



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA

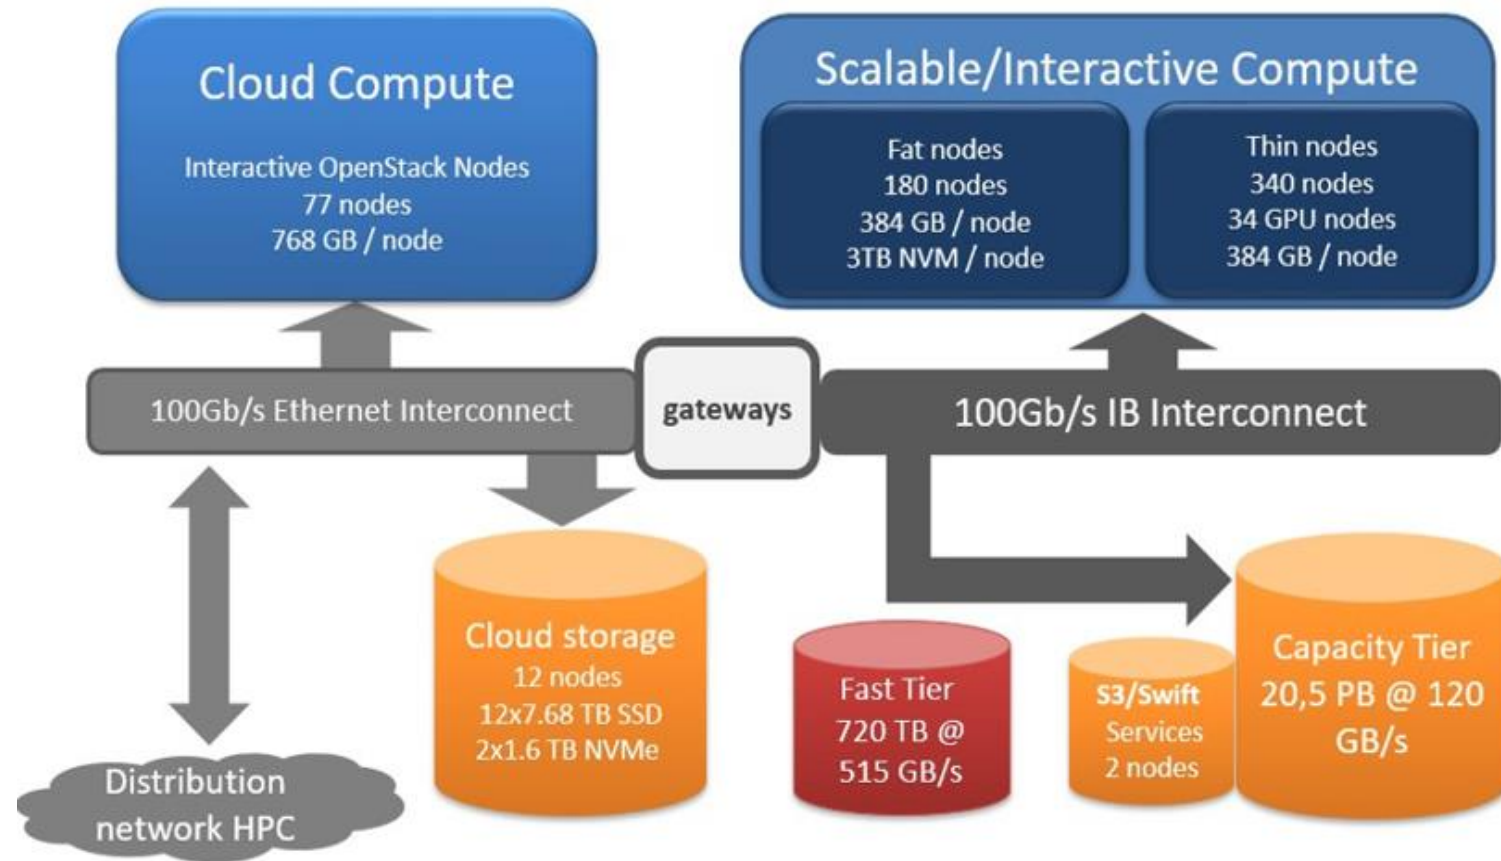




# Servizi e risorse di PRACE-Italy

Giuseppa Muscianisi

Bologna, 26 Giugno 2024

Conferenza TeRABIT





# GALILEO100

<https://www.hpc.cineca.it/systems/hardware/galileo100/>

<https://www.hpc.cineca.it/systems/hardware/ada-cloud/>



### Cloud

- OpenStack cloud
- ~ 7000 vCPUs
- 77 x (2x CPU 8260, 768 GB RAM)
- 100Gbs Ethernet interconnection.

### Storage services

- File (POSIX), Block, Object
- ~ 20 PB parallel FS storage
- ~700 TB block storage for cloud (full flash)
- ~ 700 TB IME burst storage

### Low- mid HPC


- HPC resources
- Fat nodes
- CPU/GPU nodes
- Fast interconnect hi-bw low latency (NVIDIA HDR)

### Interactive comp.

- Prompt response
- Dedicated resources (GPU/CPU)
- Close to data repository
- Framework for processing data (Jupyter, R, Julia, ...)

## CLOUD targets

- Data processing and analysis services **complementing** High Performance Computing services
- **Bridging** HPC infrastructure (Cloud to host front-end services for the user)
- Data management services **exposing** simulation data to the web via community-defined services
- Workload **requiring ISO27001** certification (processing sensitive data)
- Flexible and **automated deployment** via Kubernetes and containerized workflows
- **Collaborative** infrastructure deployment (Infrastructure as Code)
- Everything that requires **performance** and **flexibility**



## Traditional HPC vs CLOUD - HPC

	Traditional HPC resources	CLOUD-HPC integration
<b>Performance</b>	Target the highest possible	Depend on workload, but generally virtualization has small impact
<b>User access</b>	CINECA staff authorization	Once a project is granted, it is managed by the user
<b>Operating System</b>	It is chosen by CINECA staff given the HW constraints. Security updates are managed by CINECA.	Selected by the user. Security patch and updates are managed by the user.
<b>Software stack</b>	Mostly installed by CINECA staff. User can install its own without administration (root) privilege. The environment is provided "as is"	User is "root" on the VMs and can install all the required software stack. Users can modify the environment to suit its needs.
<b>Snapshot of the environment</b>	Cannot be done	User can save snapshot images of the VMs
<b>Running simulations</b>	Users are provided with a job scheduler (SLURM)	Users can plan their workload execution

# CINECA CLOUD HPC HISTORY

PICO



**2015**

**1'600** vCPUs,  
10 TB of storage

OpenStack  
version "**Liberty**"

MEUCCI



**2020**

**5'000** vCPUs,  
200 TB of storage

OpenStack  
version "**Train**"

ADA



**2021**

**7'000** vCPUs,  
1'000 TB of storage

OpenStack  
version "**Zed**"

## CLOUD - HPC Installation

- **Cloud nodes:**


- Zed OpenStack version
- 3 physical nodes as Cloud Controller, HA
- 76 Compute nodes:
  - 71 physical nodes, HT enabled, 96 vCPU each and 798 GB of DDR4 RAM
  - 5 nodes, HT enabled, 96 vCPU each, 384 GB DDR4 RAM, equipped with 2xGPU nVidia V100 PCIe3 with 32GB RAM
- 1024 allocable public IPs
- High-bandwidth low-latency 100 Gb/s ethernet network interconnection
- Installation made by Kolla-Ansible system, based on Ubuntu 22.04 with QEMU-KVM hypervisor

- **Storage, fully dedicated to the cloud infrastructure:**

- Based on CEPH version 15.2.13 (Octopus)
- 12 servers equipped with 1 PB of total raw disk, full SSD/NVMe
  - 750 TB "Raid 6" full flash (SSD/NVMe)
  - Ideal to host IOPS workloads (e.g. DBs, AI)

- **Additional:**

- HW and SW architectures are implemented fully redundant, to avoid SPoF
- No oversubscription of resources is allowed, to address HPC.



7500 vCPUs



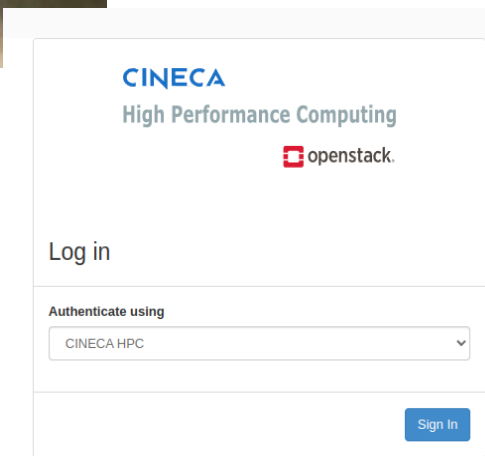
79 Cascade Lake servers  
(2x) upgradable to 150



750 TB Raid 6 CEPH  
20 PB via client/NFS (based on  
DDN Exascaler)

## CLOUD - HPC User Info

- Cloud Compute name: **ADA cloud**
- OpenStack Horizon dashboard: <https://adacloud.hpc.cineca.it>
- User authentication via CINECA IdP, based on keycloak OpenID.
  - 2FA enabled to login in ADA cloud dashboard
  - Additional Fenix IdP is going to be available
- User Guide: [ADA CLOUD CINECA HPC](#)
- Multiple cloud images provided, but users can use their own
- Multiple cloud flavors provided, spanned from 1 up to 96 vCPUs, to allocate the entire compute node capability with a single VM. This is to address power demanding user workflows

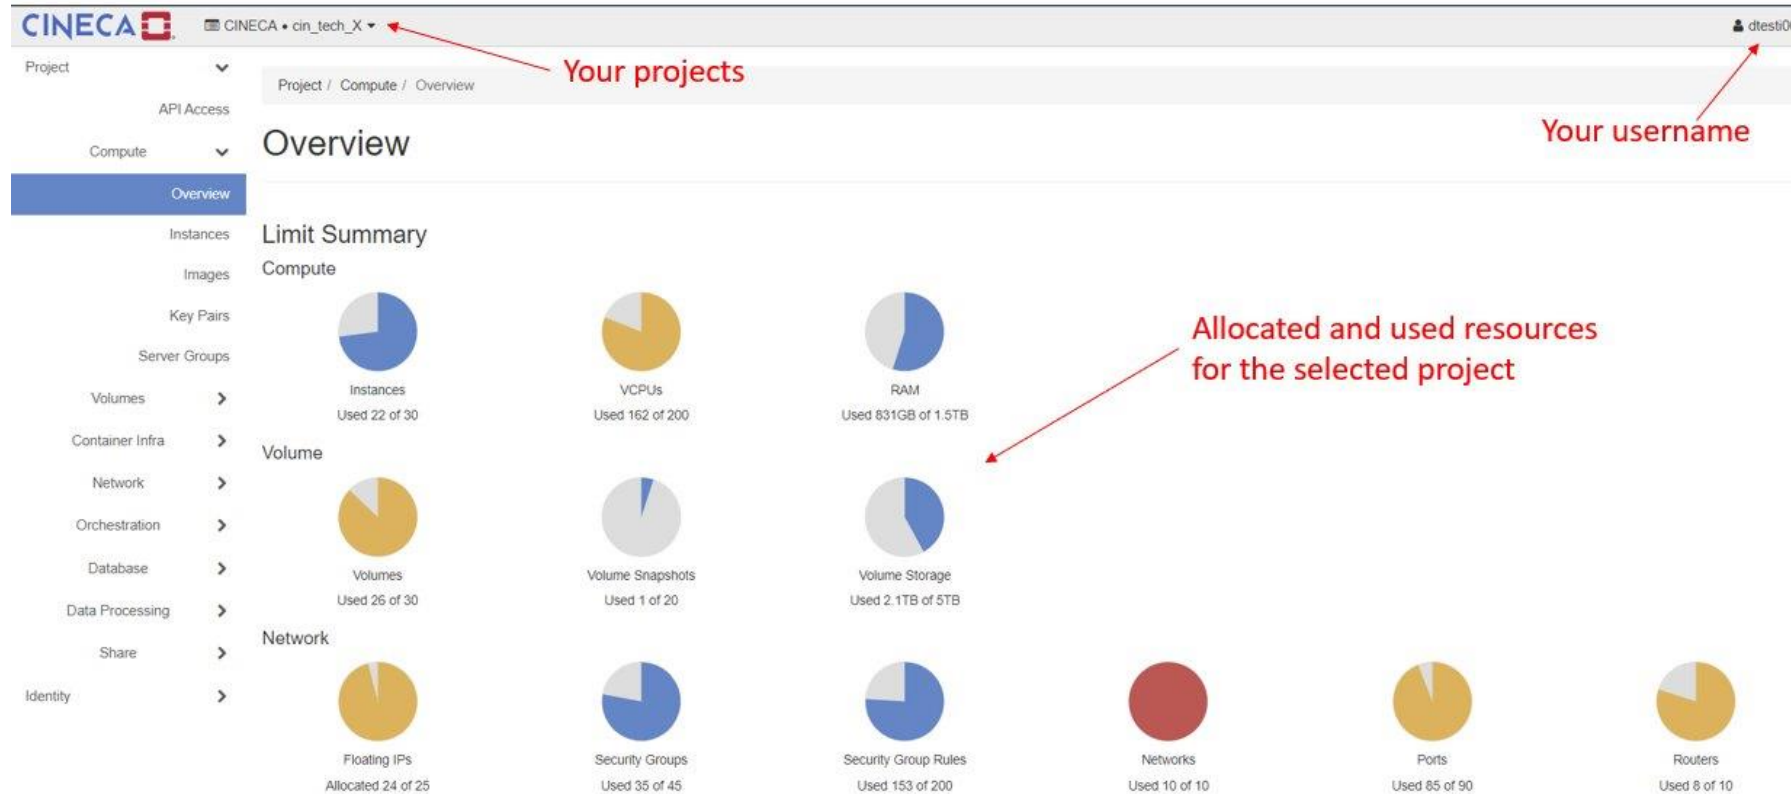


CINECA  
High Performance Computing  
openstack.

Log in

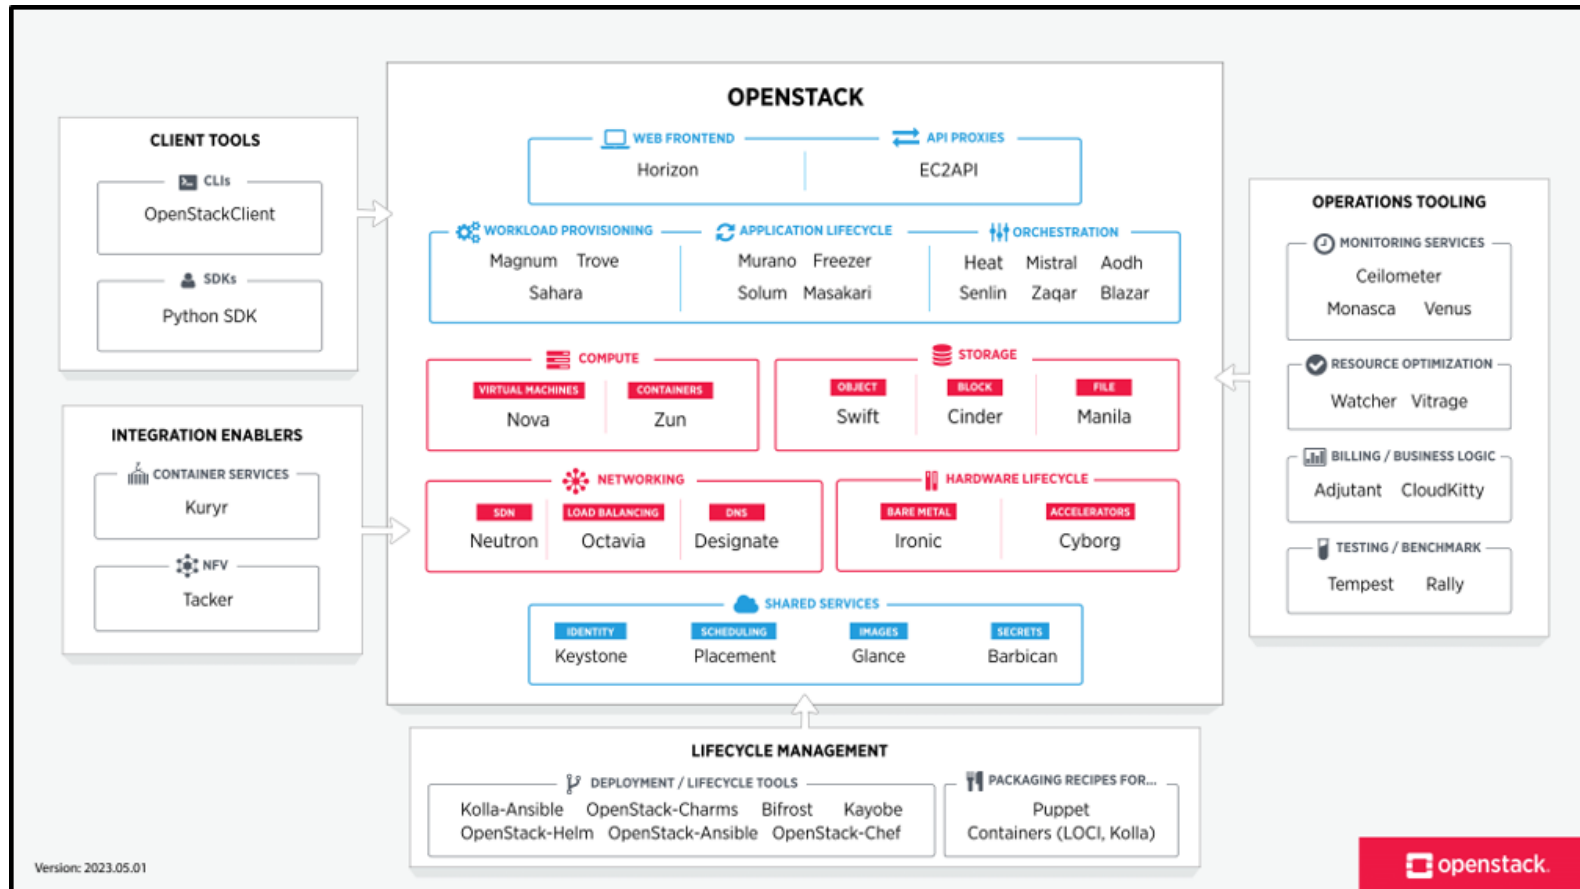
Authenticate using  
CINECA.HPC

Sign In



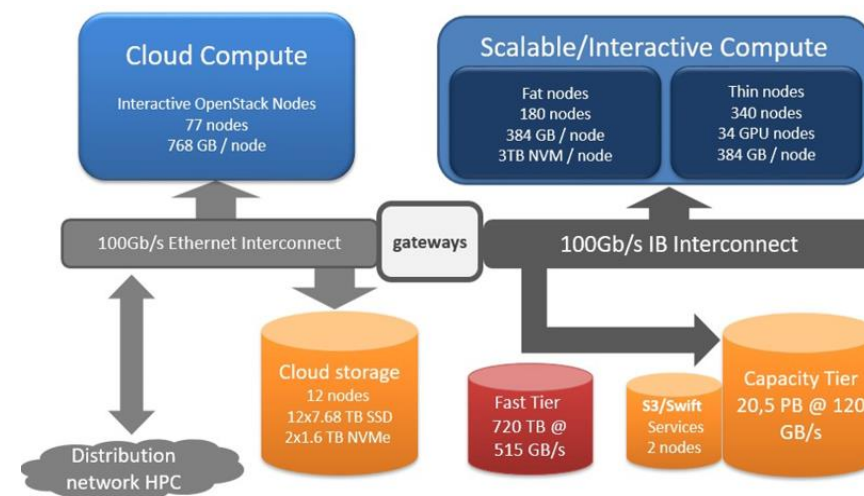


# OPENSTACK services



## GALILEO100 / ADA CLOUD FS SHARING

- Possibility to share a fileset of the LUSTRE 20 PB hot filesystem among Scalable/Interactive Compute and VMs in ADA cloud
- Possibility to move GPU nodes from the Scalable/Interactive Compute partition to the Cloud one to increase flexibility





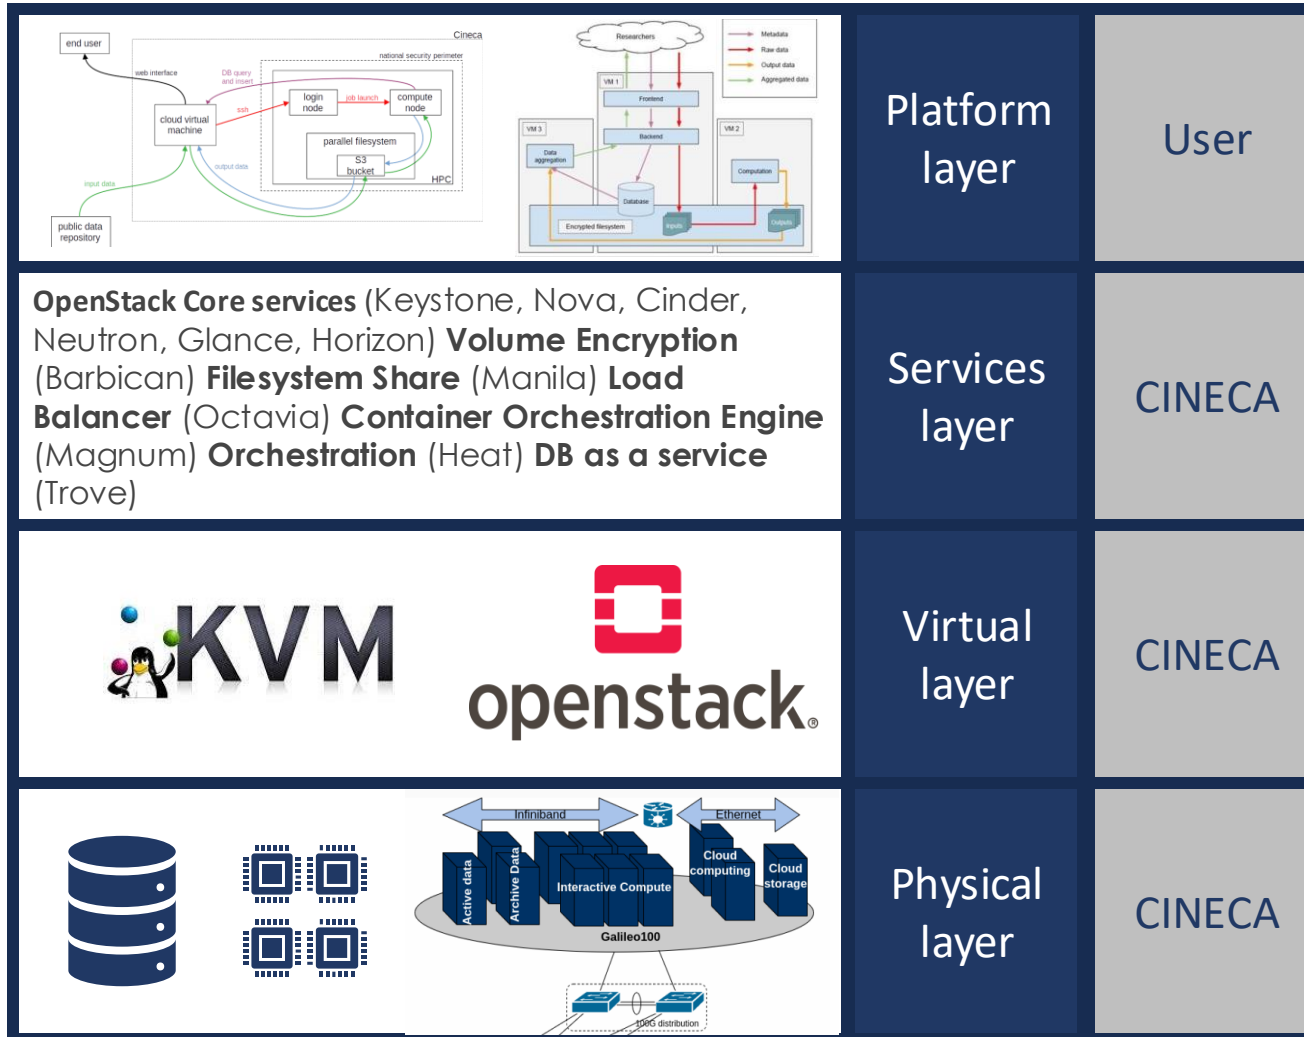
Finanziato dall'Unione europea  
NextGenerationEU



Ministero dell'Università e della Ricerca



Italiadomani  
PIANO NAZIONALE DI RIPRESA E RESILIENZA





Finanziato dall'Unione europea  
NextGenerationEU




Ministero dell'Università e della Ricerca



Italiadomani  
PIANO NAZIONALE DI RIPRESA E RESILIENZA



## RELEVANT USE CASES



Highlander  
High performance computing to support smart land services



Finanziato dall'Unione europea  
NextGenerationEU



Ministero dell'Università e della Ricerca



Italiadomani  
PIANO NAZIONALE DI RIPRESA E RESILIENZA



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



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



- Data processing and analysis services **complementing** High Performance Computing services
- **Bridging** HPC infrastructure (Cloud to host front-end services for the user)
- Data management services **exposing** simulation data to the web via community-defined services
- Workload **requiring ISO27001** certification (processing sensitive data)
- Flexible and **automated deployment** via Kubernetes and containerized workflows
- **Collaborative** infrastructure deployment (Infrastructure as Code)

# WEATHER FORECAST

## COSMO simulations (Fortran+MPI)

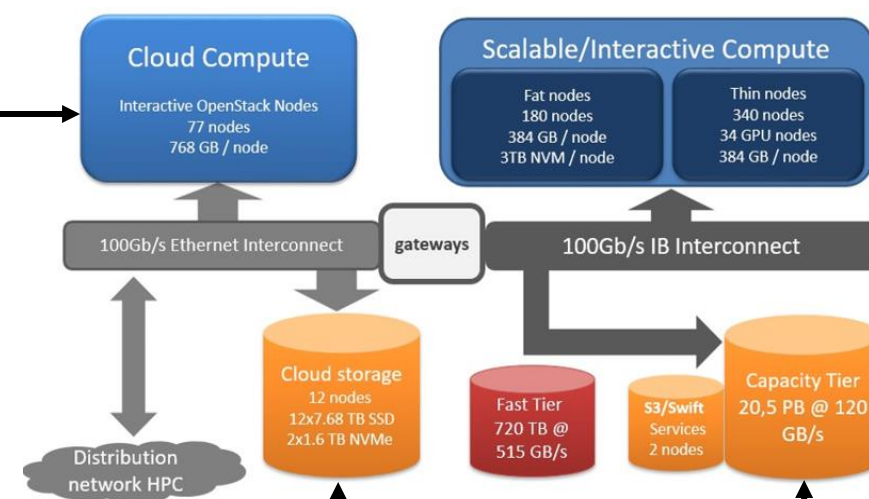
- Grid: 5km Mediterranean, 2,2km Italy
- Simulation time 18-72 hours
- 2 up to 24 runs every day, depend on the model's configuration used

## Hybrid simulations

- Application ported on NVIDIA GPUs
- Evaluation of the stability of application/model still required
- Results need to be periodically checked with CPU simulations
- Post elaboration to create the forecast maps

## Virtual hosting

- Dissemination of data/maps
- VMs + IO
- 250 TB of archival space, tightly integrated
- <https://meteohub.mistralportal.it>





Finanziato dall'Unione europea  
NextGenerationEU



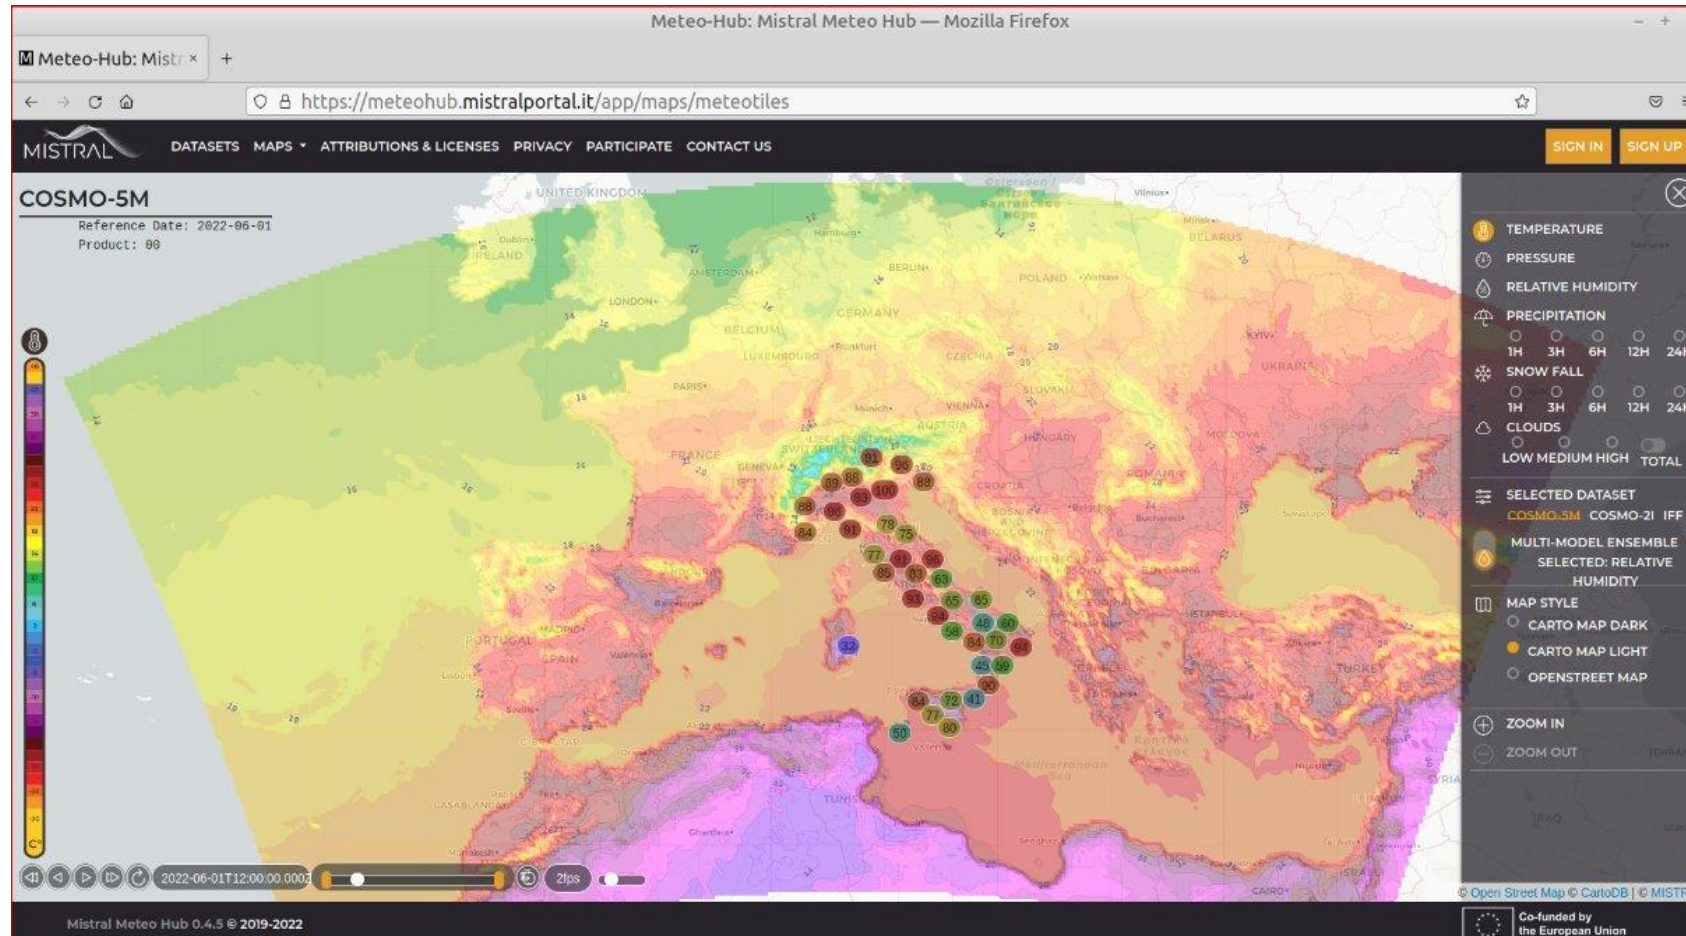
Ministero dell'Università e della Ricerca



Italiadomani  
PIANO NAZIONALE DI RIPRESA E RESILIENZA



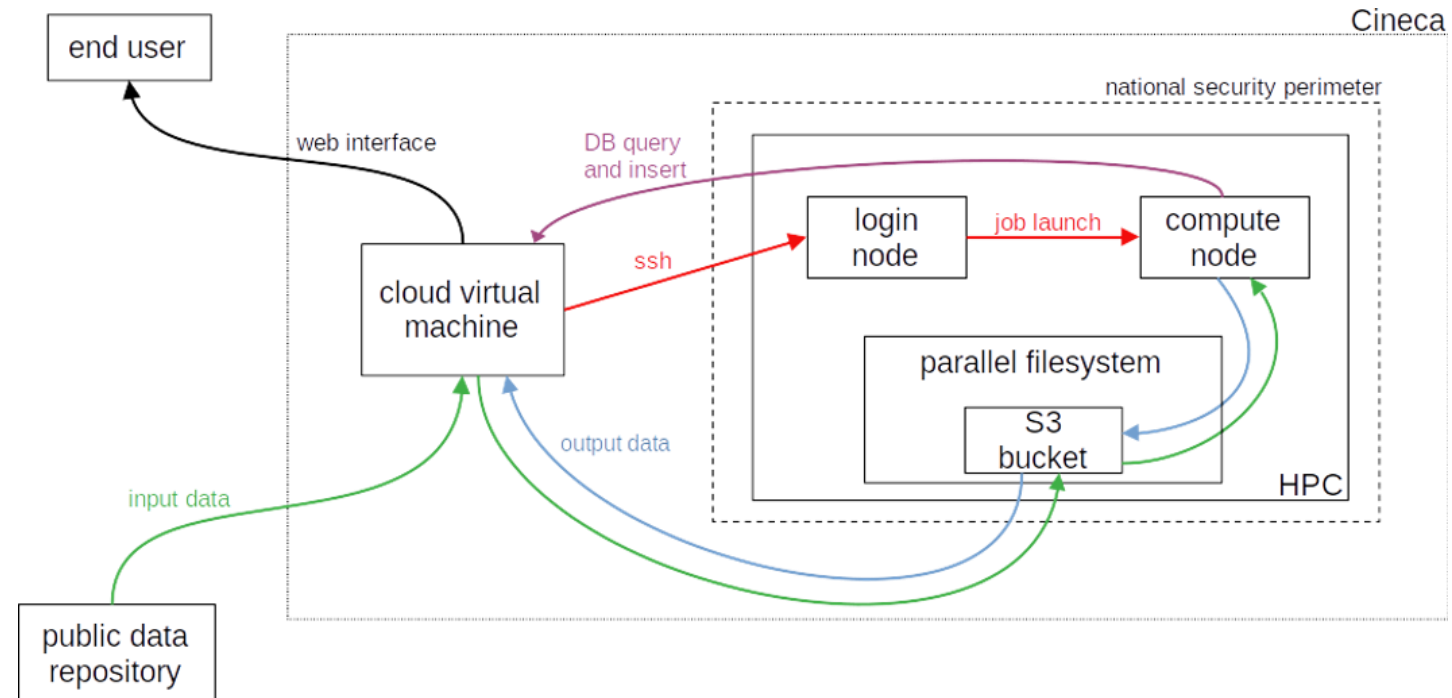
# WEATHER FORECAST






## NOWCASTING

- Hybrid workflow combining **HPC and cloud services** that avoids data duplication thanks to S3
- A virtual machine runs three Docker containers orchestrated by **Docker compose**:
  - one container manages the **data download** and the connection to the HPC
  - one container runs the **database**
  - one container hosts a **website** to display the latest available forecast
- The workflow **downloads weather data** from a remote repository and **submits jobs** to the HPC through ssh





## NOWCASTING

- Dimension of data is **~150MB every 5 minutes**
- **S3** is used to read/write input and output data from both the HPC and the cloud
- **Metadata about weather images** is stored in the non-relational database running in the virtual machine and can be **queried by the HPC job** to obtain desired input data
- The platform is designed to **retrain the artificial intelligence model on demand** with the new data in a perspective of online learning

## Environment Flexibility - InterActiveComputing Service

- Traditional approach to HPC resources:
  - Command line interface on remote machine
  - Scheduler for jobs running the core of the computation on a queue-based system
  - Output is obtained at the end of the job
  - Optimizes resources usage, but not flexible (schedule time unknown, difficult visualization, CLI is not user friendly)
- Interactive Computing approach
  - Main idea: interaction on the fly during workflow execution
  - User can easily check intermediate results and change the workflow accordingly
  - Web interface to allow access via browser to HPC resources and enhance the user experience
- An HPC infrastructure alone is not flexible enough to allow an interactive computing approach, thus it requires an interaction with the cloud environment






Finanziato  
dall'Unione europea  
NextGenerationEU



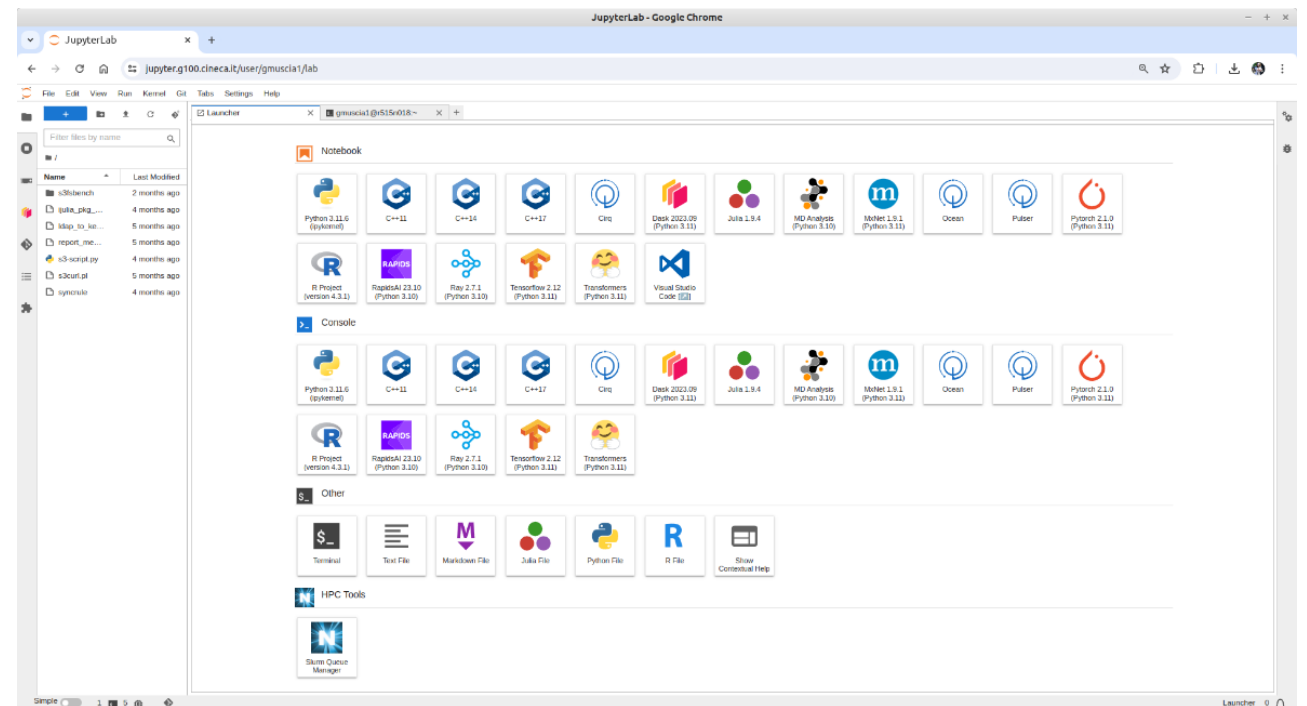
Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA

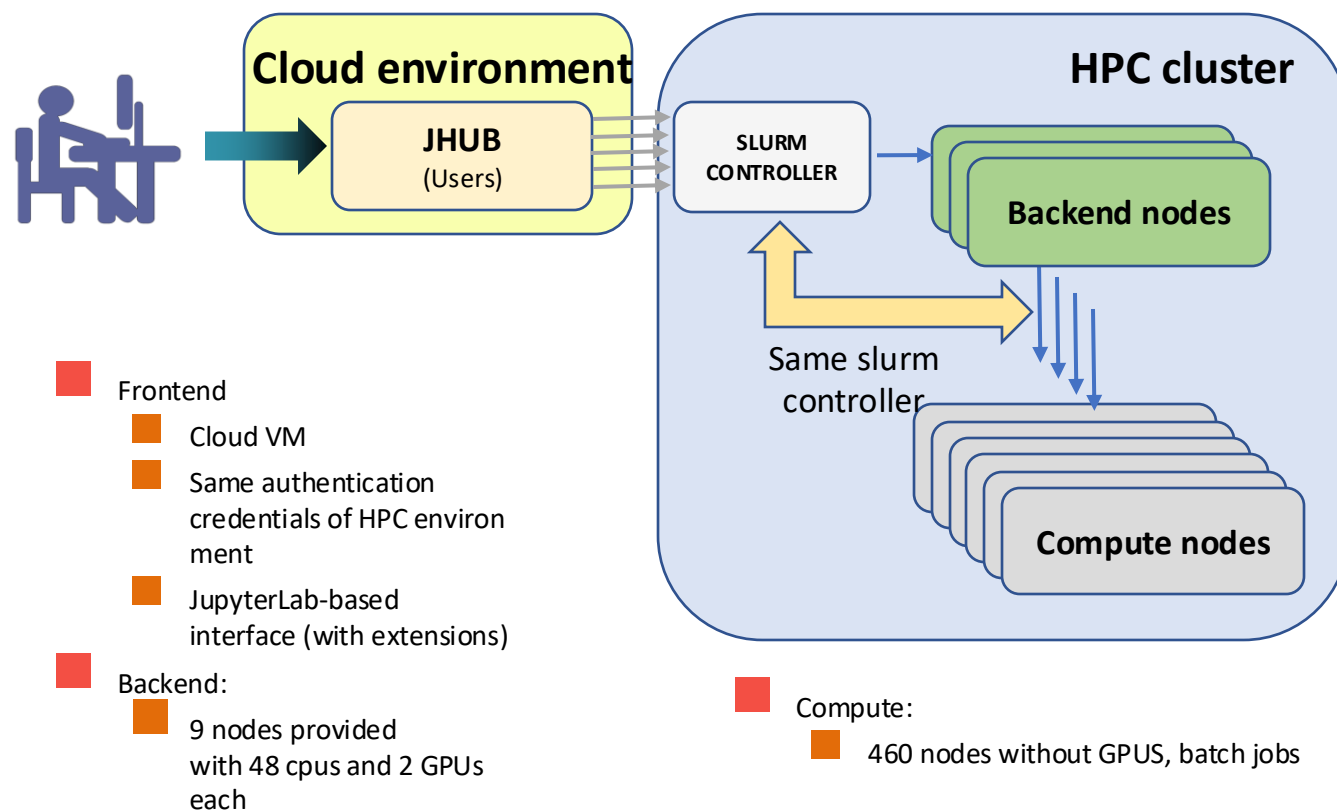


# Environment Flexibility - InterActiveComputing Service



## Environment Flexibility - InterActiveComputing Service

- The framework involves three main components:
  - a frontend layer, based on Jupyterlab
  - a backend layer, running on the HPC cluster
  - Interaction between these two layers happens behind the scenes using Slurm scheduler





Finanziato  
dall'Unione europea  
NextGenerationEU




Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



- Data processing and analysis services **complementing** High Performance Computing services
- **Bridging** HPC infrastructure (Cloud to host front-end services for the user)
- Data management services **exposing** simulation data to the web via community-defined services
- Workload **requiring ISO27001** certification (processing sensitive data)
- Flexible and **automated deployment** via Kubernetes and containerized workflows
- **Collaborative** infrastructure deployment (Infrastructure as Code)



## Sensitive Data – NIG



### Network of different Italian hospitals for Italian Genomes

#### AIMS

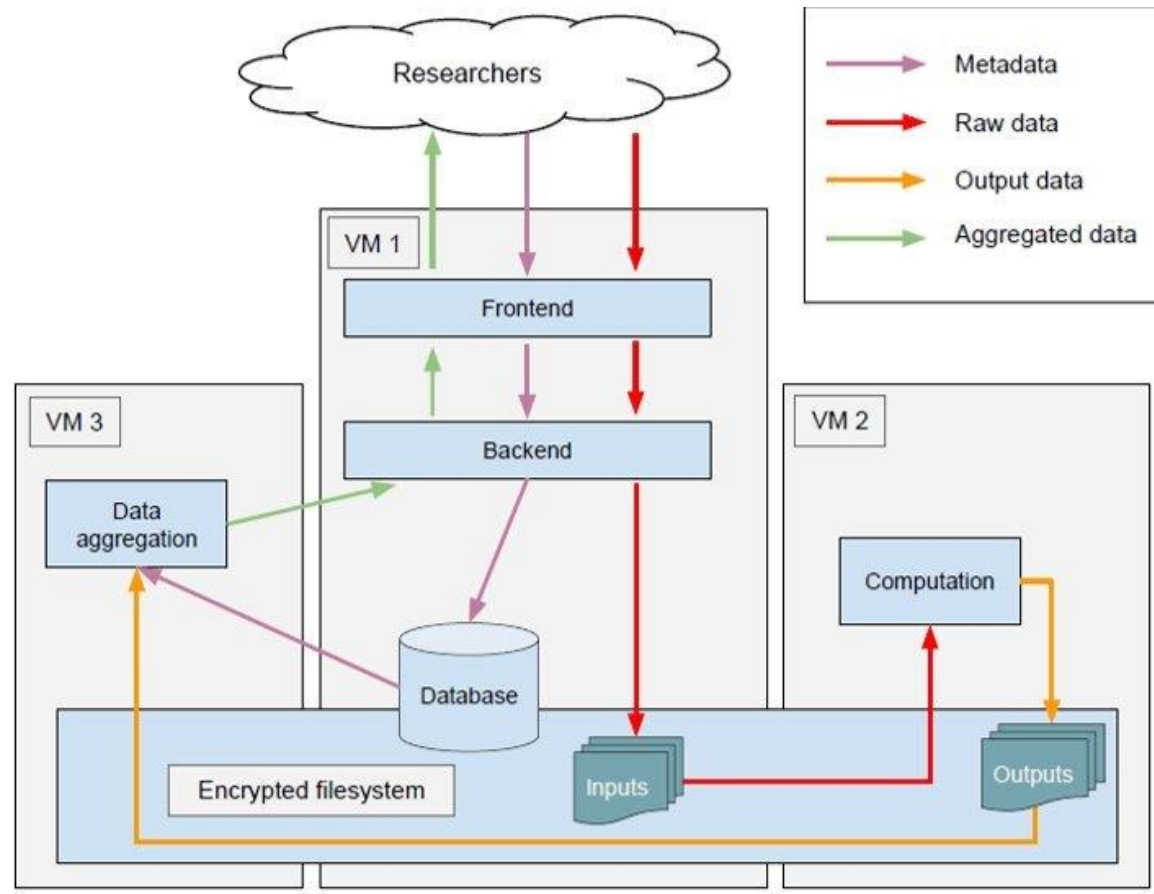
- Centralized repository for human genomic variant data produced by various Centers in Italy
- Characterization of the genomic variability of the Italian population by means of the aggregation of genomic data.
- NIG's interest concerns both pathological and non-pathological variants. The latter are equally important to be able to obtain the typical genome of the Italian population to be used as a reference.
- The data that are noted are: place of birth, sex, age of the patient, comorbidities

#### FORMAL DOCUMENTS

- Specific agreement is signed between CINECA and each single Center
- CINECA has to be appointed as data processor by data controller



## Sensitive Data – NIG



### SECURITY MEASURES ADOPTED

- Minimization of data
- Avoid identifying information (name, birth date, ...)
- Virtual Machine with encrypted volume
- Two factors authentication
- Protection against malware, loss or exposure
- Measures for data breaches
- SSH access only from selected internal CINECA IPs

## Sensitive Data – ORCHESTRA



### Connecting European Cohorts to Increase Common and Effective Response to SARS-Co V-2 Pandemic

The ORCHESTRA project aims to deliver **sound scientific evidence** for the **prevention and treatment** of the infections caused by SARS-CoV-2 assessing epidemiological, clinical, microbiological, and genotypic aspects of population, environment and socio-economic features.

#### Challenge

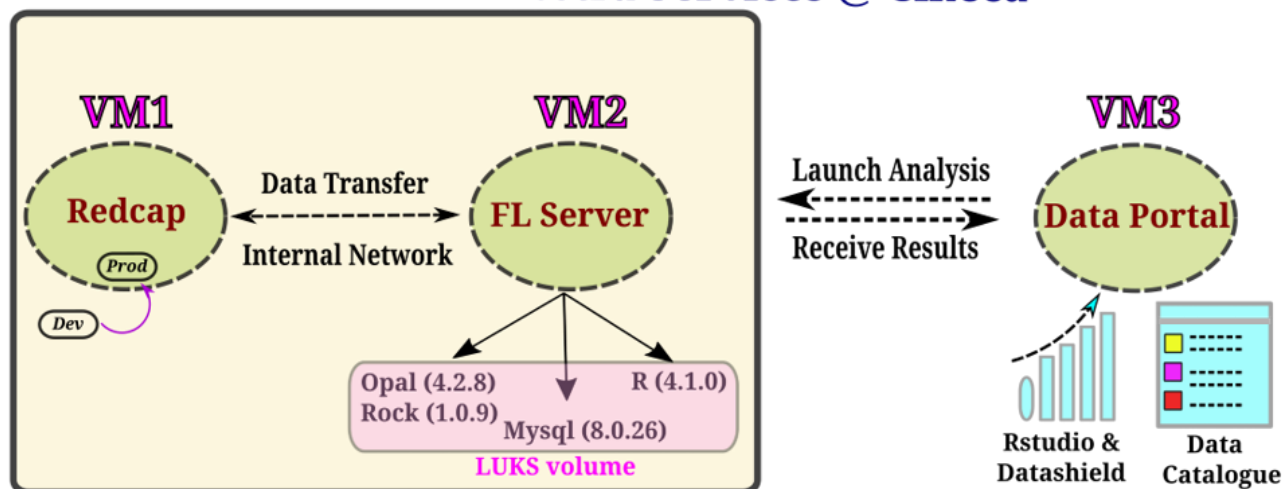
1. Collection, harmonization and standardization of large number of datasets.
2. Secure storage, access and analysis of sensitive data.

#### Technical-components

1. Redcap: Secure web application for building and managing online surveys and databases.
2. Opal: Enable federated analysis of sensitive data without having access to individual values.
3. Data portal: Interface to data analyst to perform analysis and get aggregated results.

## Sensitive Data – ORCHESTRA

### Orchestra Services @ Cineca



### Technical components

- Sensitive data archived at Redcap and Opal servers are stored in encrypted **LUKS volume** and the data transfer is enabled through the **internal network** to avoid any data leakage.
- Analysts are provided with tokens for datasets of interest to perform analysis and get aggregated statistics "without access to individual data values".



## OMICS PLATFORM

CINECA has built OMICS, an **user-friendly web interface** designed for seamless **uploading and processing of genomic** data through NVIDIA Clara Parabricks, a software suite of GPU-accelerated genome analysis tools for genomic read alignment, variant calling, and GVCF filtering and post-processing.

### Main Features

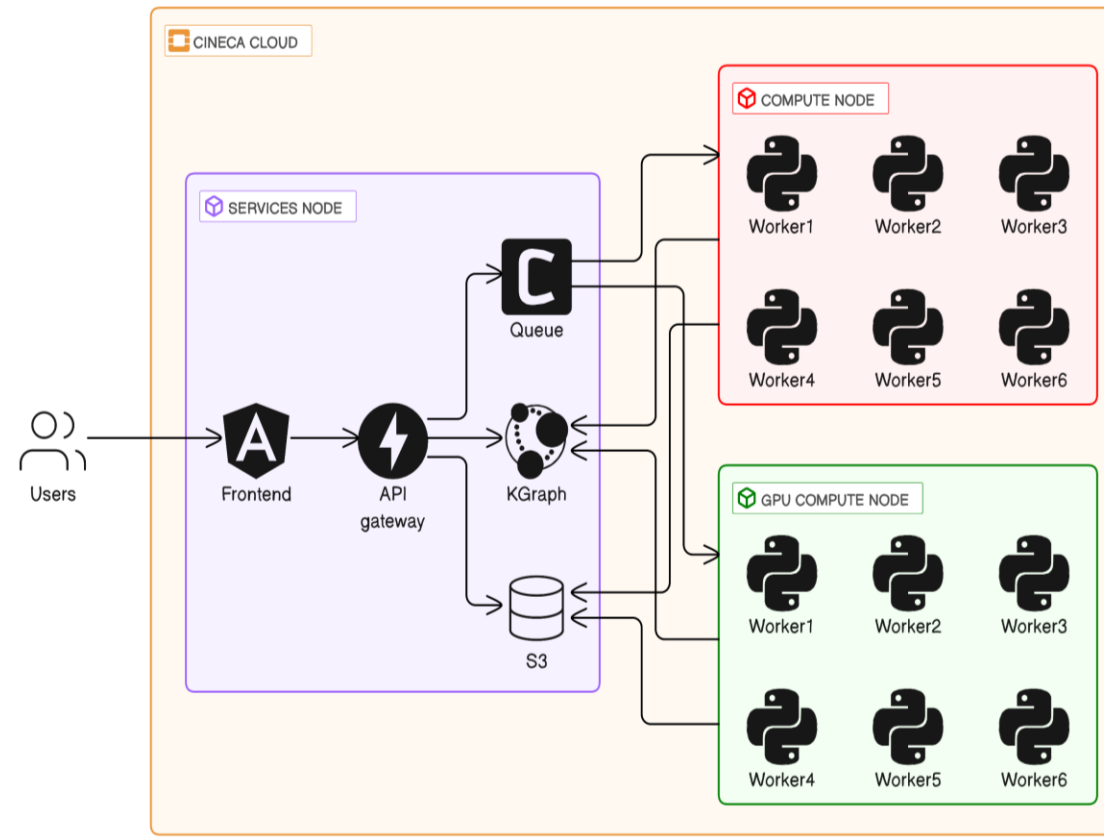
- Data security and privacy compliance with EU GDPR regulations are maintained through a secure encrypted cloud Virtual Machine (VM)
- VM is equipped with GPUs to make the analyses increasingly scalable





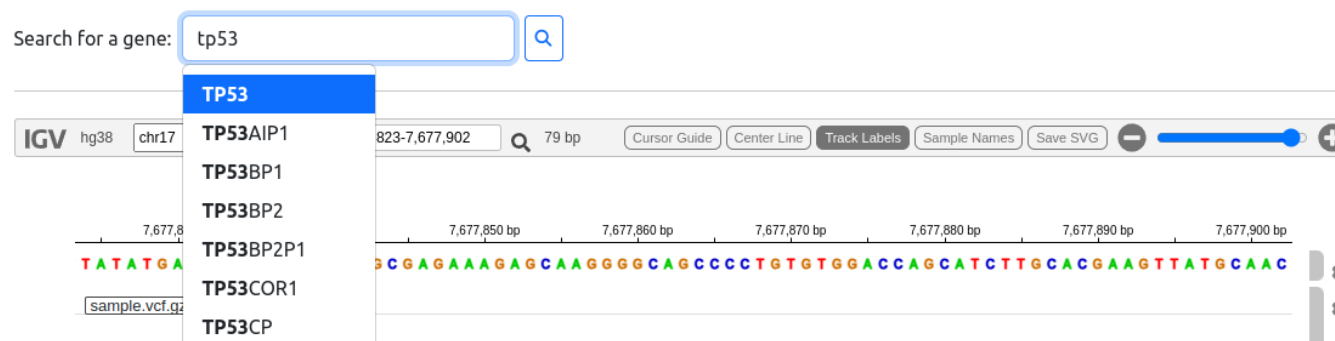
# OMICS PLATFORM Design

- **Open Source** techs
- Microservices oriented
- Cloud ready
- Easy portability
- Easy scalability
- **Private Cloud** resources



## OMICS PLATFORM

- OMICS main dashboard allows users to effortlessly upload their raw data and customize advanced options for three containerized pipelines
- Capability to swiftly upload and download large datasets (e.g. raw sequence data)




- OMICS aims to streamline WGS data analysis while prioritizing user experience, data privacy, and accelerated processing through GPU technology, aiming at supporting activities of Spoke 0 and Spoke 8 of ICSC



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



- Data processing and analysis services **complementing** High Performance Computing services
- **Bridging** HPC infrastructure (Cloud to host front-end services for the user)
- Data management services **exposing** simulation data to the web via community-defined services
- Workload **requiring ISO27001** certification (processing sensitive data)
- Flexible and **automated deployment** via Kubernetes and containerized workflows
- **Collaborative** infrastructure deployment (Infrastructure as Code)


## EO4EU - Multi-cloud Kubernetes Platform



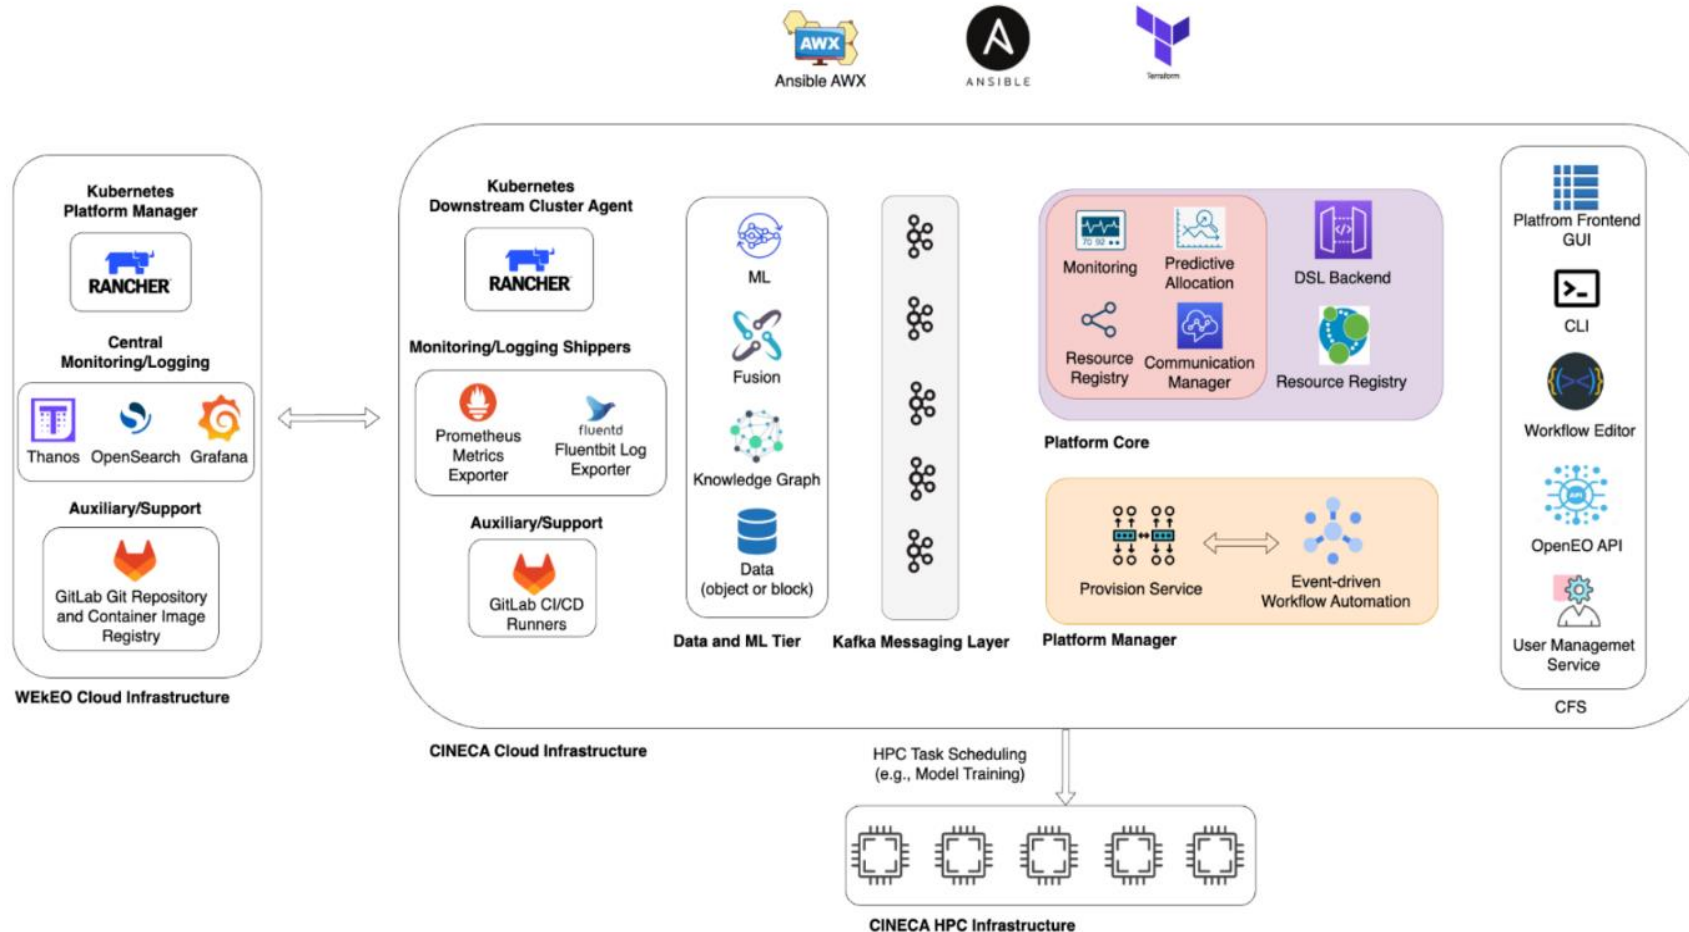
Earth Observation for EUrope

<https://www.eo4eu.eu/> | Horizon Europe 2022

- **EO4EU Platform:** AI-enabled easy access to Earth observation data.
  - Data access, processing (e.g., workflow marketplace), visualization & XR (extended reality, i.e. augmented and virtual reality).
- **ADA & WEKEO** (EU Copernicus DIAS, ECMWF) cloud infrastructures.
- **Deployment automation:**
  - **IaC:** Terraform, Ansible and Kubernetes
  - **& GitOps methodologies:** GitLab pipelines (CI) and Rancher fleet for the deployment of the Kubernetes cluster (CD)



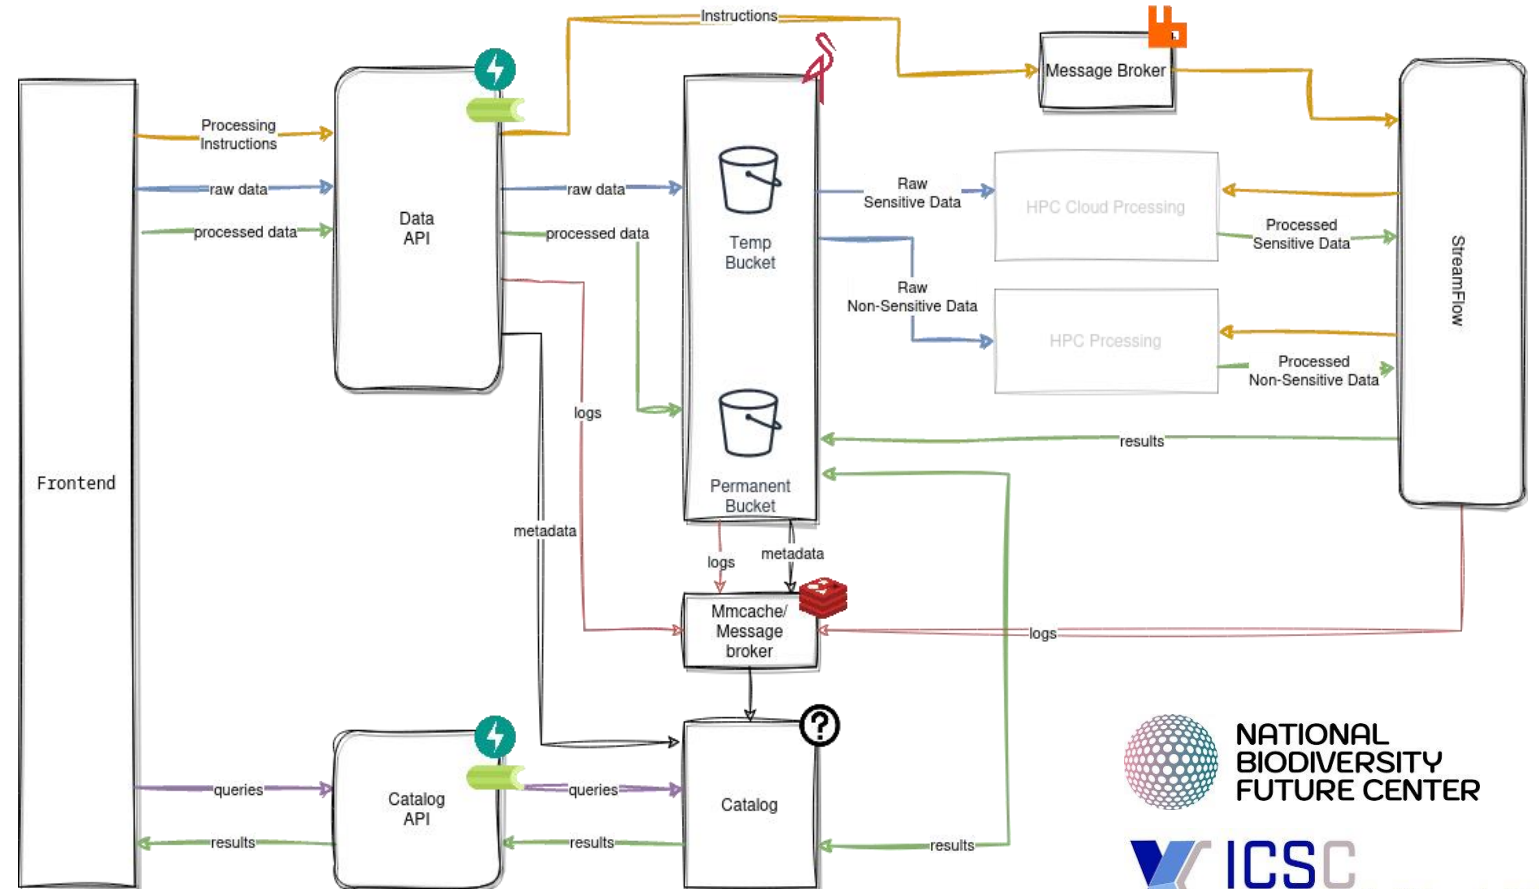
# EO4EU - Multi-cloud Kubernetes Platform





## NBFC & CNHPC Scalable Cloud Platforms

- **NBFC Platform:** Biodiversity Platform to analyse and store a wide range of data types
- **CNHPC Platform:** Platform to process and analyse sensible Genetic Data
- Cloud **certified** for storage of highly **sensible data** (DNA)
- Use of **GPUs** on Ada Cloud for processing of DNA
- Seamless deployment through the use of **Ansible** and **Gitlab-Runners** of all infrastructure (incl. GPUs and LUKS volumes)






Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



# Thanks for your attention

CINECA HPC Cloud Support gang

