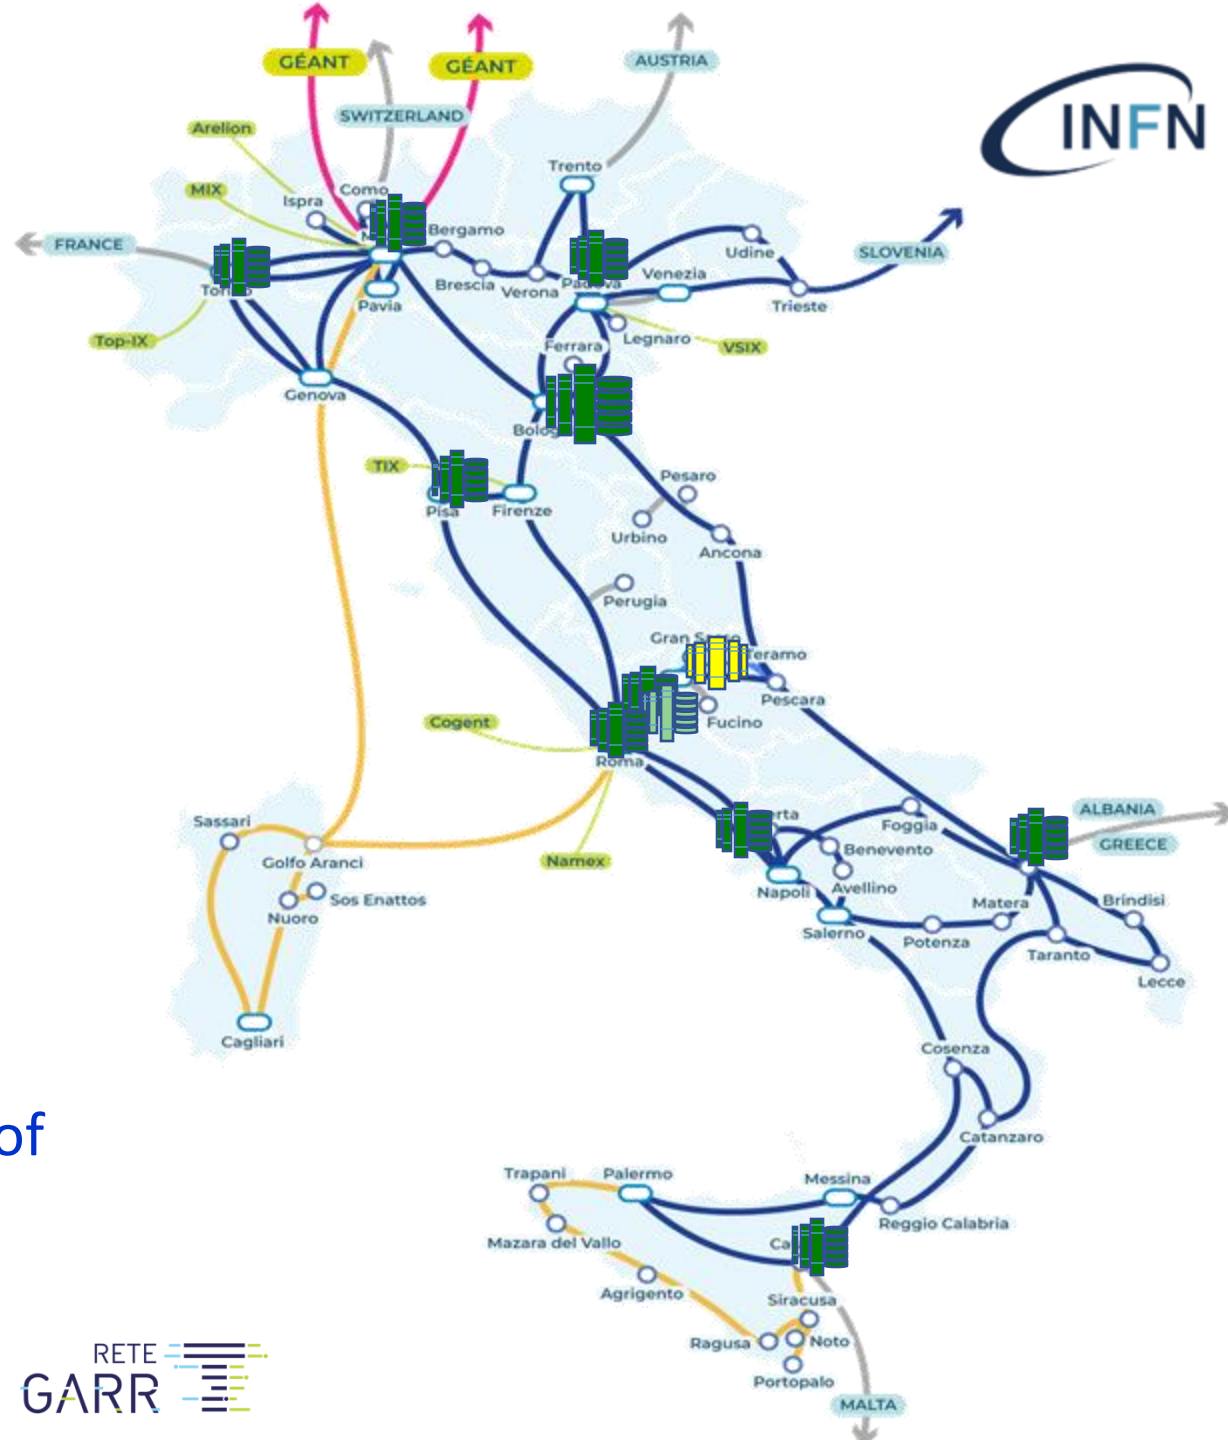# INFN DataCloud Infrastructure for Scientific Computing

- Tier-1 (CNAF)
- Tier-2's (BA, CT, LNF, LNL/PD, NA, MI, PI, RM1, TO)
- INFN Cloud
  - Backbone and federated clouds
- HPC4DR (LNGS)
- (Tier-3)

DataCloud was born to address the needs of INFN research activities, but it is serving several external projects

Building the Italian Cloud Federation

In the framework of the NRRP projects, in particular ICSC and TeRABIT, the INFN DataCloud model is the basis for the creation of the Italian Cloud Federation

Special focus on enhanced privacy platforms in DARE

The goal is to access all Italian scientific computing resources through uniform interfaces
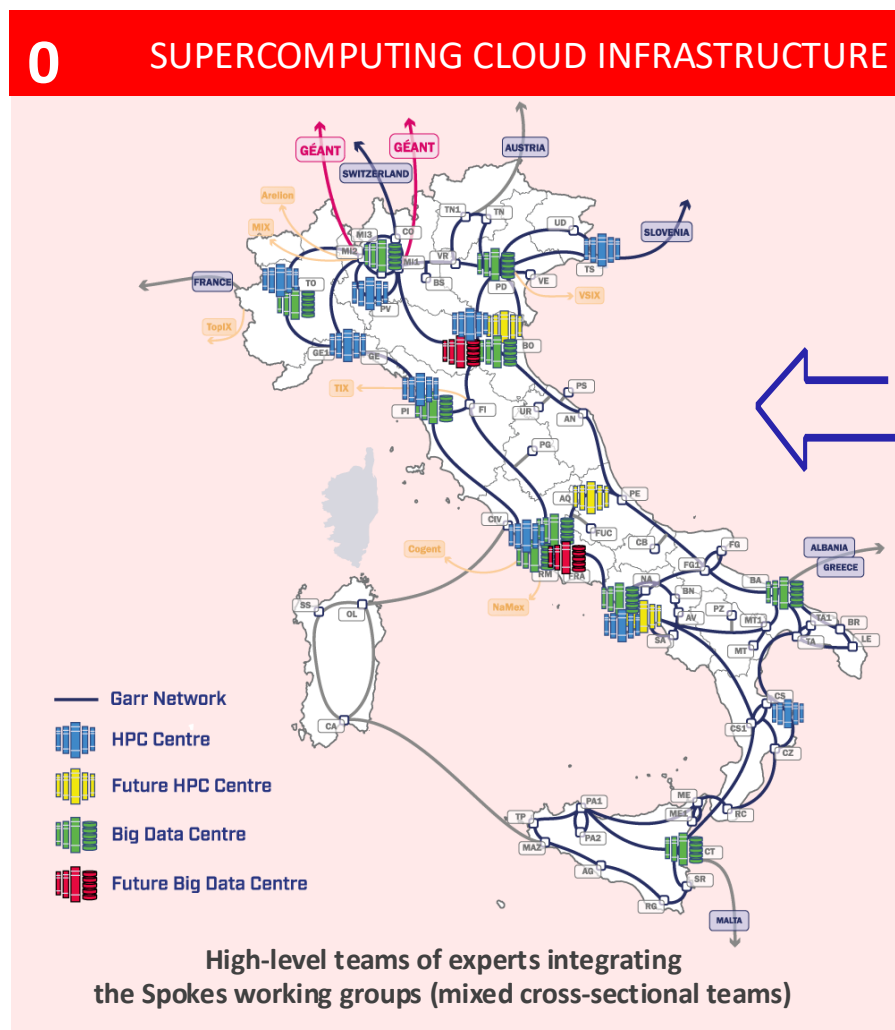
Main players: CINECA, GARR, INFN

National Research Centre in HPC, Big Data and Quantum Computing

https://www.supercomputing-icsc.it/en/icsc-home/

10 thematic spokes
1 infrastructure spoke
(CINECA, GARR, INFN)

25 universities
12 research institutes
14 strategic private companies

320 M€ budget

0 SUPERCOMPUTING CLOUD INFRASTRUCTURE

- Garr Network
- HPC Centre
- Future HPC Centre
- Big Data Centre
- Future Big Data Centre

High-level teams of experts integrating the Spokes working groups (mixed cross-sectional teams)

EDUCATION & TRAINING, ENTREPRENEURSHIP, KNOWLEDGE TRANSFER, 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

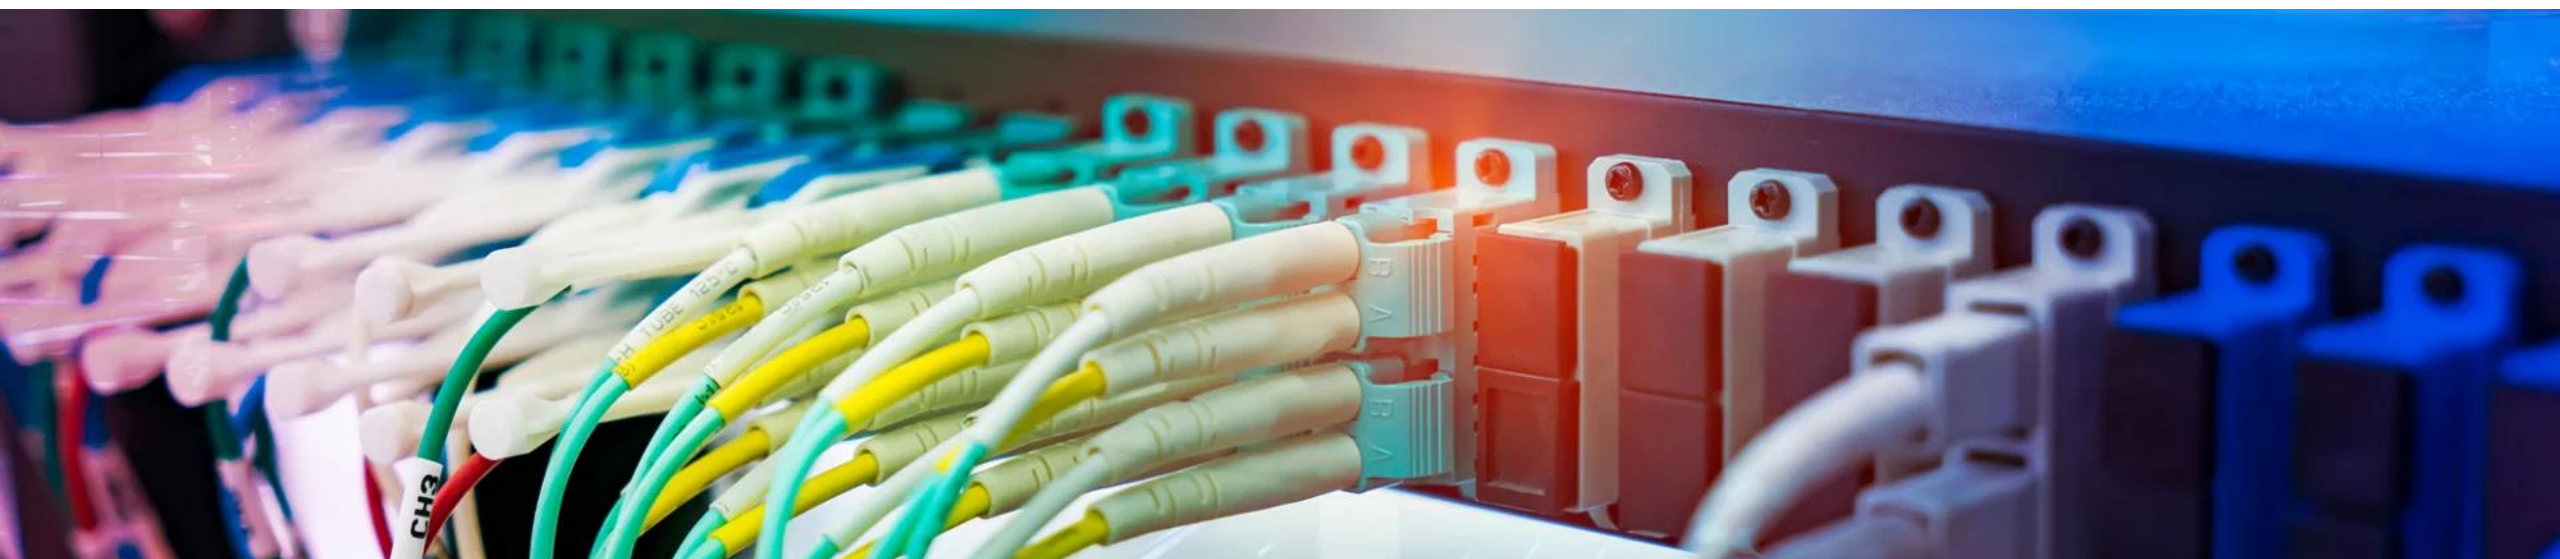# TeRABIT: Terabit Network for Research and Academic Big Data in ITaly

TeRABIT is a Research Infrastructure project synergic with ICSC

Partners are the same of the ICSC Spoke-0 (Supercomputing Cloud Infrastructure):

INFN, CINECA and GARR

Covers areas complementary to those of the ICSC infrastructure

41 M€ budget

DARE (DigitAl lifelong pRevEntion)

DARE is an initiative that aims to leverage data to enhance health promotion and prevention throughout the lifespan. It aims to use digital technologies to create a data-driven healthcare ecosystem

Funded through the "*Piano nazionale per investimenti complementari al PNRR*"

124 M€ budget (3.5 M€ for INFN)

INFN has the responsibility to deploy and manage part of the computing infrastructure

Leverages the EPIC (Enhanced PrIvacy and Compliance Cloud) INFN infrastructure

Acquired resources

HTC servers and storage by ICSC:

- ~600 kHS06, ~120 PBN
- Library + N*100 PB tapes

Coming:

- ~300 kHS06, ~50 PBN

HPC Bubbles by TeRABIT, ICSC, DARE

160 CPU nodes

61 GPU nodes

10 FPGA nodes

118 storage nodes

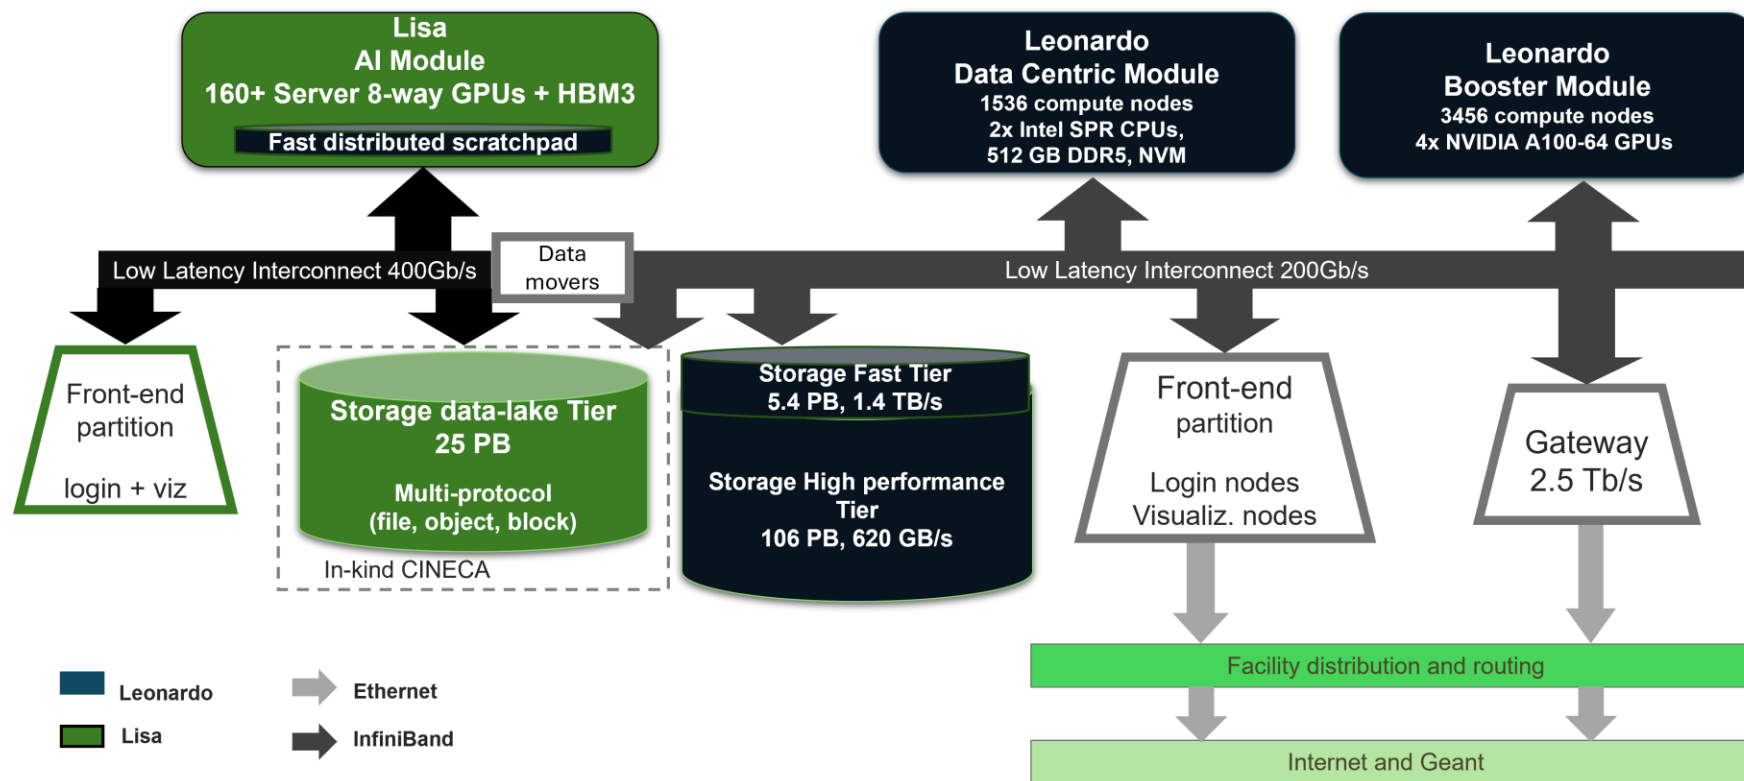| Node	Specifications
CPU node	192 physical cores 1.5TB RAM DDR5 IB NDR 400G 20TBL (SSD) + system disks
GPU node	As CPU + 4x NVIDIA H100 SXM5 with 80GB and HBM2e memory
FPGA node	32core RAM 768GB DDR5 IB NDR 440G 4 x XILINX U55C o 4 x TerasicP0701
Storage node (CEPH Bricks)	64 physical cores 1TB RAM DDR5 384 TBL HDD + 25.6 TBL NVMe
Accessories	Switch IB, Switch ETH cables IB, cables ETH Transceivers

Leonardo + Lisa High level architecture

Total ~100 Pflop/s (theoretical peak)

1 compute module
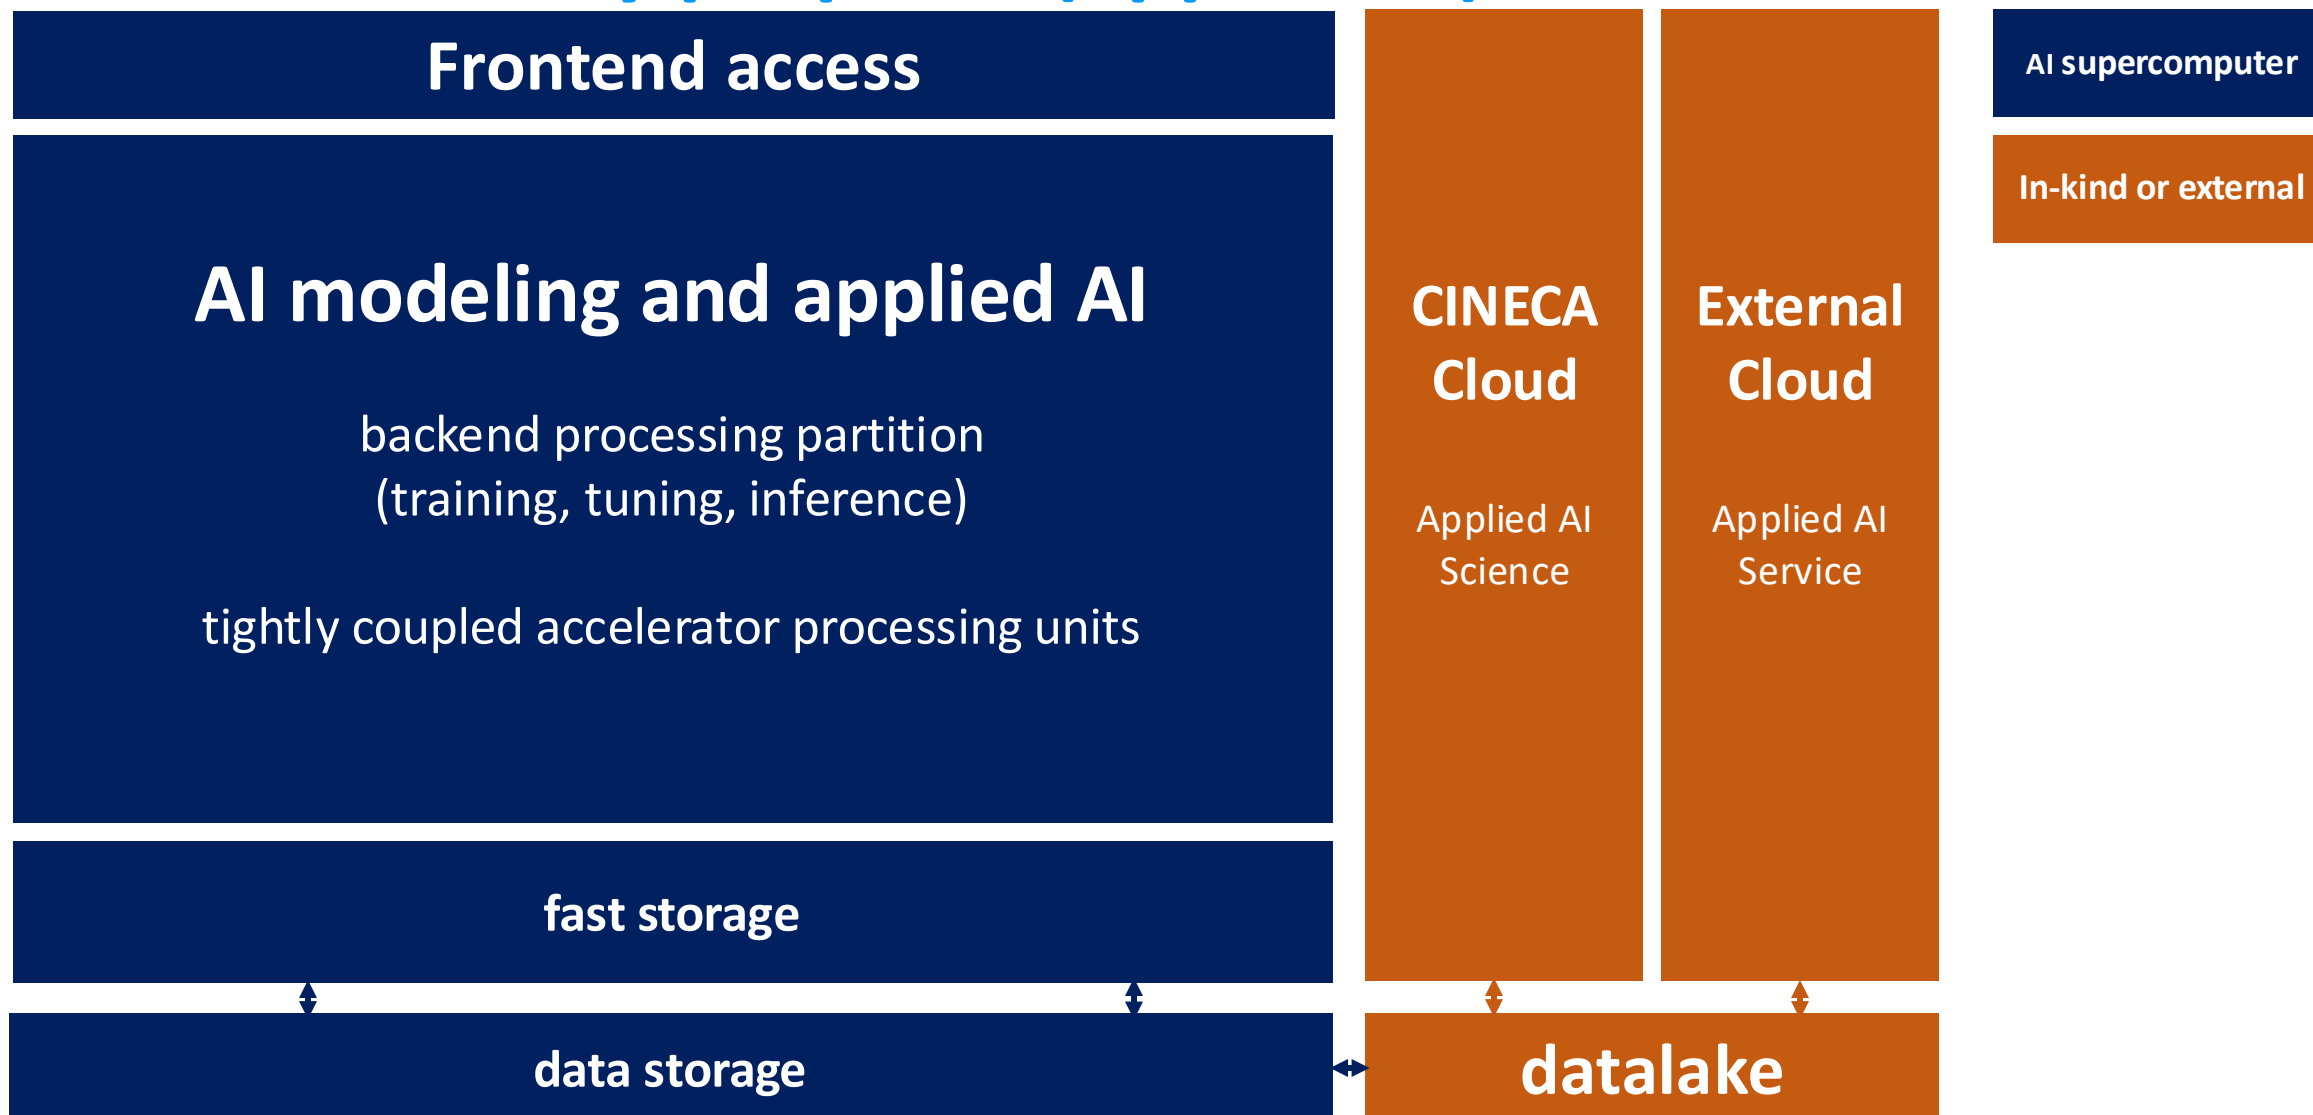
Best of market Interconnect

In-kind data-lake storage

AI oriented benchmarks

Lisa
AI Module
160+ Server 8-way GPUs + HBM3
Fast distributed scratchpad

Leonardo
Data Centric Module
1536 compute nodes
2x Intel SPR CPUs,
512 GB DDR5, NVM

Leonardo
Booster Module
3456 compute nodes
4x NVIDIA A100-64 GPUs

Low Latency Interconnect 400Gb/s

Data movers

Low Latency Interconnect 200Gb/s

Front-end partition

login + viz

Storage data-lake Tier
25 PB

Multi-protocol (file, object, block)

In-kind CINECA

Storage Fast Tier
5.4 PB, 1.4 TB/s

Storage High performance Tier
106 PB, 620 GB/s

Front-end partition

Login nodes Visualiz. nodes

Gateway 2.5 Tb/s

Facility distribution and routing

Internet and Geant

| | Leonardo | | Ethernet |
| | Lisa | | InfiniBand |

CINECA AI Factory proposal (approved)

Frontend access

AI modeling and applied AI

backend processing partition
(training, tuning, inference)

tightly coupled accelerator processing units

fast storage

data storage

CINECA Cloud

Applied AI Science

External Cloud

Applied AI Service

datalake

AI supercomputer

In-kind or external

CINECA AI Factory proposal (approved)

Front end access

AI modeling and AI applied partition

batch processing (training, tuning, inference)

tightly coupled accelerator processing units

storage scratch/pad

datalake

750-1000 PFLOPS

50-60 EFLOPS FP8

2000-3000 servers

120-160 racks

13-17 MW peak

8-10 MW ops

DLC/AIR: 70/30 – 90/10

32-26 °C

Network

GARR-T evolution:

Reach new areas: Sardinia and Abruzzo

Upgrade and integrate network in the South of Italy

GARR-T reaches the goal to become a fully unified and pervasive network for R&E community in the whole country

End User Benefits

9 New POPs

Doubling of POPs in 6 cities

Distributed 100G access, with 400G available
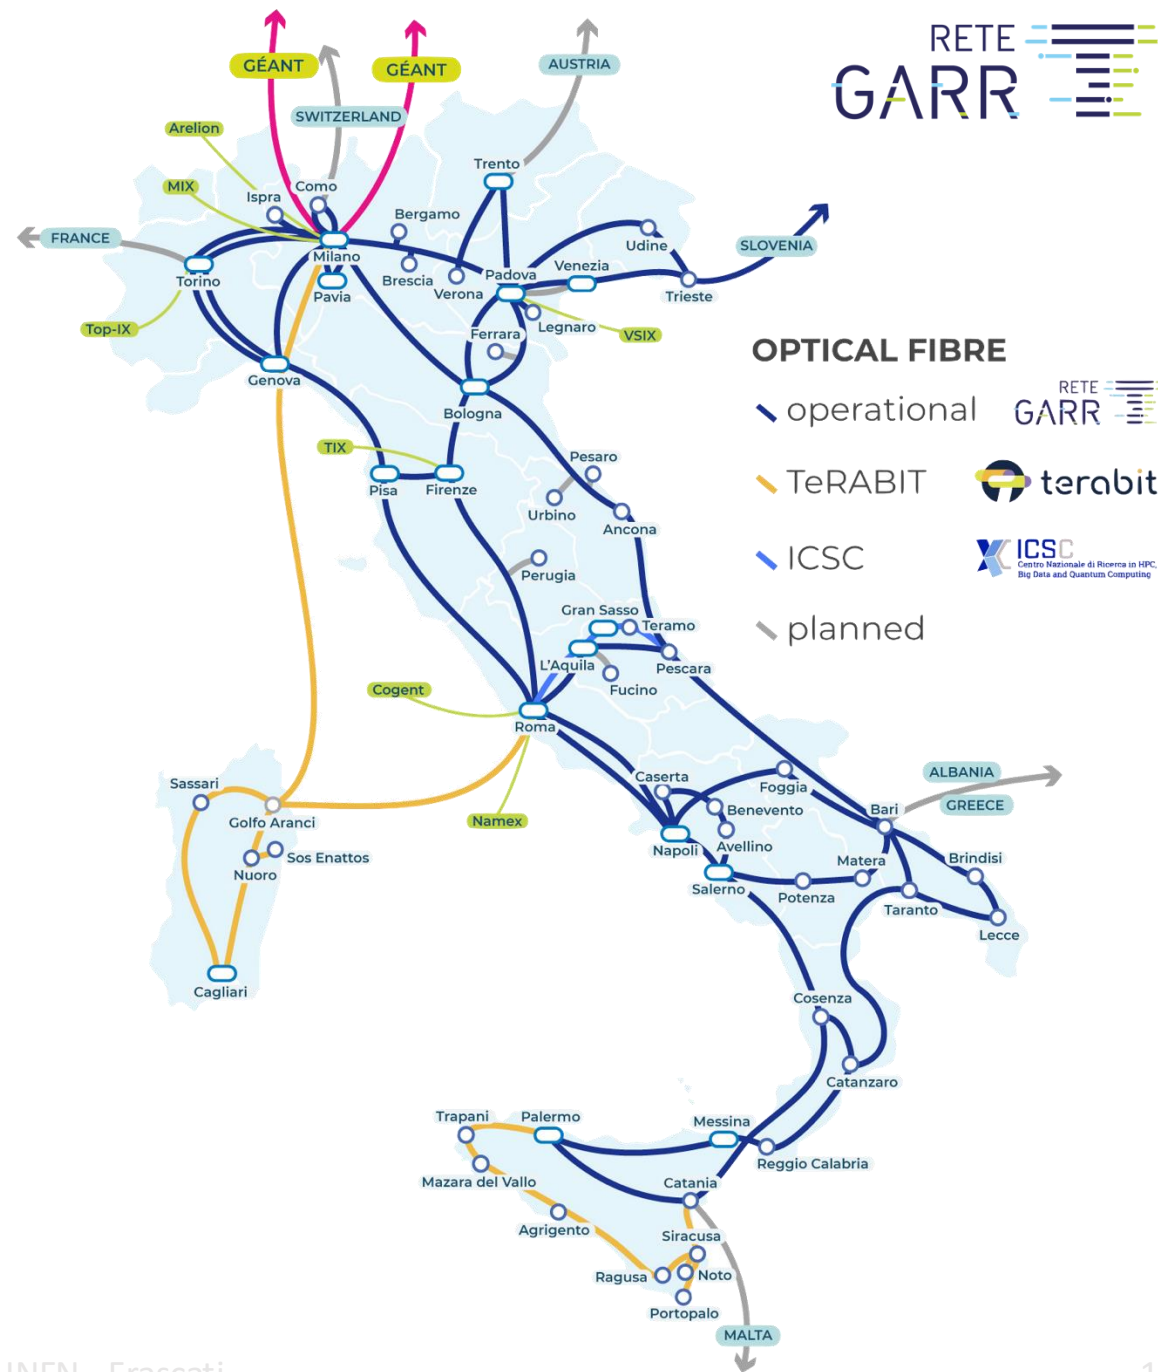
Heterogeneity

Enables to efficiently tackle diverse applications

- complex simulations

- high-throughput data processing

- Interactive analysis

- data storage and data access

- IOT integration

- critical data management

- …

There is no «one size fits all»!
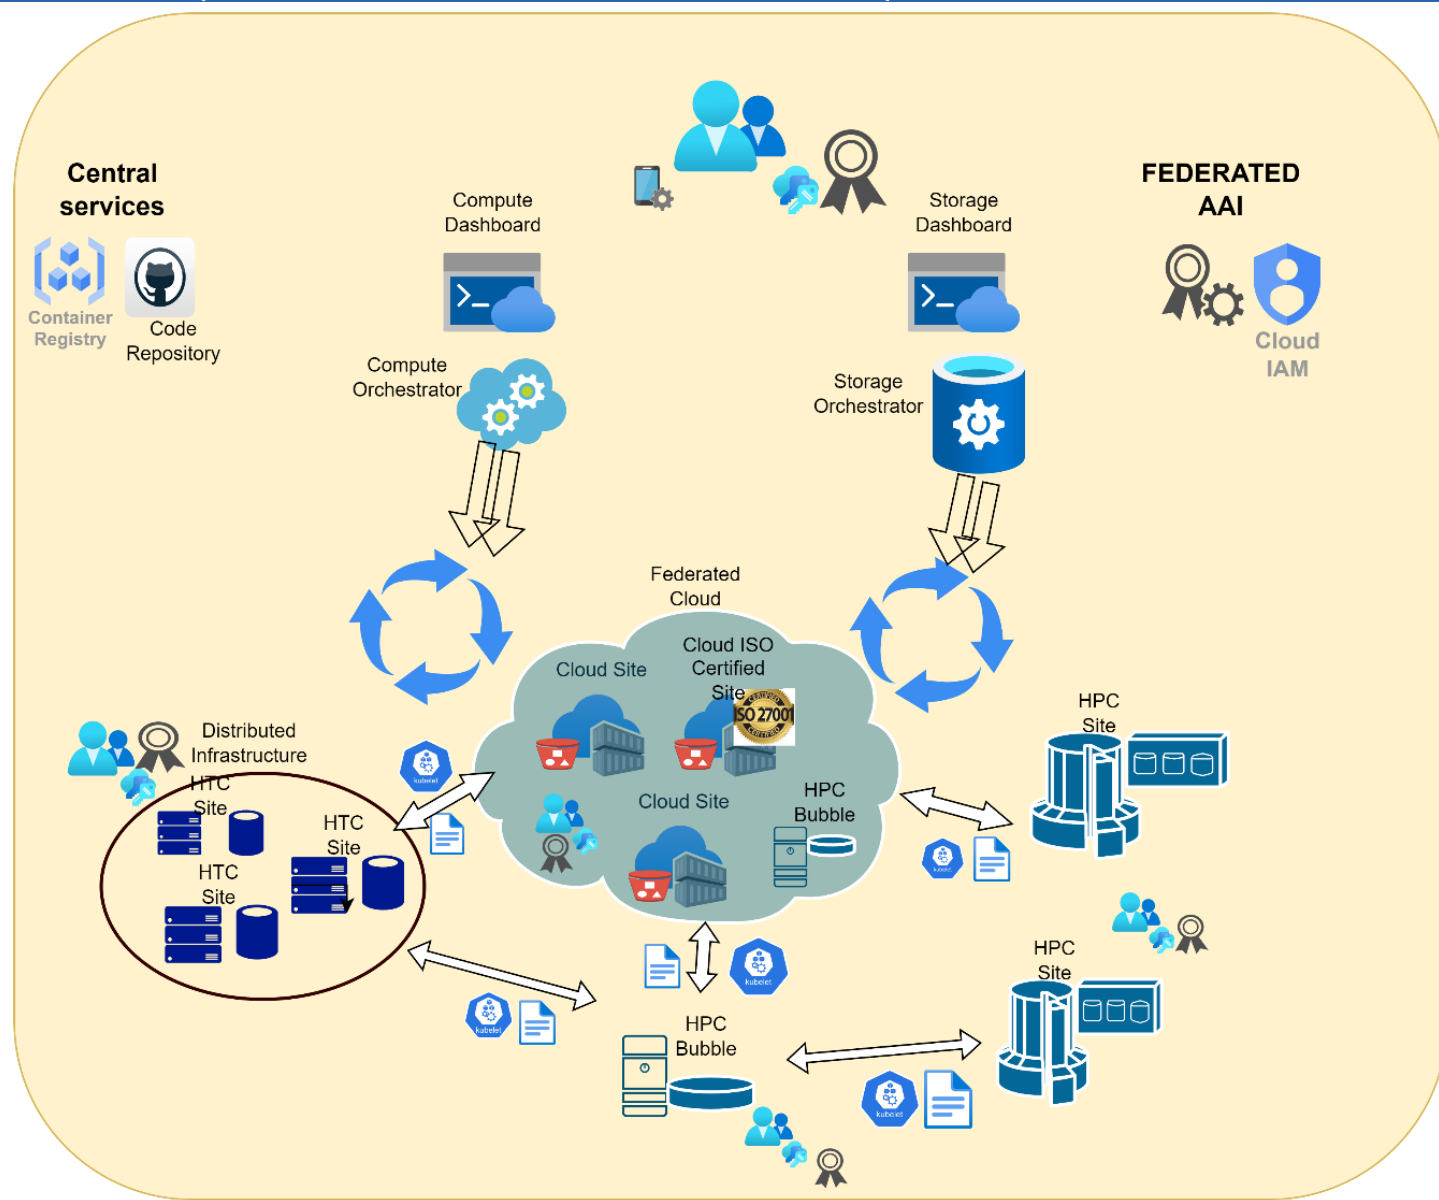
A data lake for research

Existing infrastructures aggregation, upgraded and made available to scientific domains

A dynamic model, where infrastructures and domains can also be temporary

A clear separation between the physical and the logical levels

A high-speed network interconnection to hide the actual resource locations

A unified vision (when needed) of an Italian research data-lake

Data Lake Entry Point

User Community Services

Individual communities services

General Services

Generic services

Infrastructure sites

HPC

Cloud Provider

Tier1

Tier2

Data centres and resource providers
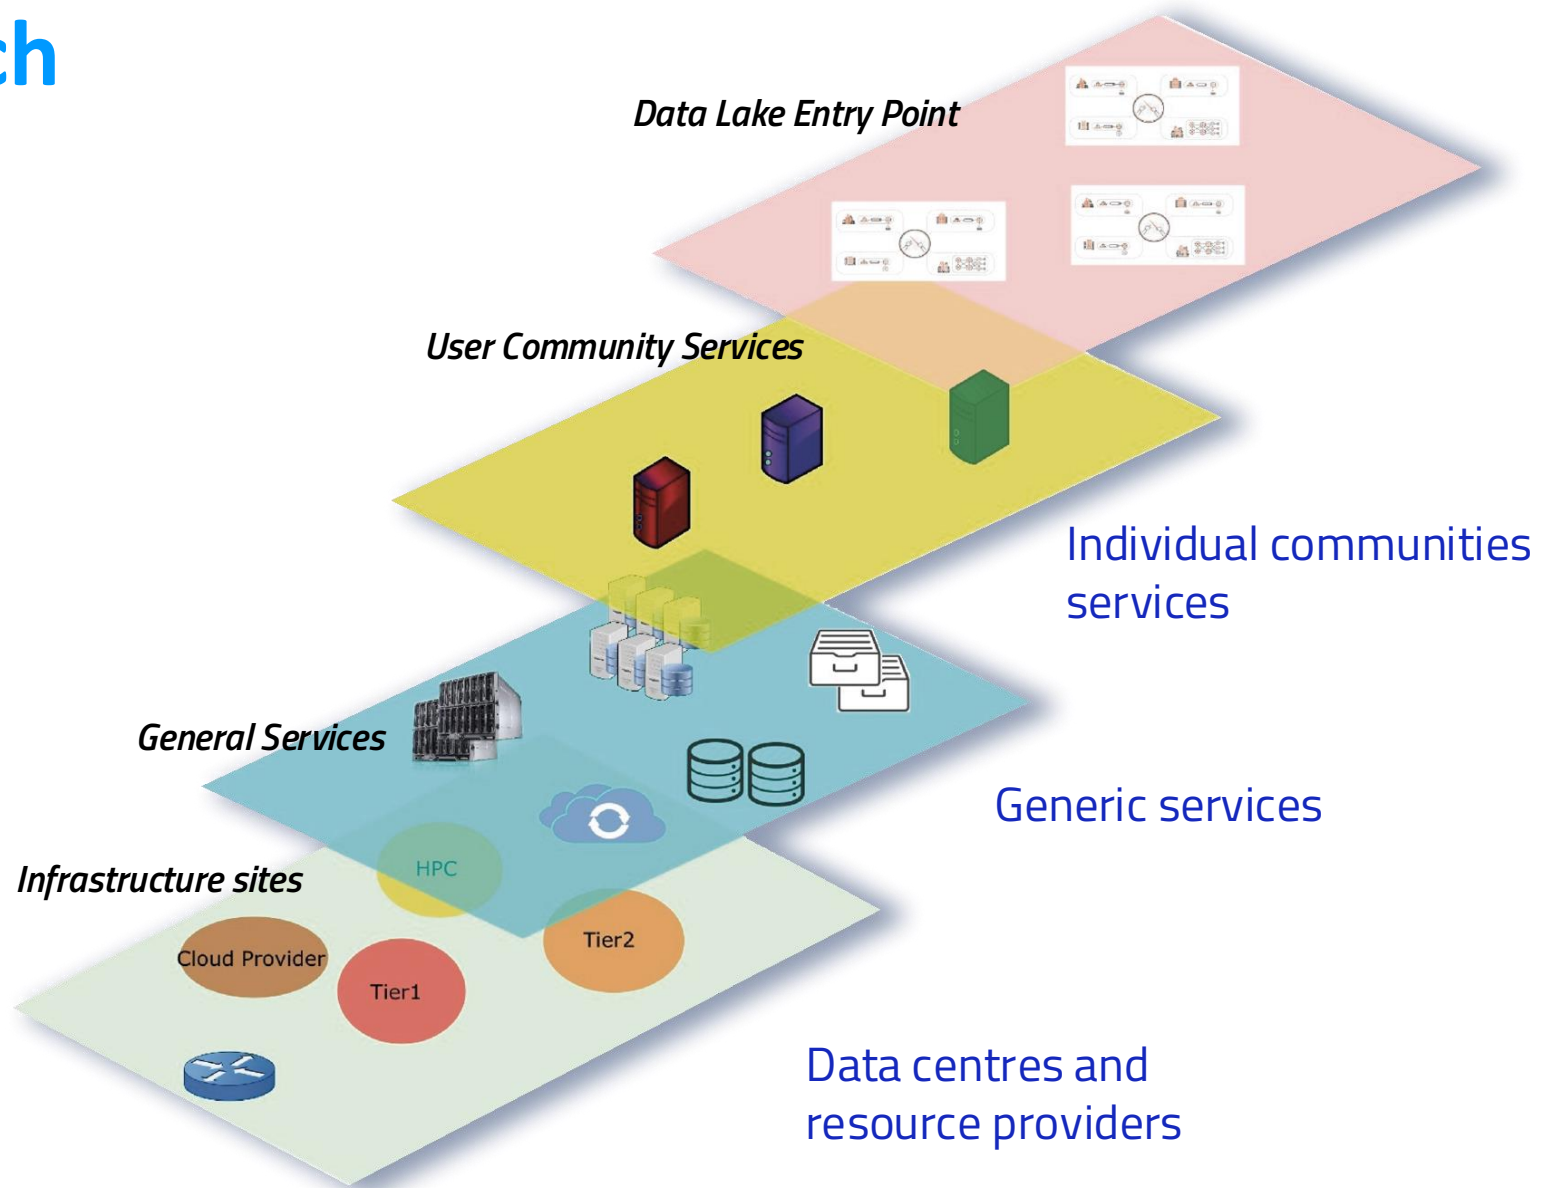
Data-centric model

Decouple storage and CPU

Storage nodes interconnected with high bandwidth network

Heterogeneous computing nodes can access data wherever they are
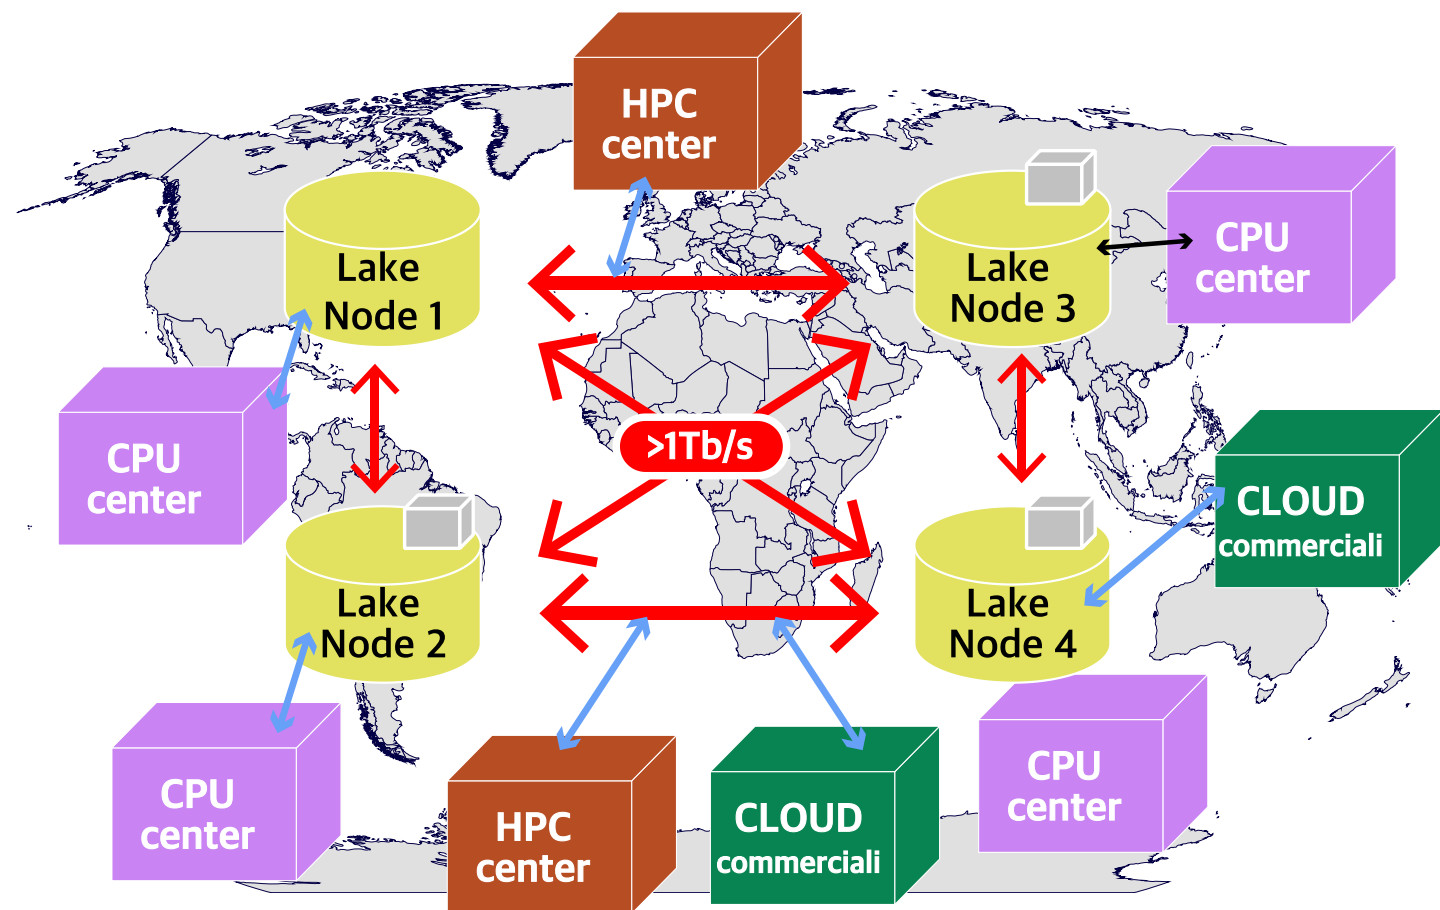
Inclusivity

The federation will include data centres that are already in production, and part of international communities

The procedures for joining the federation must be non-intrusive

Standard must be used whenever possible, and developed when missing

The federation will serve users of several fields and organizations

The procedures for user's onboarding must be as simple as possible

E.g.: use of Identity Federations

INDIGO IAM
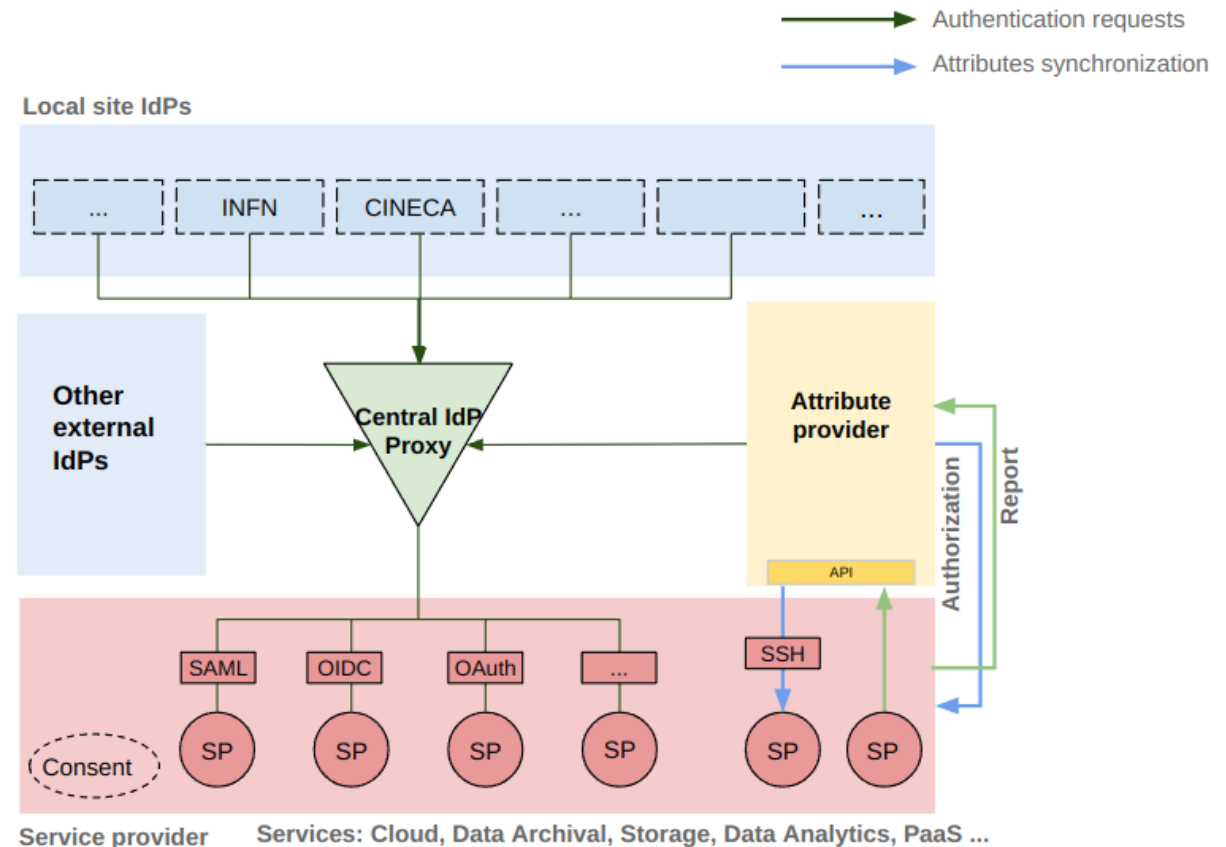
Ease of use

The federation will serve users with different computing competences

Complexity of the underlying infrastructure hidden to the end user

Support field experts in developing platforms that enable the effective exploitation of the infrastructure through composition of services and resources
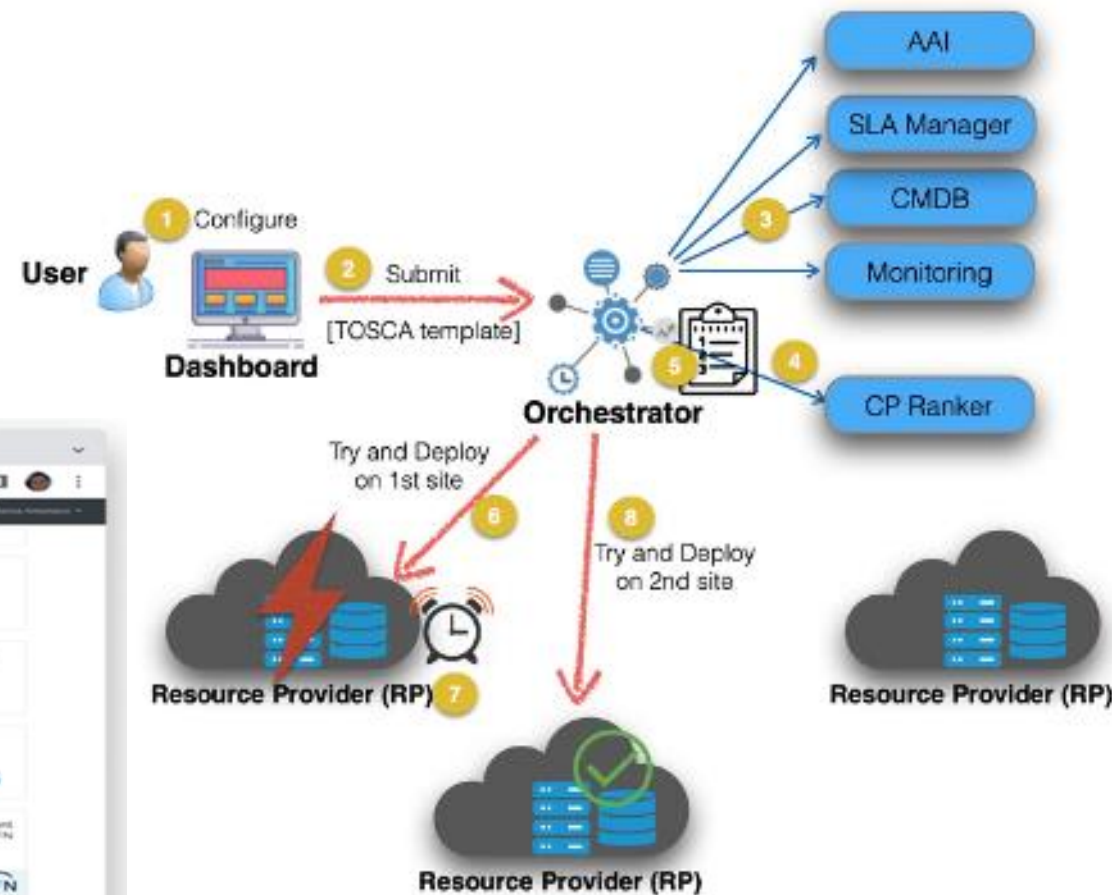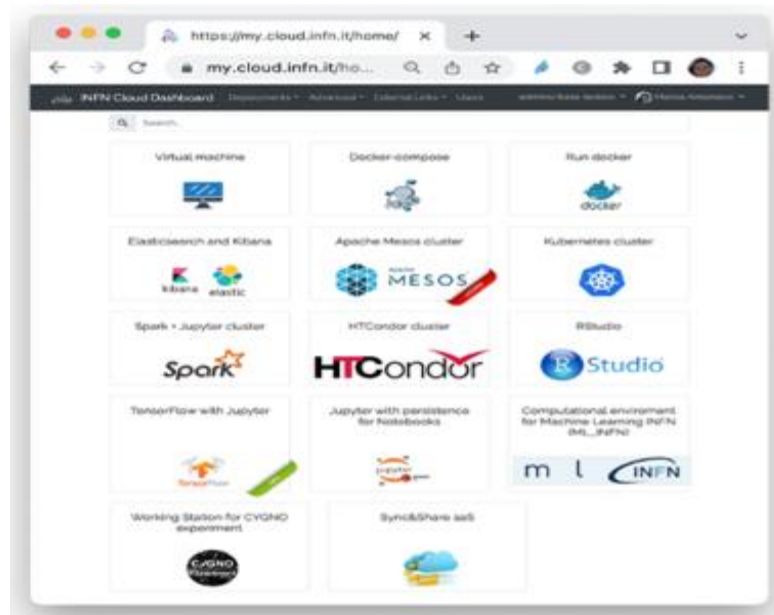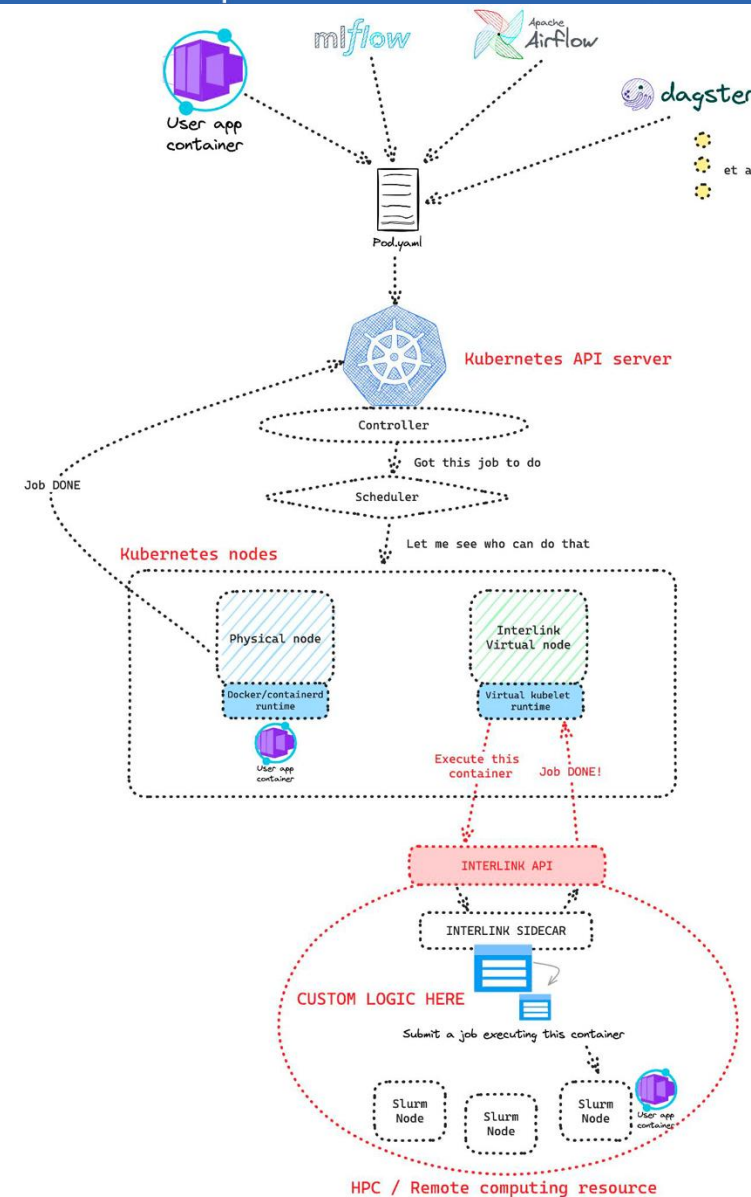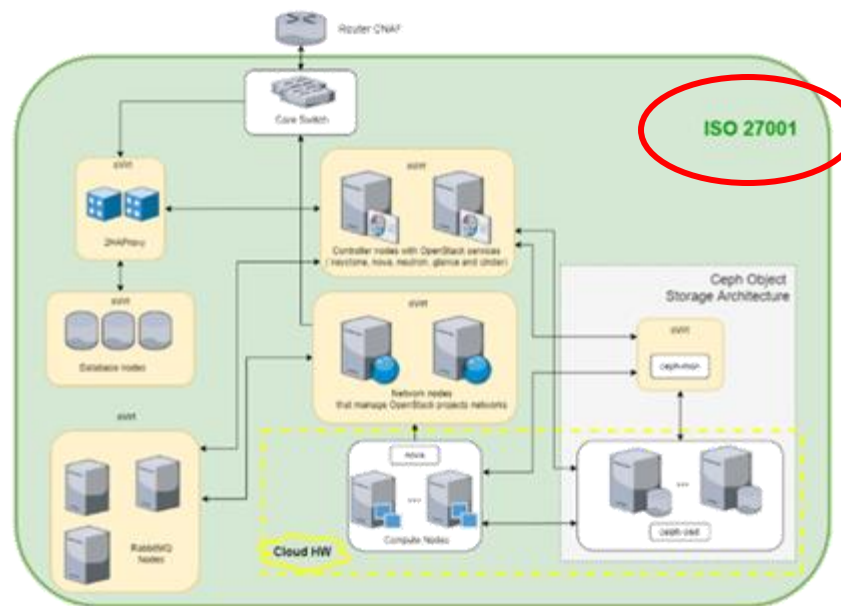
INDIGO PaaS orchestrator

Flexibility

Support multiple access methods to the resources, oriented to:

a. Transparency and ease of use

b. Efficiency and effectiveness

Support application-specific requirements

E.g. enhanced privacy
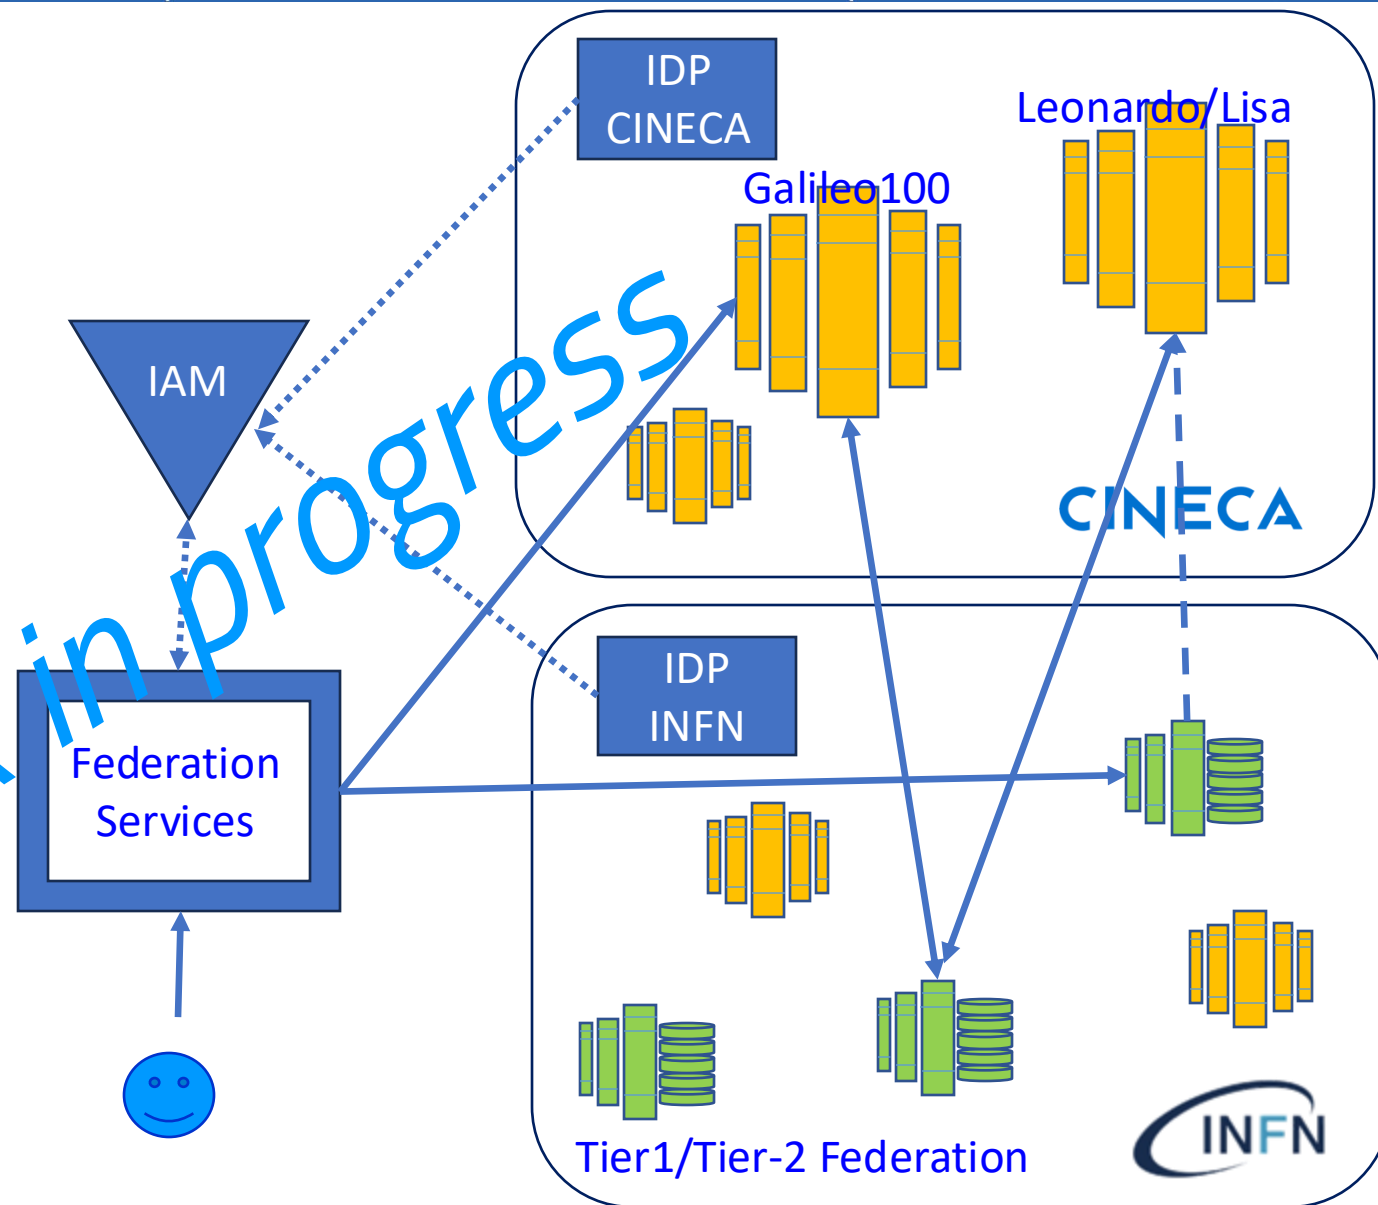
Conclusions

The PNRR projects brought new resources (inside and outside INFN) usable also by INFN research projects but brought also new user communities

Increased heterogeneity

  of resources → flexible allocation model

  of users → user support reorganization

DataCloud is evolving to be able to address this new scenario

  Cloud model to be able to hide complexity to the vast majority of end users

  Support for different allocation models to maintain efficiency and effectiveness