



Finanziato dall'Unione europea
NextGenerationEU



Ministero dell'Università e della Ricerca



Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA



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

The evolution of INFN's Cloud Platform: improvements in Orchestration and User Experience

Main author:	Marica Antonacci
Co authors:	Barbara Martelli Daniele Spiga Enrico Vianello Ettore Serra Giacinto Donvito Giovanni Savarese Luca Giommi
Presenter:	Claudio Grandi

October 19 - 25, 2024

CHEP 2024

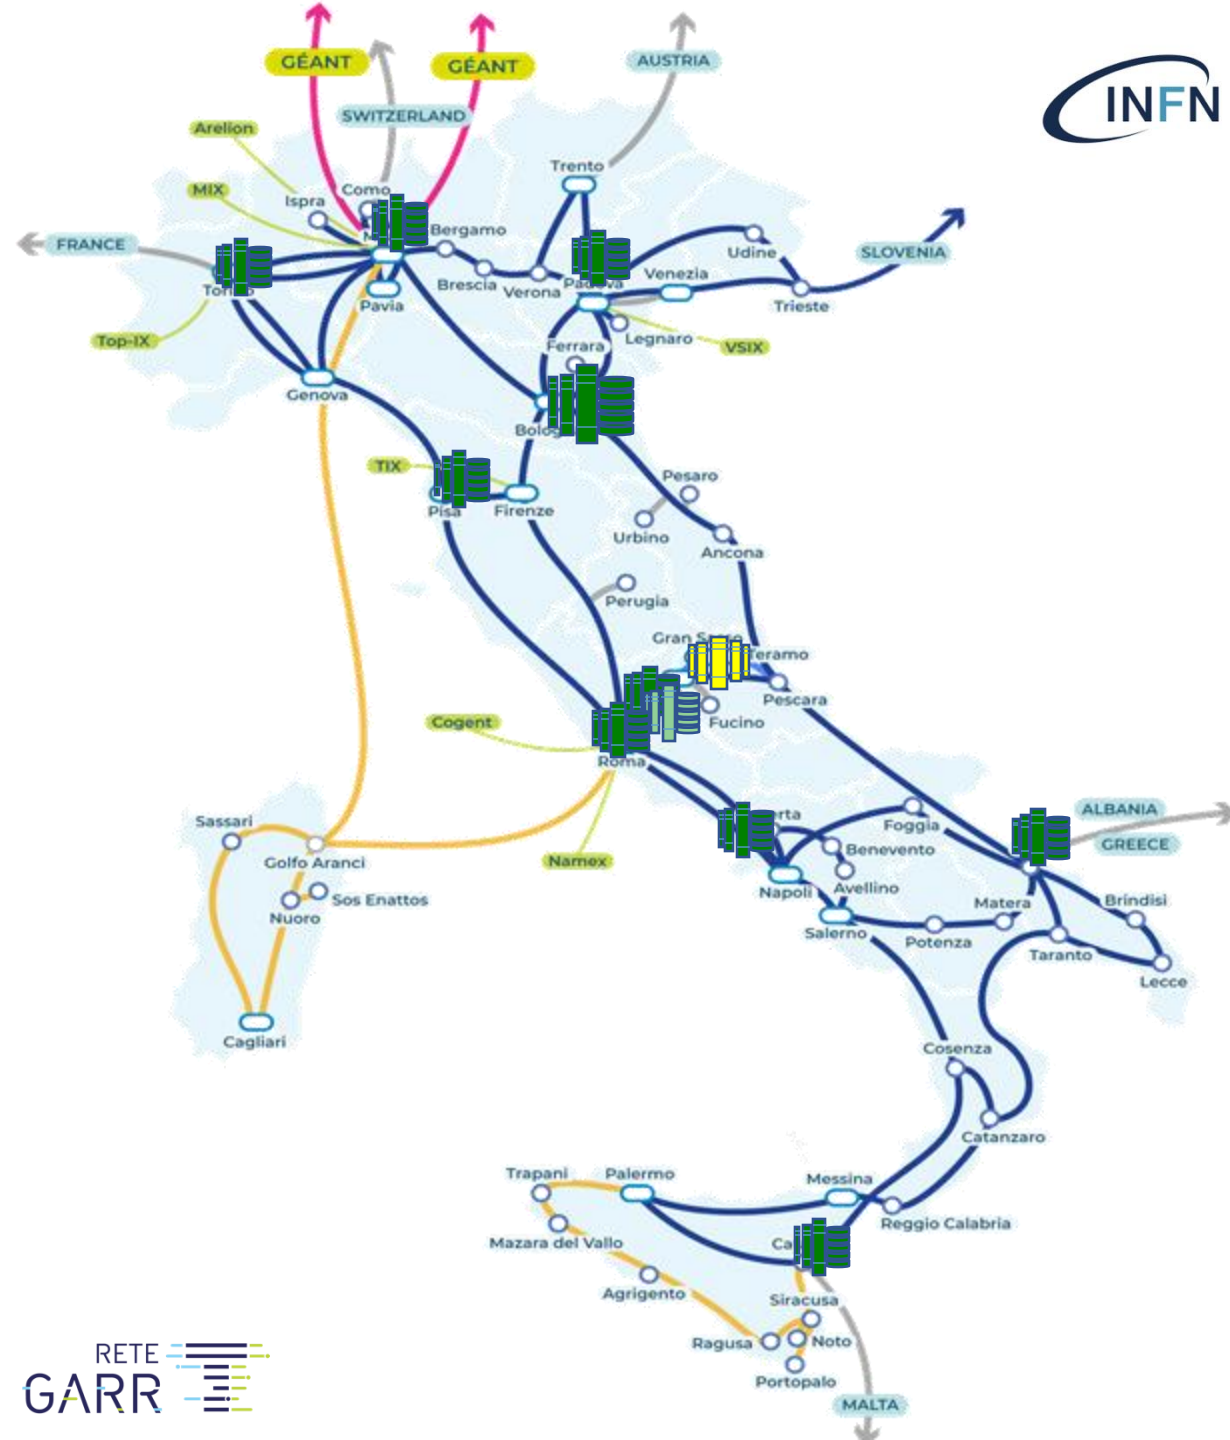




DataCloud is the Infrastructure for INFN Scientific Computing

- 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 addresses the needs of INFN research projects





DataCloud is evolving into a Cloud Federation

Following the INFN Cloud model, resources are being made available through Cloud interfaces

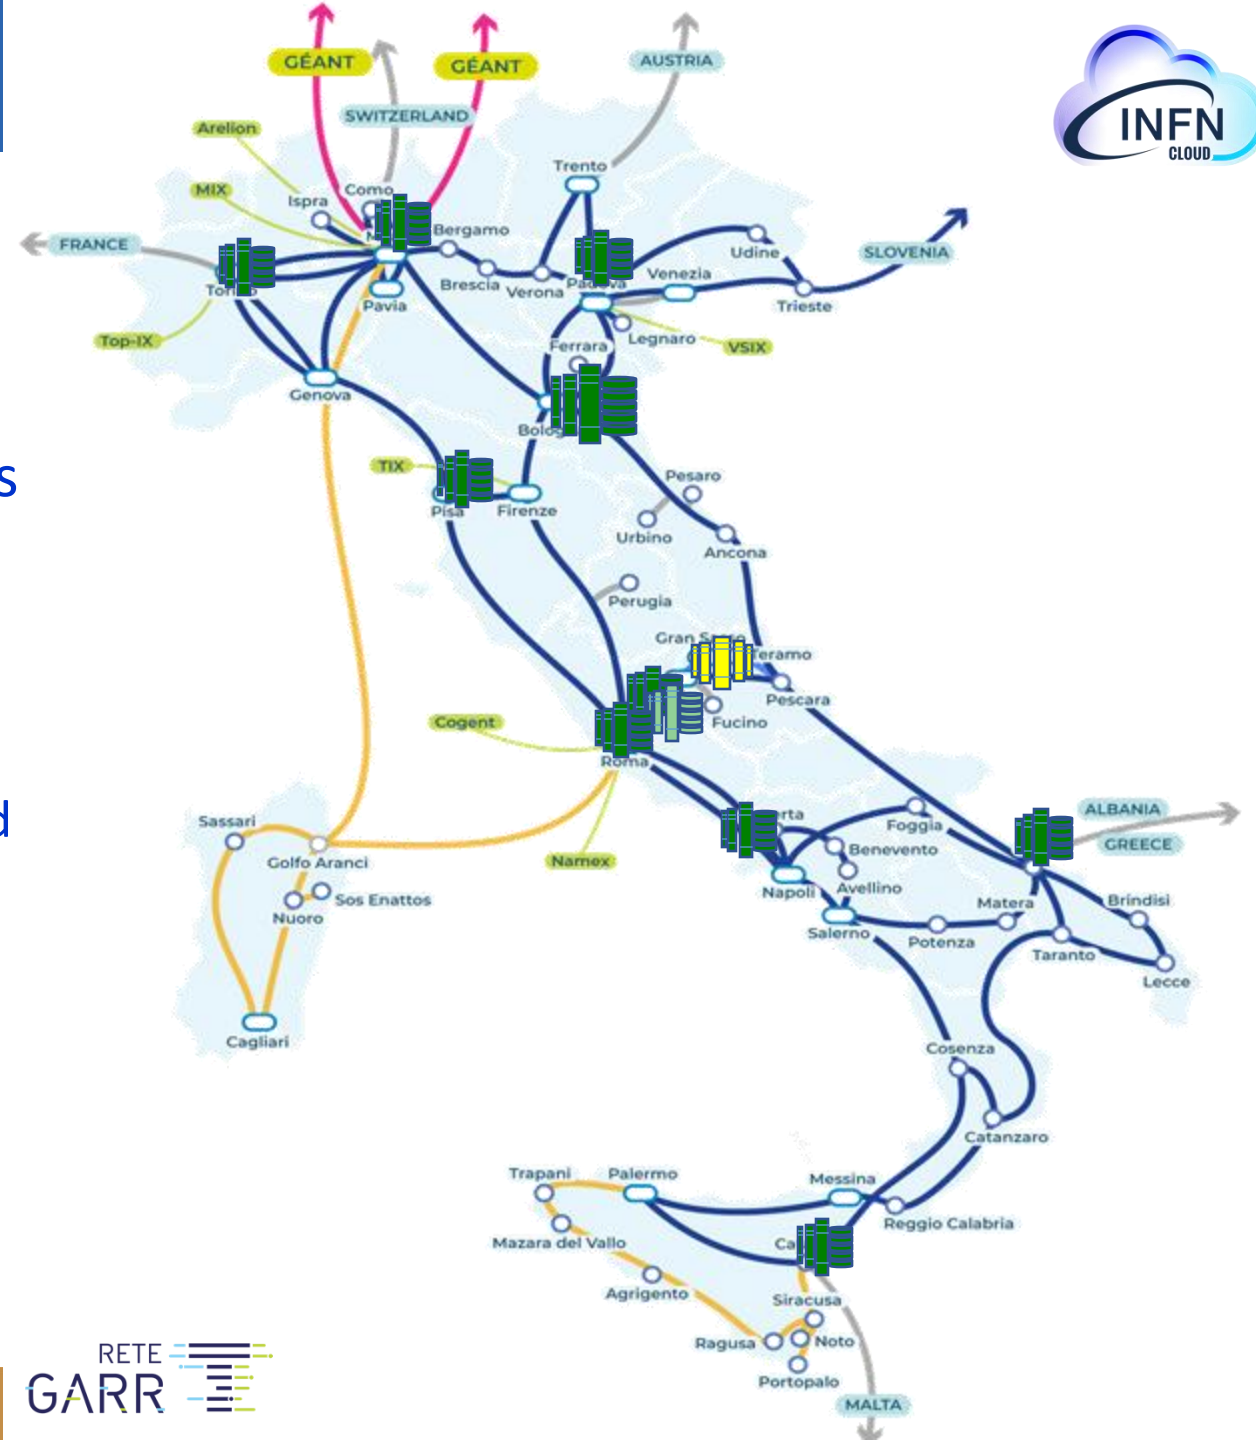
Inclusivity, through a lightweight federation model and the adoption of standards

Ease of use, through the PaaS orchestrator and dashboard

Flexibility, thanks to hybrid resource allocation mechanisms

Traditional (Grid and batch system) access remains as needed and when convenient

E.g. through Virtual Kubelets, ...





DataCloud is the basis for the Italian Cloud Federation

In the framework of the current Next Generation EU projects, in particular ICSC and TeRABIT, INFN has a leading role in the creation of the Italian Cloud Federation

The goal is to access all Italian scientific computing resources through uniform interfaces

Main players: INFN, CINECA, GARR





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

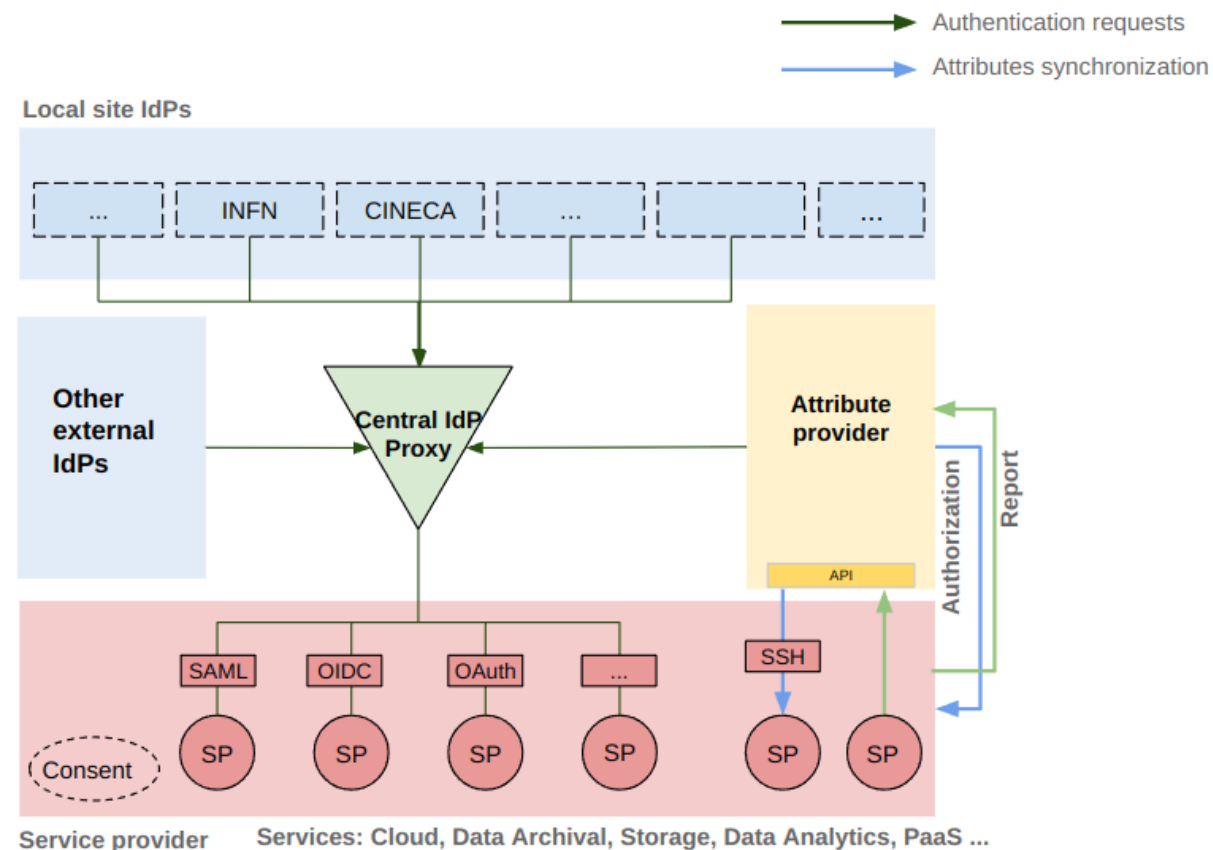
The federation will serve users of several fields and organizations

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

E.g.: use of Identity Federations

INDIGO IAM

See: *Enrico Vianello: Evolving INDIGO IAM towards the next challenges Track 4, Tuesday 17:27*

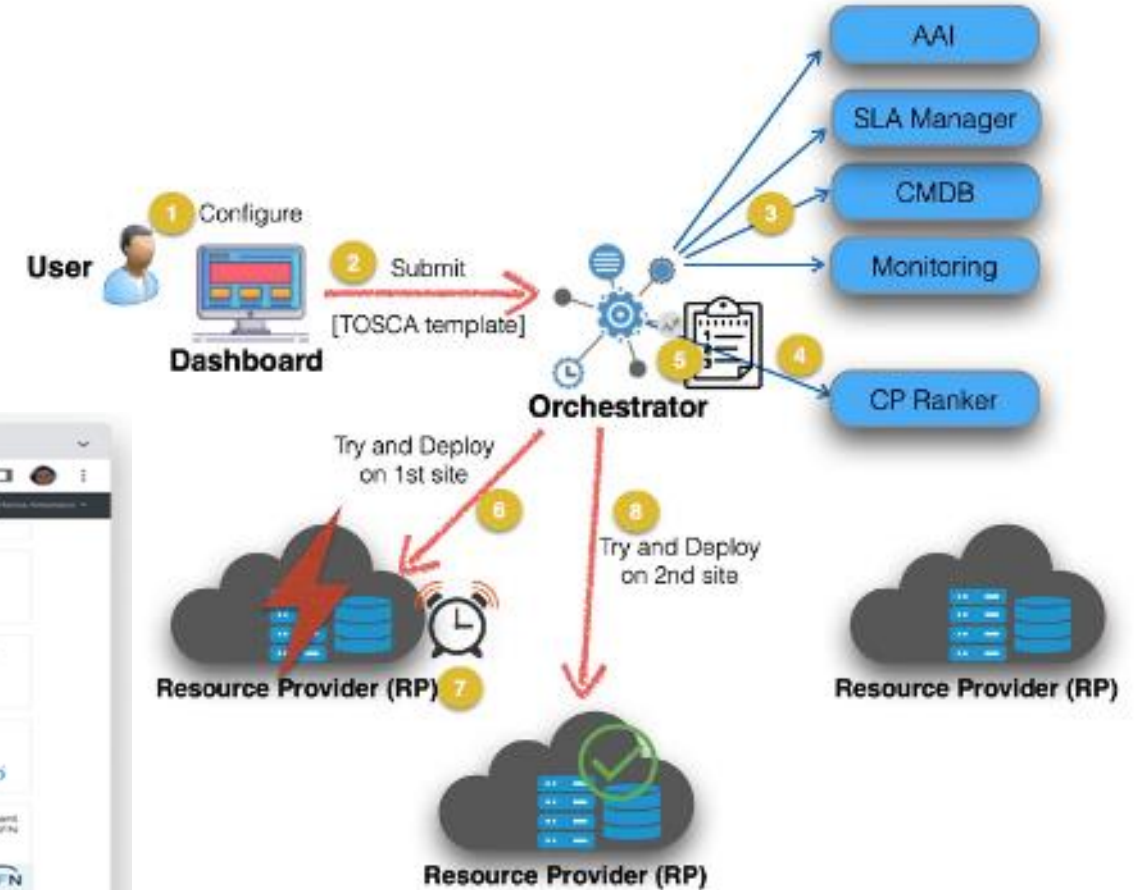
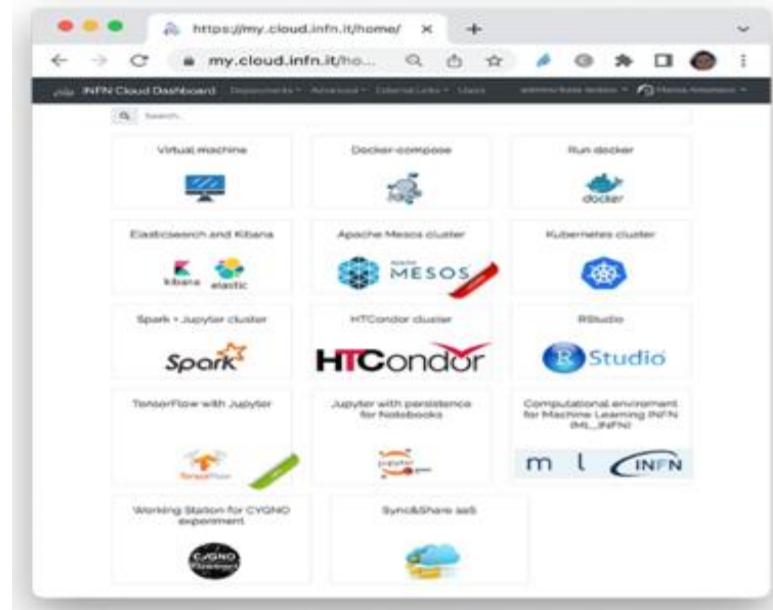


Ease of use

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

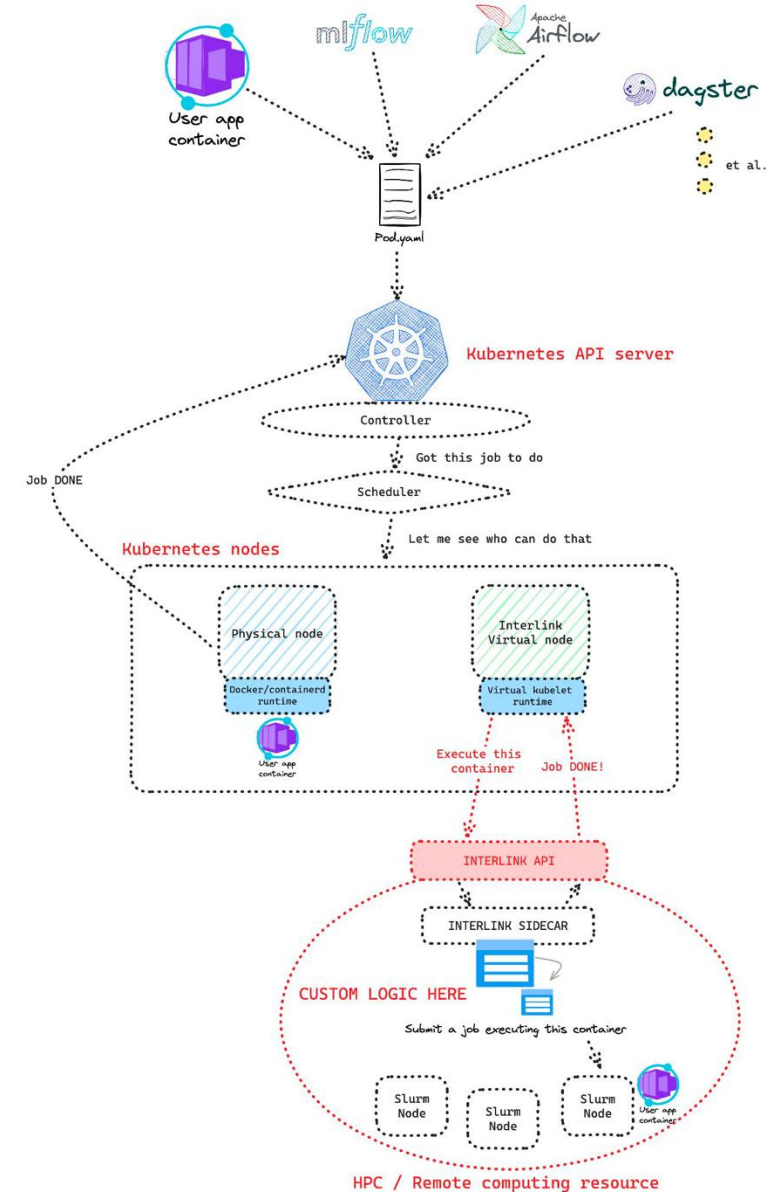
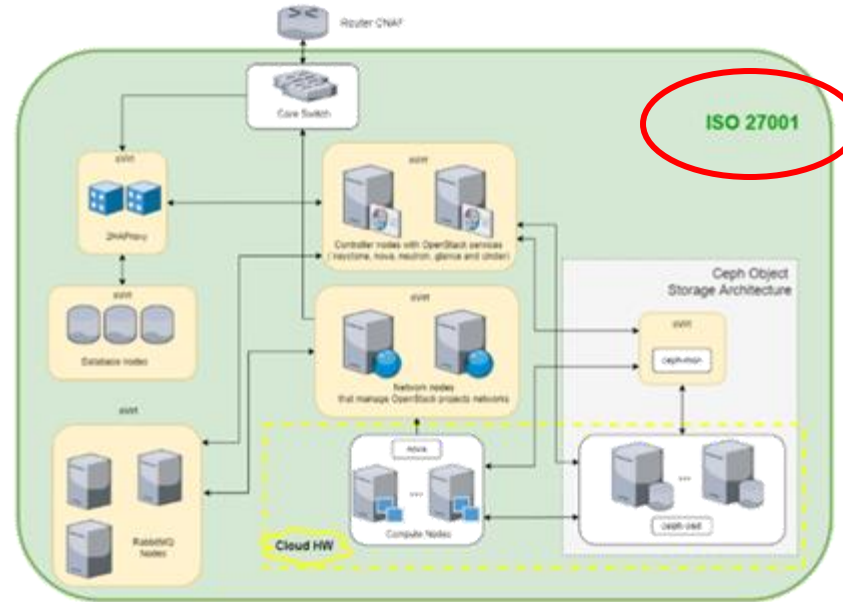
Flexibility

Support multiple access methods to the resources, oriented to:

- a. Transparency and ease of use
- b. Efficiency and effectiveness

See: Diego Ciangottini Unlocking the compute continuum: scaling out from cloud to HPC and HTC resources Track 7, Tuesday 16:51

Support application-specific requirements
E.g. enhanced privacy





First Proof of Concept of the Italian Federation


INDIGO IAM to federate CINECA and INFN IdPs


INDIGO PaaS Orchestrator to transparently access CINECA and INFN OpenStack-based resources

InterLink offloading to reach CINECA's Leonardo Supercomputing

RUCIO to federate CINECA and INFN storage systems

Working progress





Evolution of the PaaS Orchestrator

Security

Replace obsolete and vulnerable software components to address emerging threats

Robustness and scalability

Exploit modern and robust components that can cope with heterogenous resources and scale with increasing number of users and services

User experience

Hide the complexity of the underlying system through a web dashboard



Refactoring

Main programming language:



 **Flowable** workflow engine → micro services +



message broker

New Federation Registry micro service

All components use OpenID-Connect/Oauth2

New connectors for IM/  **kubernetes**_{nd}




... and more

Leveraging AI

Master orchestration exploits AI-based techniques


More dynamic and efficient choices



Preparatory work:

Identified significant metrics and data sources

Prepared datasets to analyze the problem



See: [Enrico Vianello AI-based approach for provider selection in the INDIGO PaaS Orchestration System of INFN Cloud Poster session, Wednesday \(#321\)](#)



Dashboard redesign

Users' entry point

Graphical representation of TOSCA templates

Service Catalog

Fully-managed centralised services

Self-managed on-demand services

INFN Cloud object storage
the centrally managed service based on Ceph Rados-Gateway

GO TO SERVICE →

INFN Cloud Registry
The centrally managed service INFN Cloud Registry, based on Harbor

GO TO SERVICE →

INFN-Cloud monitoring
the INFN-Cloud monitoring service

GO TO SERVICE →

Notebooks as a Service (NaaS)
Jupyter Notebooks as a Service
Scope: admins/beta-testers

GO TO SERVICE →

Virtual machine
Launch a compute node getting the IP and SSH credentials to access via ssh

CONFIGURE →

Spark + Jupyter cluster
Deploy a complete Spark + Jupyter Notebook on top of a Kubernetes (K8s) computing cluster

CONFIGURE →

Docker compose
Deploy a virtual machine with docker engine and docker compose pre-installed. Optionally run a docker compose file fetched from the specified URL.

CONFIGURE →

HTCondor mini
Deploy HTCondor mini, a technology preview of an all-in-one ("minicondor") HTCondor. This type of install is useful for testing and experimentation.

CONFIGURE →

Run docker
Run a docker container

CONFIGURE →

HTCondor cluster
Deploy a complete HTCondor cluster

CONFIGURE →

INDIGO IAM as a Service
The on-demand deployment service for the INDIGO IAM provides a quick and easy way for organizations to deploy their own instance of the INDIGO IAM, which is an open-source Identity and Access Management system...

CONFIGURE →

Jupyter with persistence for Notebooks
Run Jupyter on a single VM enabling Notebooks persistence

CONFIGURE →

Elasticsearch and Kibana
Deploy a virtual machine pre-configured with the Elasticsearch search and analytics engine and with Kibana for simple visualization of data with charts and graphs in Elasticsearch

CONFIGURE →

Jupyter + Matlab (with persistence for Notebo...
Run Jupyter on a single VM enabling Notebooks persistence and Matlab integration

CONFIGURE →

Kubernetes cluster
Deploy a single-master Kubernetes cluster

CONFIGURE →

Sync&Share aaS
The INFN Cloud Sync&Share aaS is based on popular storage solutions such as openCloud and NetStorage. INFN Cloud users have full control over the configuration parameters of their Cloud...

CONFIGURE →

Dashboard

DEPLOYMENTS
ADVANCED
EXTERNAL LINKS
ADMIN

ON-DEMAND SERVICES

Virtual machine
Launch a compute node getting the IP and SSH credentials to access via ssh

CONFIGURE →

Docker compose
Deploy a virtual machine with docker engine and docker compose pre-installed. Optionally run a docker compose file fetched from the specified URL.

CONFIGURE →

Run docker
Run a docker container

CONFIGURE →

INDIGO IAM as a Service
The on-demand deployment service for the INDIGO IAM provides a quick and easy way for organizations to deploy their own instance of the INDIGO IAM, which is an open-source Identity and Access Management system...

CONFIGURE →

Elasticsearch and Kibana
Deploy a virtual machine pre-configured with the Elasticsearch search and analytics engine and with Kibana for simple visualization of data with charts and graphs in Elasticsearch

CONFIGURE →

Kubernetes cluster
Deploy a single-master Kubernetes cluster

CONFIGURE →

Settings
Help

Marica Antonacci
admins/beta-testers

Dashboard

Marica Antonacci
admins/beta-testers

SSH keys management

SSH keys allow you to establish a secure connection between your computer and your virtual server(s).

UPLOAD SSH PUBLIC KEY
Paste your public SSH key, which is usually contained in the file ~/.ssh/id_ed25519.pub or ~/.ssh/id_rsa.pub and begins with 'ssh-ed25519' or 'ssh-rsa'. Don't use your private SSH key.

UPLOAD

CREATE NEW KEY PAIR
SSH key pair will be created from scratch. The private key will be safely stored in the Vault, while the public key will be stored in the Dashboard database.

+ CREATE NEW SSH KEY PAIR

My deployments

Refresh New deployment +

Show 10 entries Search:

DESCRIPTION	DEPLOYMENT IDENTIFIER	STATUS	CREATION TIME	DEPLOYED AT	ACTIONS
test	11ef6f81-8f57-e13d-8583-22533e954eeb	CREATE_COMPLETE	2024-09-10 14:32:00	BACKBONE-CNAF	Details
test	11ef5e0b-0a6a-fe9b-b005-3a4ac66f6def	CREATE_COMPLETE	2024-08-19 09:11:00	BACKBONE-CNAF	Details
k8s	11ef2a2f-67b6-61e4-ad50-22533e954eeb	UPDATE_FAILED	2024-06-14 09:20:00	BACKBONE-CNAF	Details
docker build	11ee9e9a-dfcd-c750-9c93-0242ec4ec63f	CREATE_COMPLETE	2024-04-20 10:15:00	RECA-SBARI	Details
k8s lisa	11ee9a9a-dfcd-c750-9c93-0242ec4ec63f	UPDATE_COMPLETE	2023-12-14 16:07:00	BACKBONE-BARI	Details

Showing 1 to 5 of 5 entries

Previous 1 Next

Conclusions

The INFN scientific computing infrastructure is evolving to become interdisciplinary and to cope with heterogenous needs and resources

Strong collaboration with national HPC centres

The PaaS Orchestrator and the web dashboard are key component to provide flexibility and ease of use

The orchestrator is undergoing a complete rewriting to provide sustainability and ability to address new needs

The dashboard scope and design is being extended to support workflows that do not necessarily go through the orchestrator

*Work supported by the Italian Ministry of University and Research PNRR Mission 4, Component 2
ICSC: Investment 1.4, Project code CN00000013 - CUP I53C21000340006
TeRABIT: Investment 3.1, Project code IR0000022 - CUP I53C21000370006*

The background features a deep blue gradient. On the left side, there is a vertical axis of light trails and dots that create a sense of depth and movement, resembling a data stream or a digital tunnel. The text 'Backup slides' is positioned on the right side of the image.

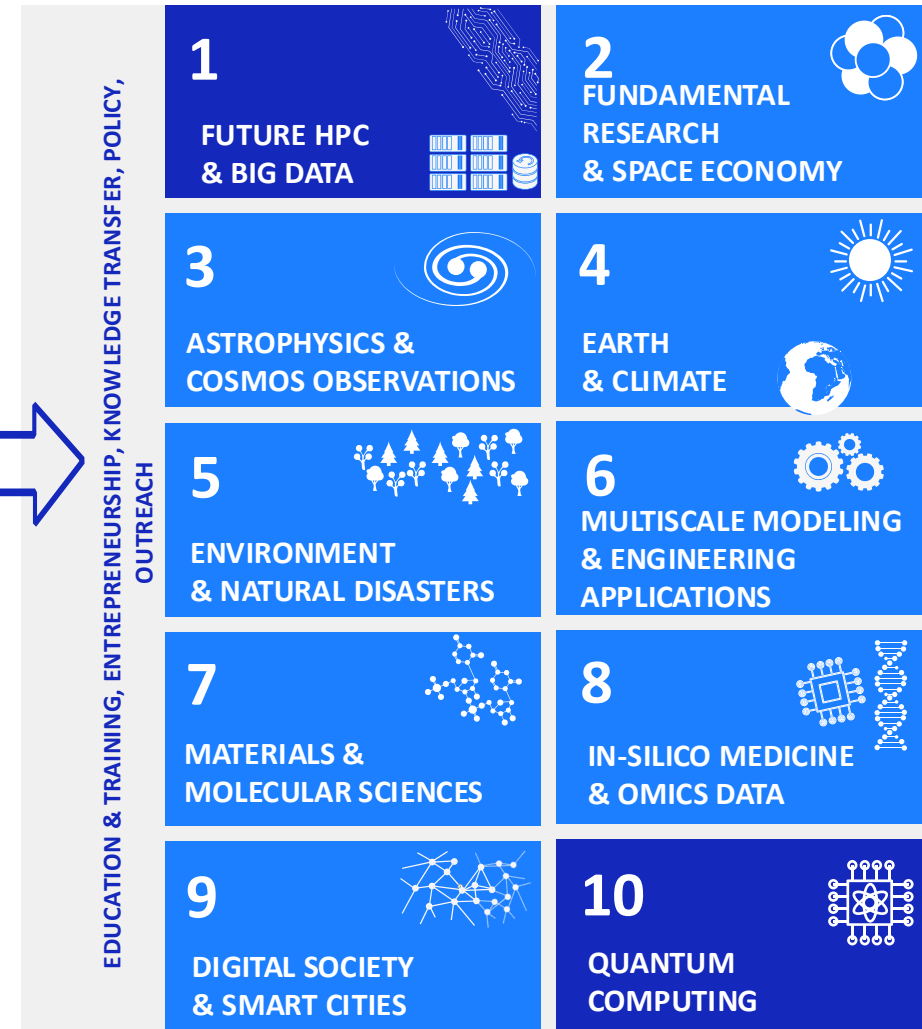
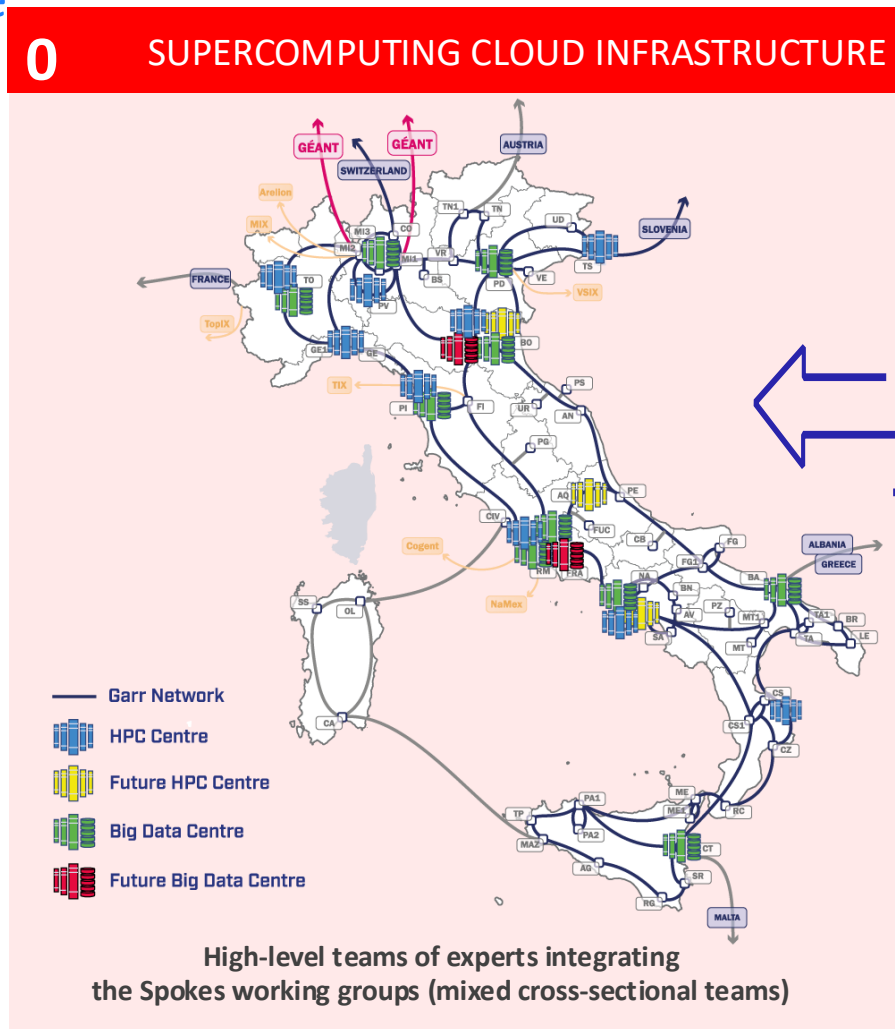
Backup slides

National Research Centre in HPC, Big Data and Quantum Computing

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

25 universities
12 research institutes
14 strategic private companies

320 M€ budget





Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA



<https://www.terabit-project.it/>

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



A data lake for research

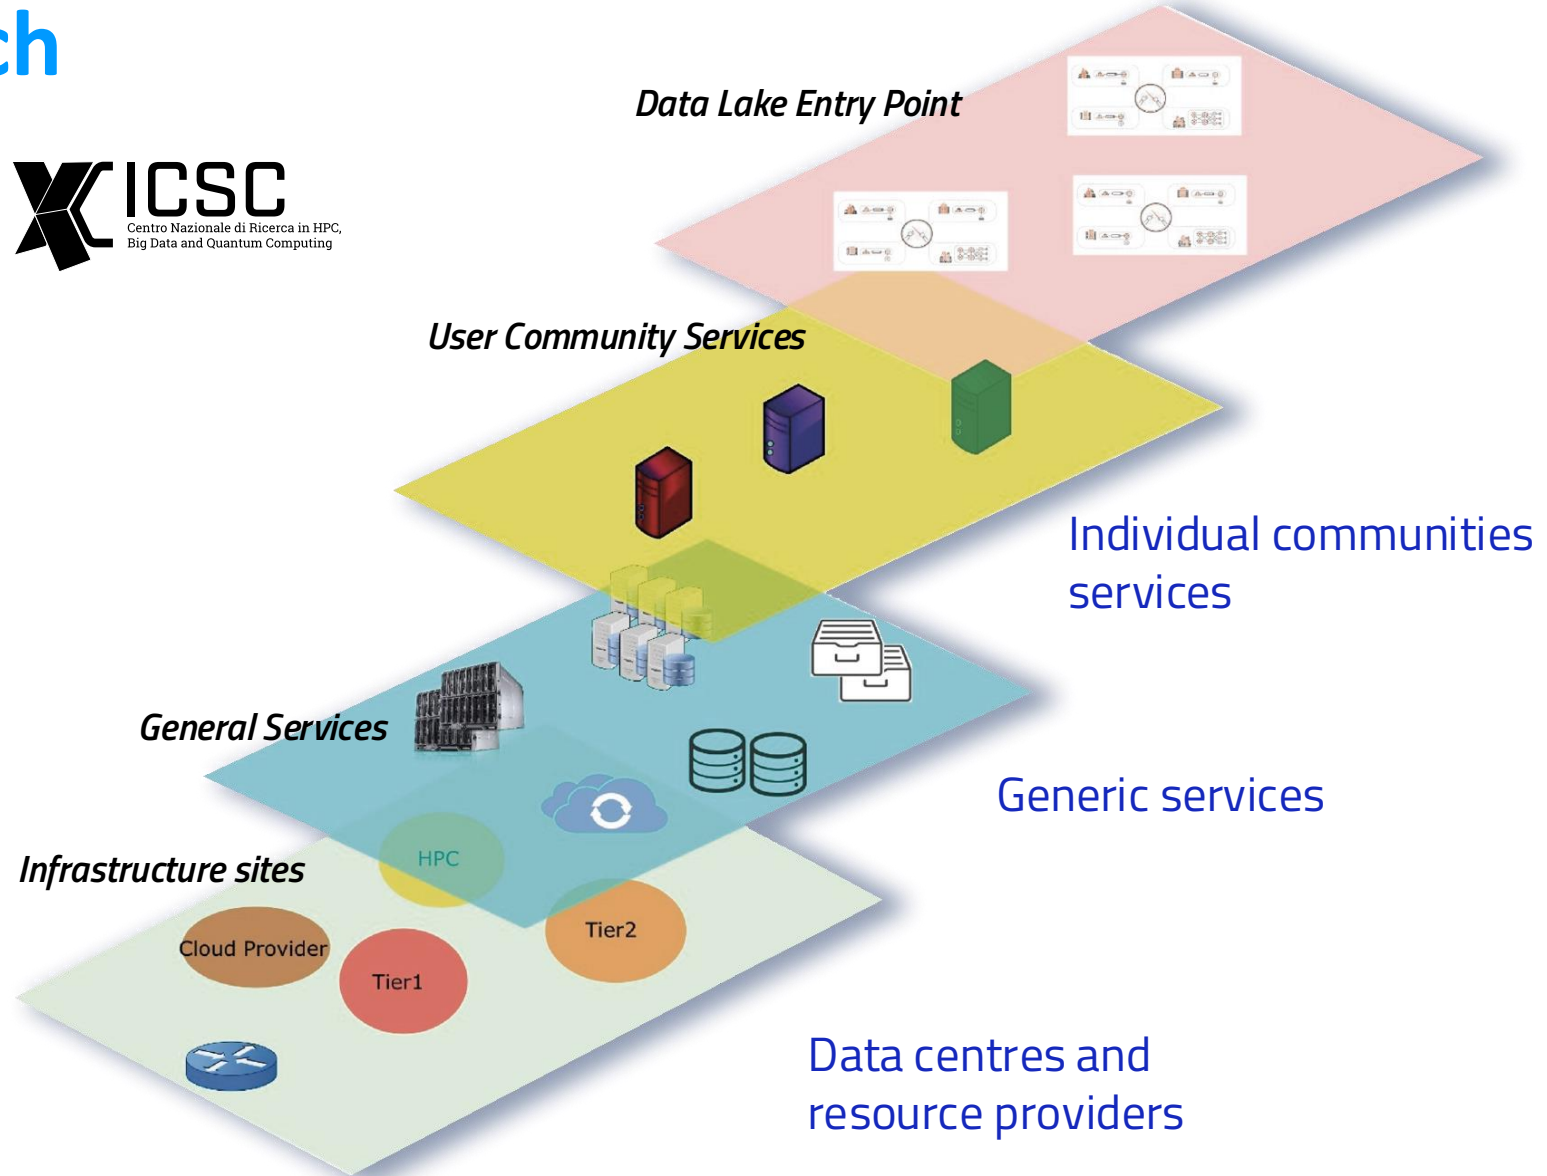
Existing infrastructures aggregation, upgraded and made available to scientific domains

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-centric model

Decouple storage and CPU

Storage nodes interconnected with high bandwidth network

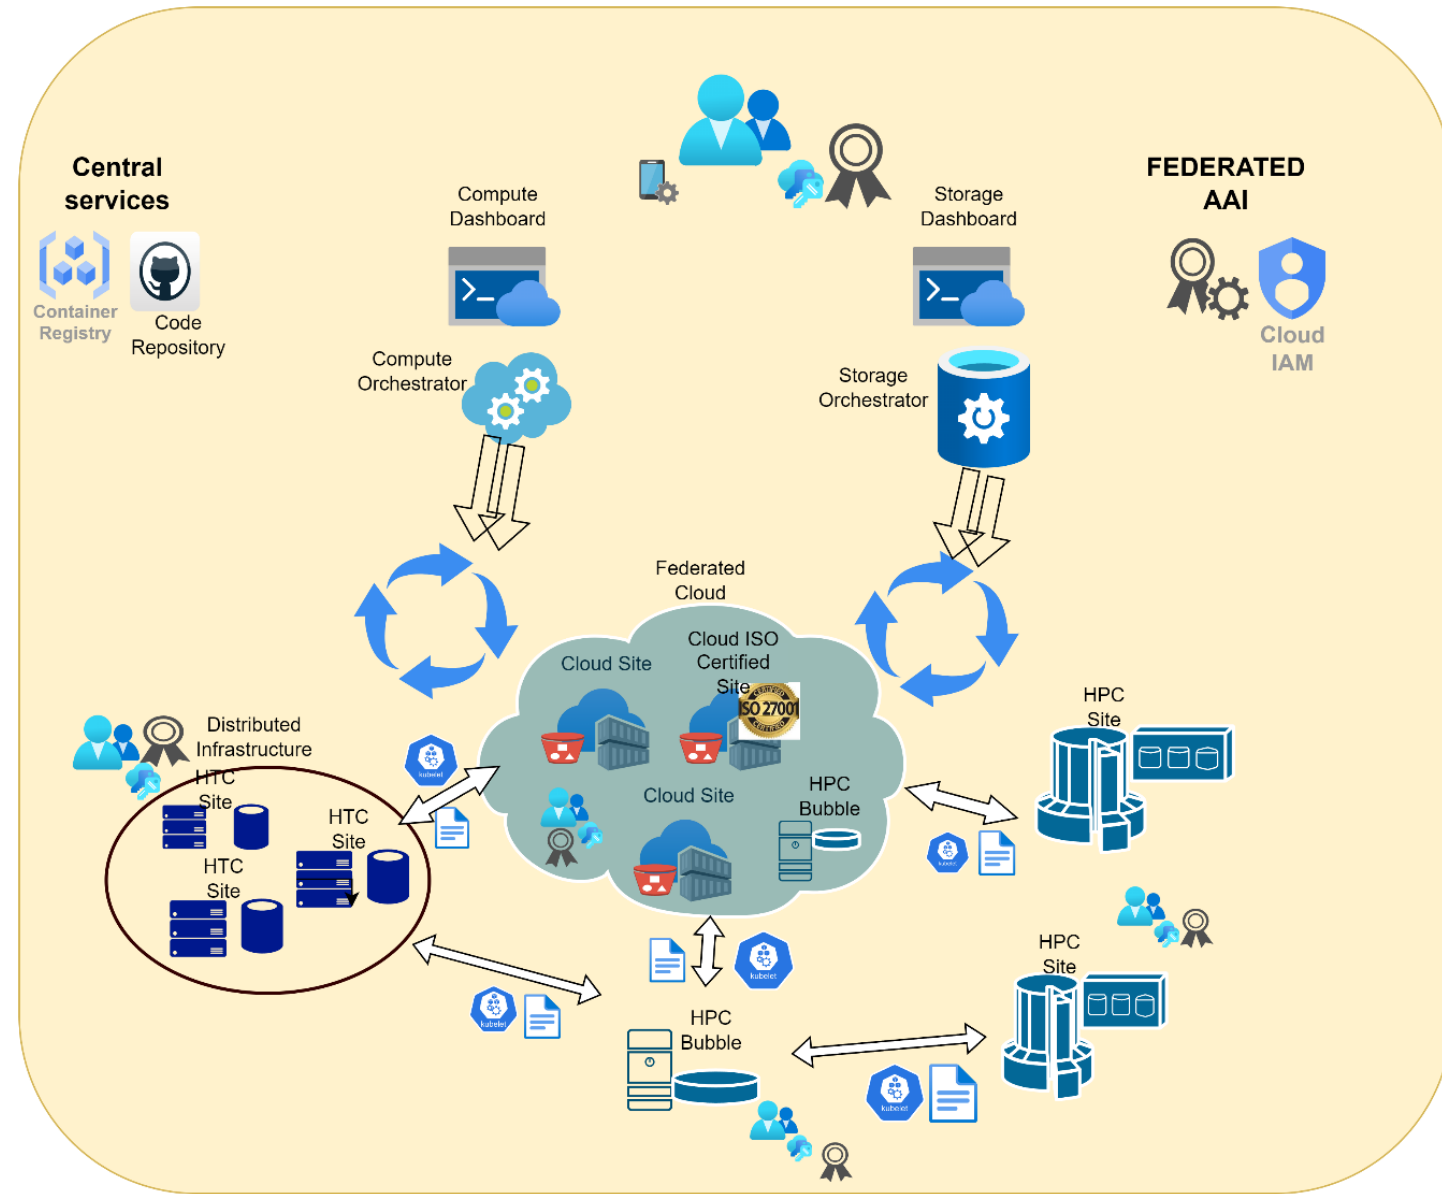
Heterogeneous computing nodes can access data wherever they are





Etherogeneity

Integration of a diverse set of resources, providers, and solutions





Current Orchestrator Architecture

