is **European Regional Development Fund** — the same instrument exists in every cohesion country
- The rule shapes the map: **over 85% of the funds must be spent in the South;** the programme allocates **90%**, four of the five nodes
- **Funding rules shape engineering:**
    - milestones are **award dates, not delivery dates**
    - **energy monitoring is contractual** — designed in from the start
    - at least one **industrial use case** — an obligation we report on

# Conclusions

- **AI_INFN**: a production AI/ML platform serving the INFN community
- CNAF design **proven replicable** — replicated at ReCaS-Bari
- First concrete prototype of **geo-distributed AI computing within INFN**
- With **CREST**, that prototype now has a **funded production target** — on the very sites we have been federating

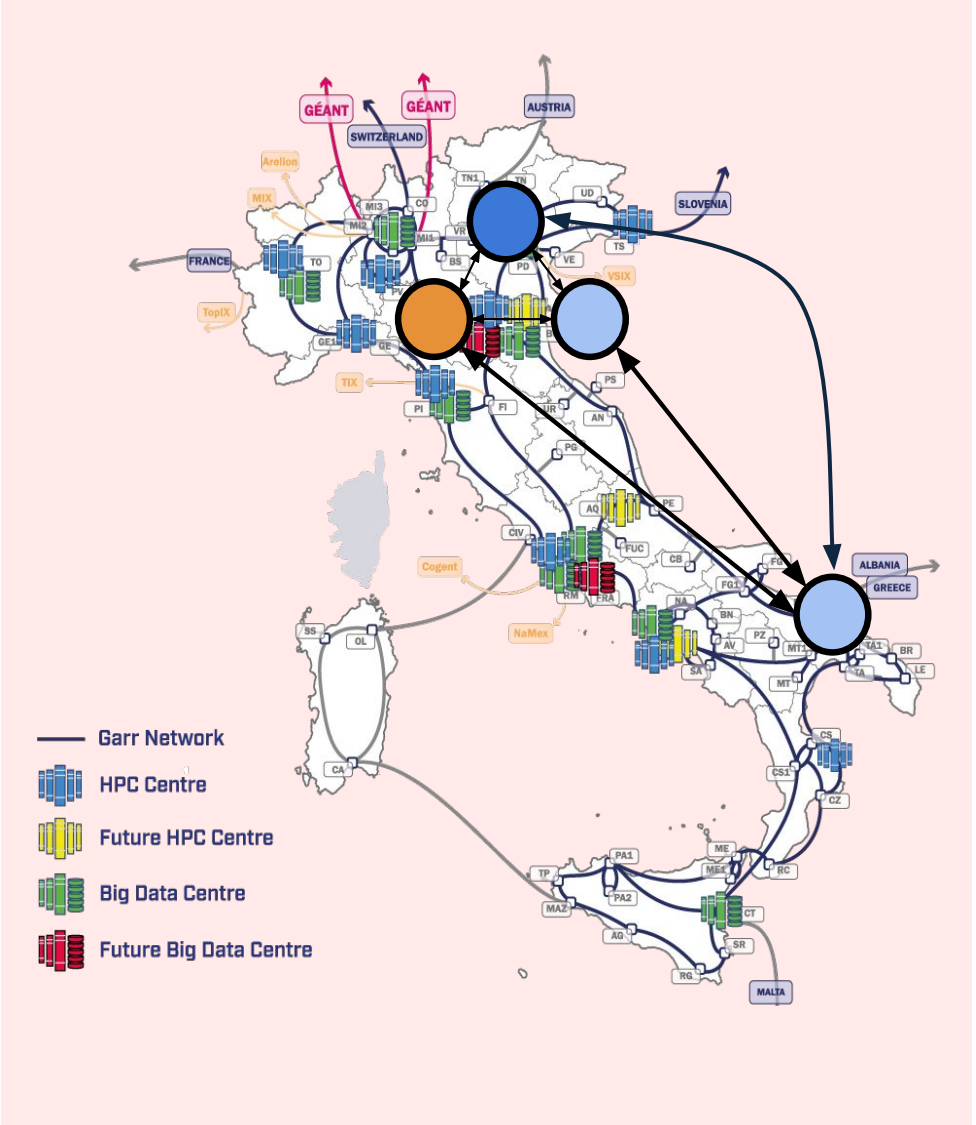

# THANKS FOR YOUR ATTENTION

**Gioacchino Vino - INFN, Unit of Bari**
**gioacchino.vino[at]infn.it**

CREST — funded by the Italian Ministry of University and Research under PN RIC 2021-2027, co-funded by the European Regional Development Fund. CUP I93C26000070007.

# BACKUP

# Where each component actually runs

*state as of September 2026*

| | CNAF | ReCaS-Bari |
|---|---|---|
| JupyterHub / Kubernetes platform | production | deployed |
| Home directories (NFS on VM) | production | production |
| Multi-tenant NFS (cloud-native) | designed | designed |
| Home backup — Borg | production | — |
| Home backup — restic + backrest | tested | production |
| POSIX mount of S3 user space (rclone) | production | production |
| Offloading via interLink | validated | commissioning |
| Kubernetes accounting (AUDITOR) | planned | in progress |
| Monitoring (Prometheus / Grafana) | production | production |

# The nine intermediate objectives

*what the Ministry monitors — and what each one actually delivers*

| | | WP | deadline | deliverable |
|---|---|---|---|---|
| **OI01** | Recruitment of the Infrastructure Manager | WP1 | **Nov 2026** | done |
| **OI02** | Plant upgrade — Bari | WP2 | **Nov 2027** | tender award |
| **OI03** | Plant upgrade — Naples | WP2 | **Nov 2027** | tender award |
| **OI04** | Federated infrastructure architecture | WP5 | **Nov 2027** | technical report |
| **OI05** | Energy-efficiency monitoring | WP7 | **May 2028** | report + pilot |
| **OI06** | CPU / GPU servers for INFN data centres | WP3 | **May 2028** | award decree |
| **OI07** | Project workshop | WP6 | **May 2028** | workshop held |
| **OI08** | Network systems (Bari, Naples) | WP4 | **Sep 2028** | award decree |
| **OI09** | Collaboration with companies and SMEs | WP8 | **May 2028** | agreements |

**For the purchase objectives the deliverable is the award decision — delivery, installation and acceptance follow, within the project (closes mid-2029).**