Running Interactive analysis on INFN Cloud Platform

Diego Ciangottini

On behalf of the INFNCloud team

Credits to **F. Fanzago and D. Salomoni** for the “inspirational” material

Outline

- What is INFN Cloud?
- An on-demand ecosystem for interactive analysis
- Implementations

Italian Institute for Nuclear Physics
INFN

The **5 research lines**

and the INFN National Scientific Committees

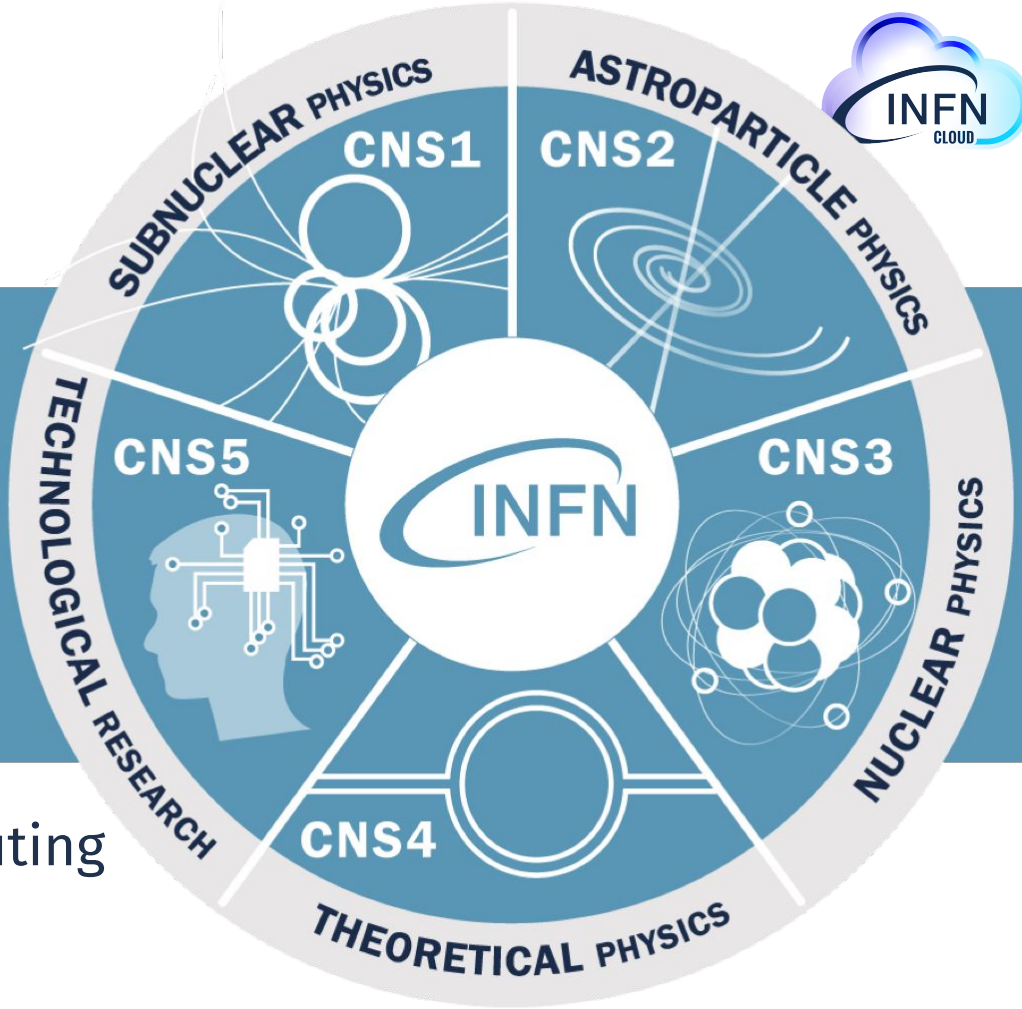

And we learnt early that computing is a transversal need....

The INFN **Facilities**

- 4 National Laboratories
- 20 Divisions
- 6 Associated groups
- 3 National Centres and Schools
- 1 International consortia

Not only LHC…
Regardless, a long time experience

- INFN Distributed Computing federation delivers the LHC experiments O(7-20)% of their computing budgets
- "non LHC" (VIRGO, Astro, Nuclear, …) is ~ 10-20% of the total
- Sites are of top quality among their peers, and have worked uninterruptedly for the last ~15 years
- In many cases, the infrastructures are close to those deployed in 2005 or slightly after
  - They could be bigger (CNAF uses ~1 MW, more recent centers are leaning towards ~10 MW)
  - They could be "greener"
    - Free cooling, direct cooling, …
  - They could use more recent technologies, hardware and software
    - Go towards a national cloud, implement a datalake model for storage, …

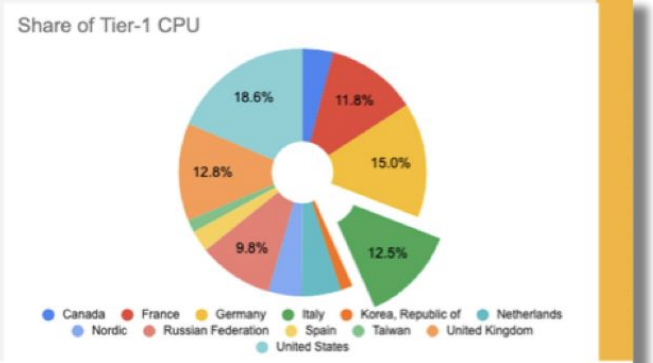

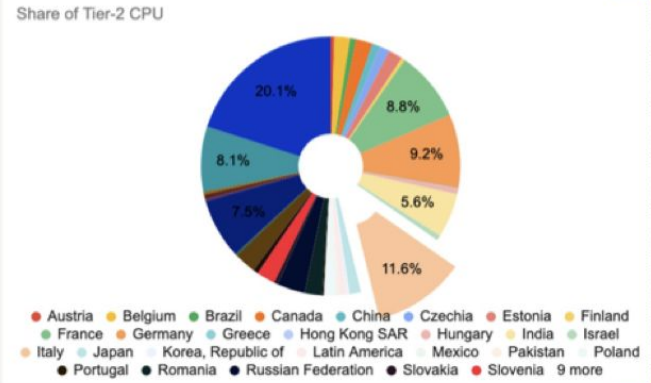

An ever growing asset

- Gradually shifting from “driven by in-home communities” to interdisciplinary projects
- One main goal: explore and adopt the most effective tools for our users to work with

Of course, at this point, cloud computing could not be left behind.

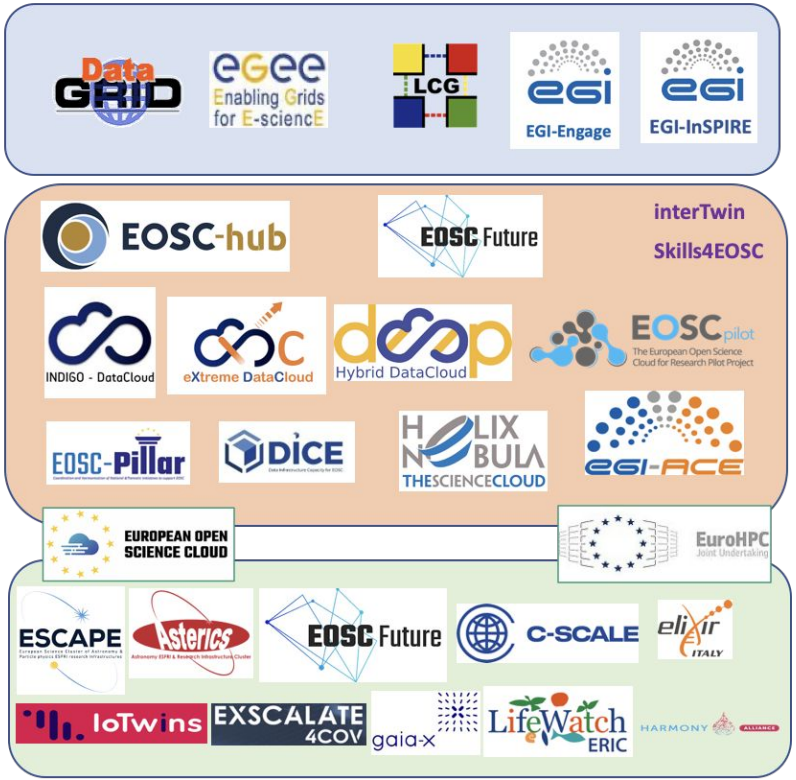

And then, the challenge

Extrapolations for High-Luminosity LHC

- A 20% yearly increase in performance from 2021 to ~2028 gives us a speed-up factor from technology advancements of $(1.2)^7$ **~3.6x**.
- Back-of-the-envelope linear estimates foresee increased needs for computing at High-Luminosity LHC of **~75x** (box on the right).
- Therefore, **we miss a factor 20x**!
- How can we cope with this?
  - Get 20x more money from the Funding Agencies (unrealistic, the ballpark figure would be >> 1 BEur/y)
  - Find ways to reduce the needs…

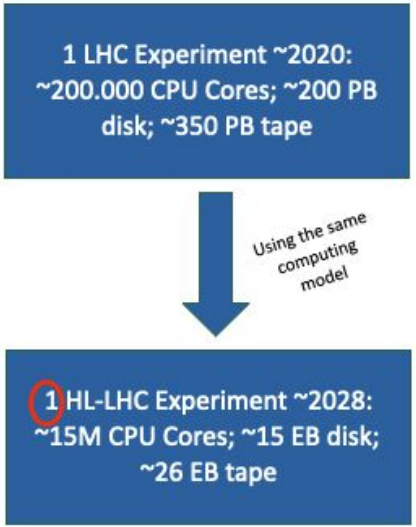

And then, the challenge(s)

Extrapolations for High-Luminosity LHC

- **DUNE**: ProtoDUNE in 2019 collected 3 GB/s (same as CMS at the same time); real DUNE expected 80x at the end of the 2020s.
- **SKA**: up to 2 PB/day (CMS ~3) , to be collected and processed at "complex" locations.
- **Genomics**: a single genome ~100 GB. Any population study (>1M people) over 100 PB.
- **CTA**: ~ 10 PB/y in 2025+.
- **Virgo**: ~10% of a LHC experiment.
- **ET** aiming at being ~10% of a HL-LHC experiment.

One umbrella to collect the answers….

INFN Cloud, https://www.cloud.infn.it/

- In **production** since March 2021.
- The **initial seed** of a National Datalake for research and beyond, building on (existing|renewed|new) e-Infrastructures.
- The **base of the evolution** of the INFN Distributed Computing vision.
- Built on a **thin middleware layer** running on top of *federated clouds*, decoupling physical and logical views via a **service composition** mechanism.
- The **INFN foundation** for the NRRP computing-related initiatives

What is INFN Cloud

A production-quality set of resources and solutions providing:

- A **core backbone**, with ancillary and special-purpose services.
- A **multi-site, federated Cloud infrastructure**.
  - INFN Cloud can transparently federate INFN sites as well as public or private Clouds (e.g.: AWS, Google Compute Cloud, Microsoft Azure, and others)
- A **customizable portfolio of services** accessible via web interfaces, terminal or API.
- A **fully distributed organization for the support and management** of both infrastructure and services.
- A set of **rules that define access resources and policies**, according to INFN, national and European laws.

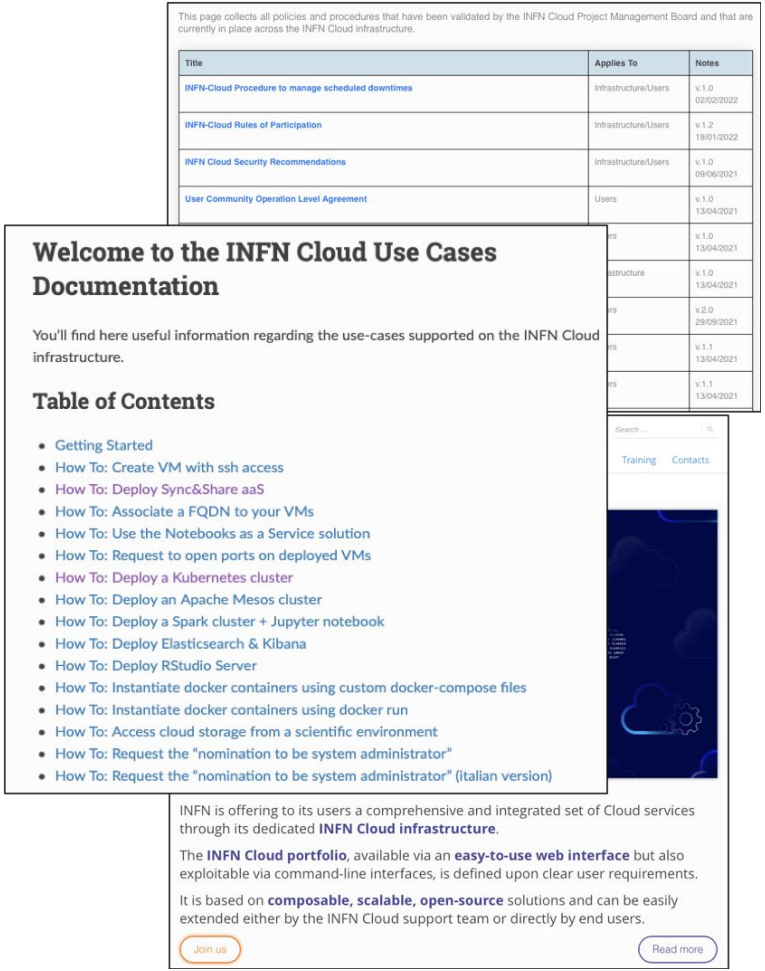

“A backbone to serve them all”

A multi site cloud infrastructure to host the core services

- Geographical HA / failover
- Automation

The idea of federated clouds:

- Resource will be attached to the backbone and made available seamlessly

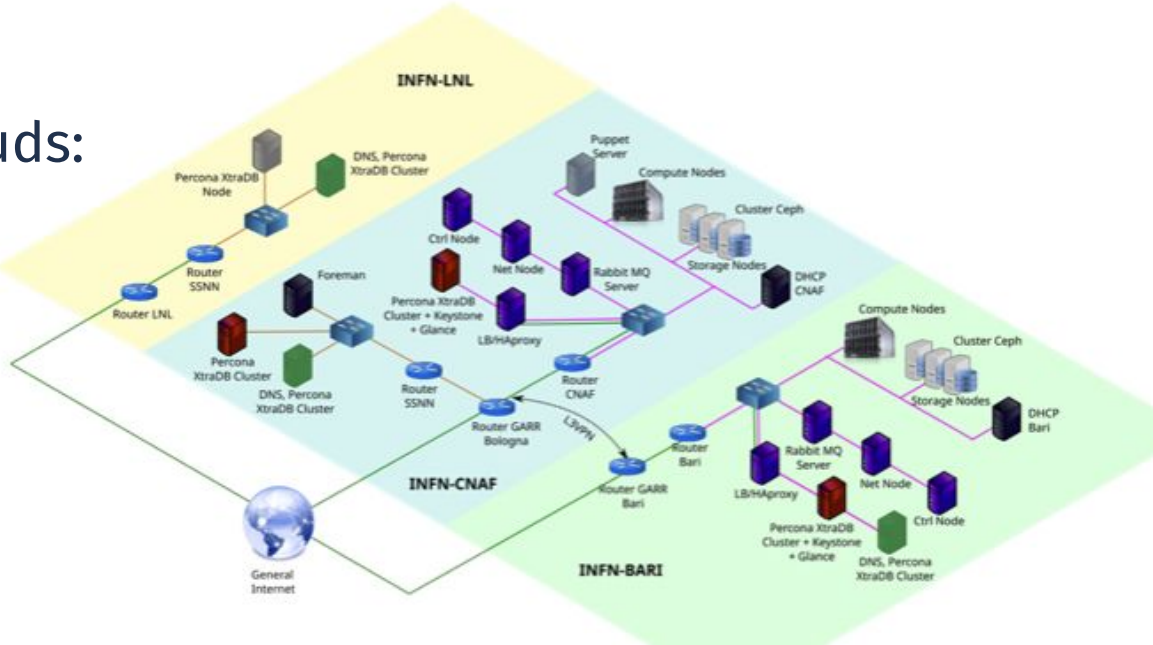

INDIGO PaaS

The INDIGO-DataCloud PaaS is rooted on:

- A distributed resource orchestration framework
- A standard—based federated solution for identity access management (INDIGO-IAM)

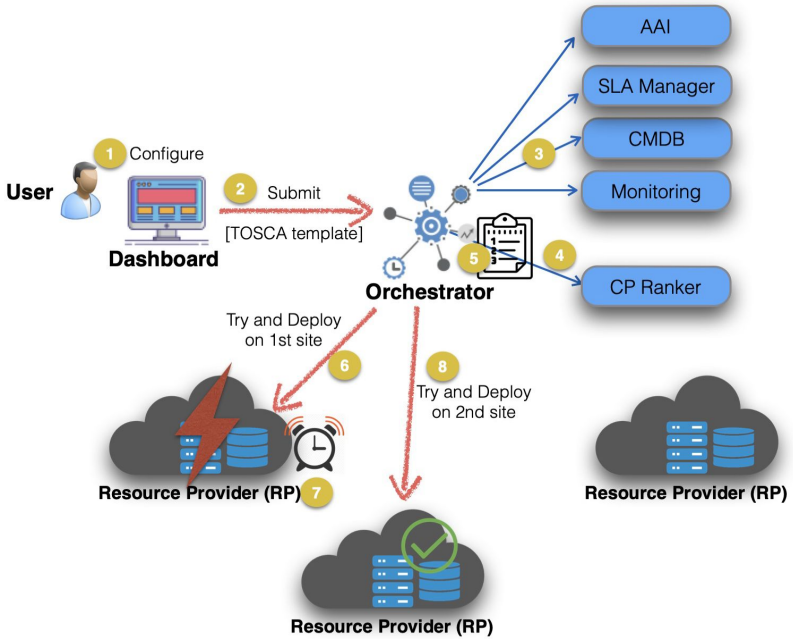

User interfaces first class citizens

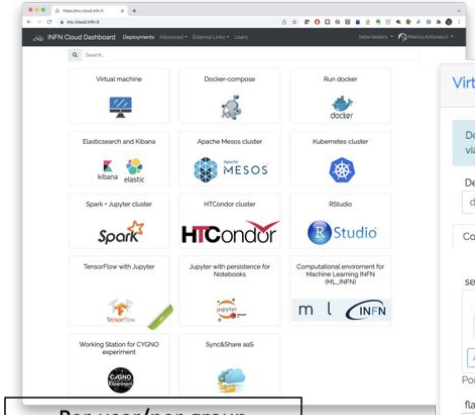

Per-user/per-group
View of the main dashboard

Basic/Advanced per-service
Configuration options

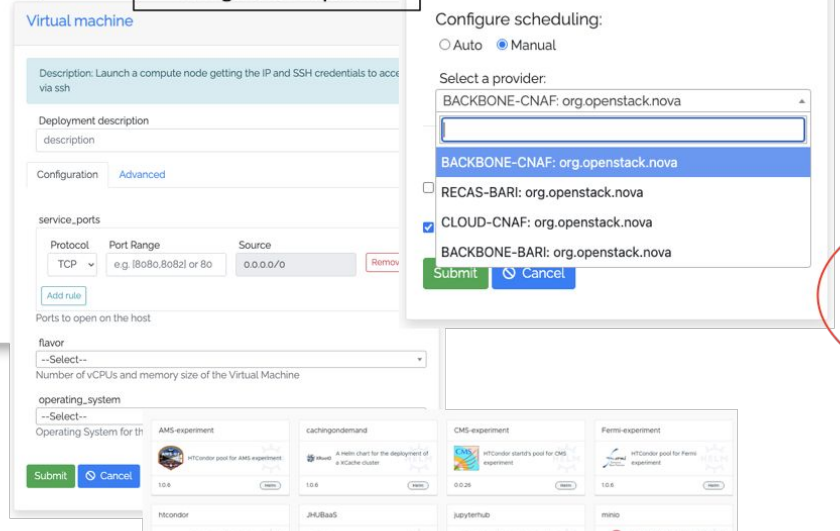

Kubeapp-level
Service dashboard

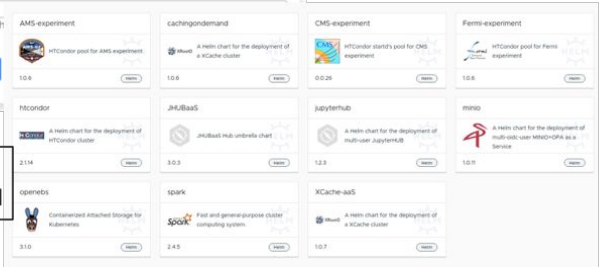

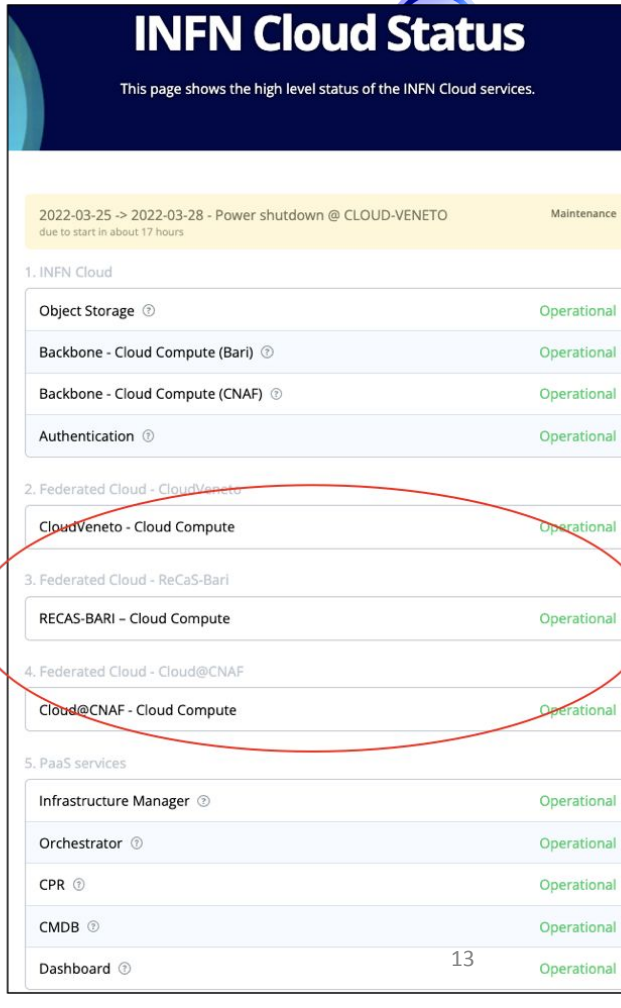

What is a user looking for?

Currently two main use cases have been raised from many fronts:

1. Interactive and graphical analysis environment
   a. colab-like, swan@cern like, in a word JupyterLab(++)
2. An easy to use storage system
   a. for data analysis (e.g. input, config reading)
   b. for sync&share (see dedicated talk on this later)
   c. Very often Posix-like access is an hard requirement

Different implementations
Slightly different needs

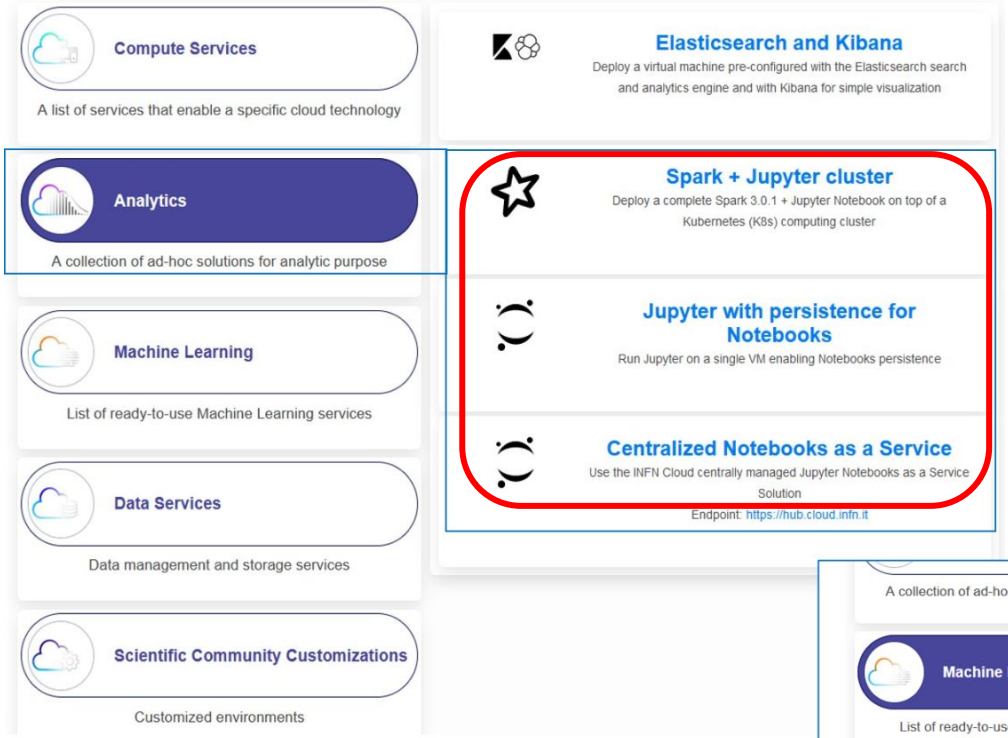

https://www.cloud.infn.it/service-catalogue

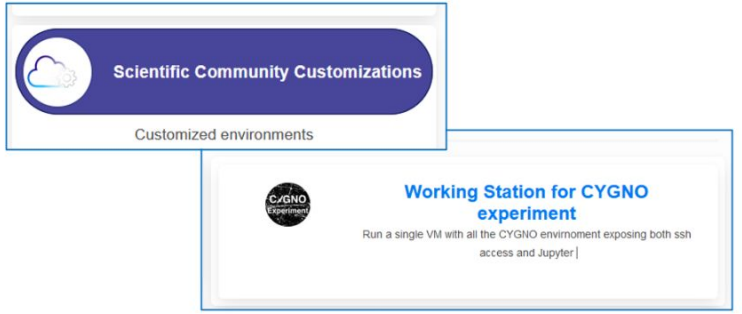

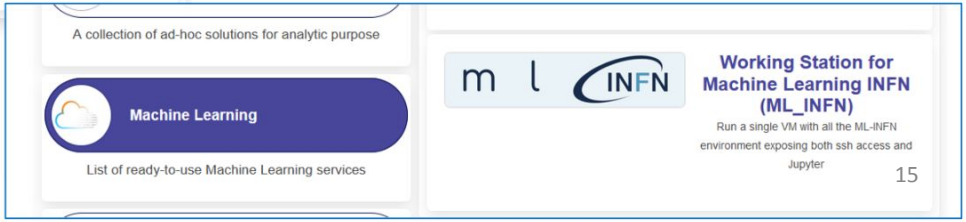

Deploy your service

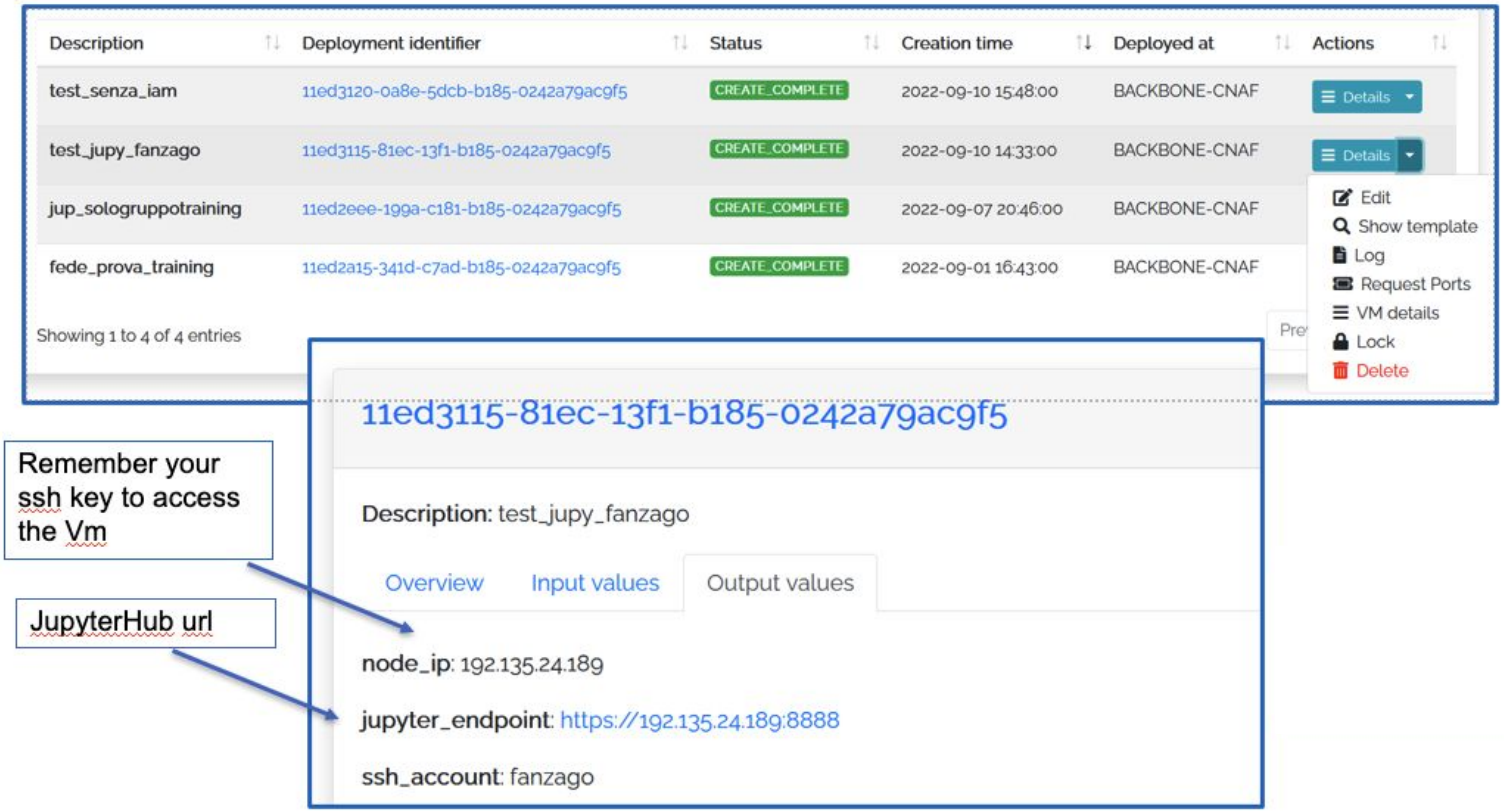

AuthN/Z: OIDC leads the integration

The entry point is a customized JupyterHUB instance integrated with INDIGO IAM:

1. Authenticate the user and forward the identity to the spawning jupyterlab instance
2. Authorize different flavor of instances based on the group membership
   a. E.g. different cpus/ram or GPU
   b. Administration rights

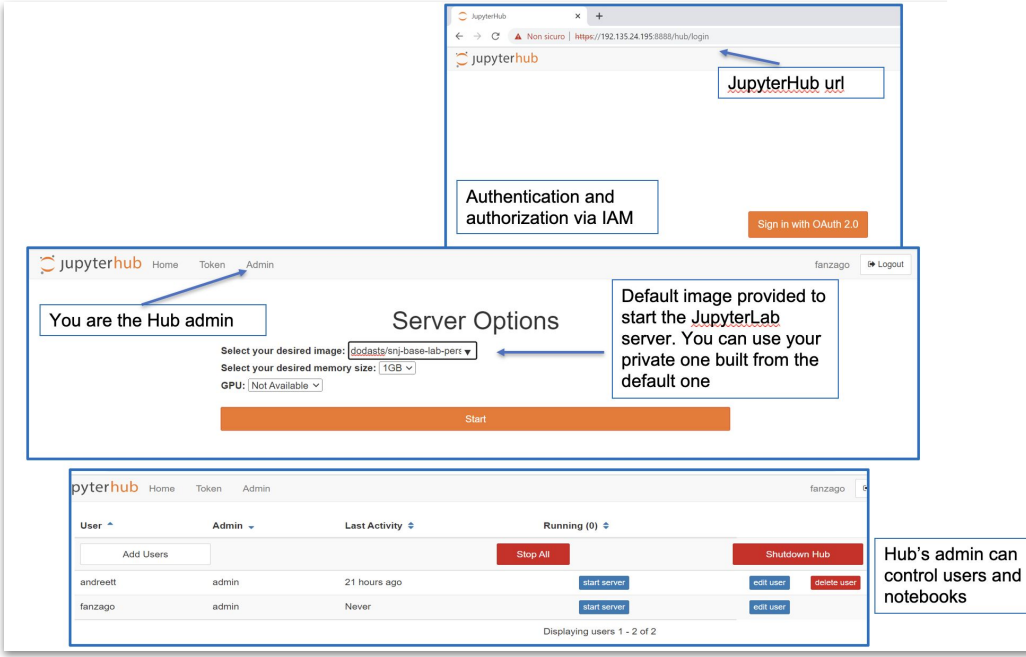

Then we needed a shareable and scalable storage…

INFN Cloud provides object storage based on Openstack Swift, as centrally managed service.

- The object storage is replicated in the two sites of backbone, Bari and CNAF
  - It guarantees the redundancy of data
- INFN Cloud is using the Minio-gateway software on Swift
  - Indigo-IAM OIDC authentication has been integrated and authorization policies are currently managed via OpenPolicyAgent
  - Ceph migration is under evaluation
    - Difficult to find good alternatives for Minio's WebUI tough
- Service URL https://minio.cloud.infn.it

Store JupyterLab persistent data

- All the services running notebooks mount user areas store on the object storage backend
  - E.g. they are visible in the file system as posix directories.
  - STS-wire: rclone mount+oidc shim library as been developed for the purpose

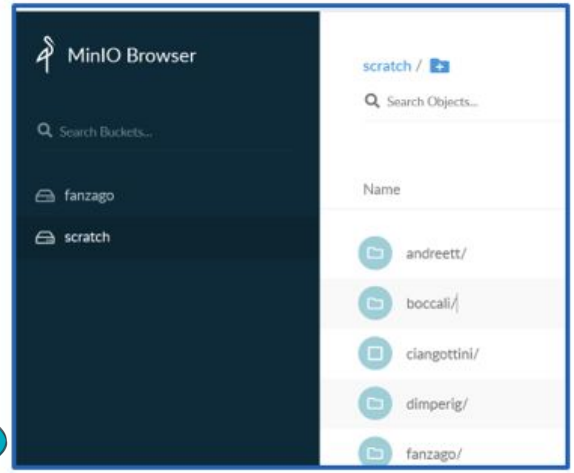

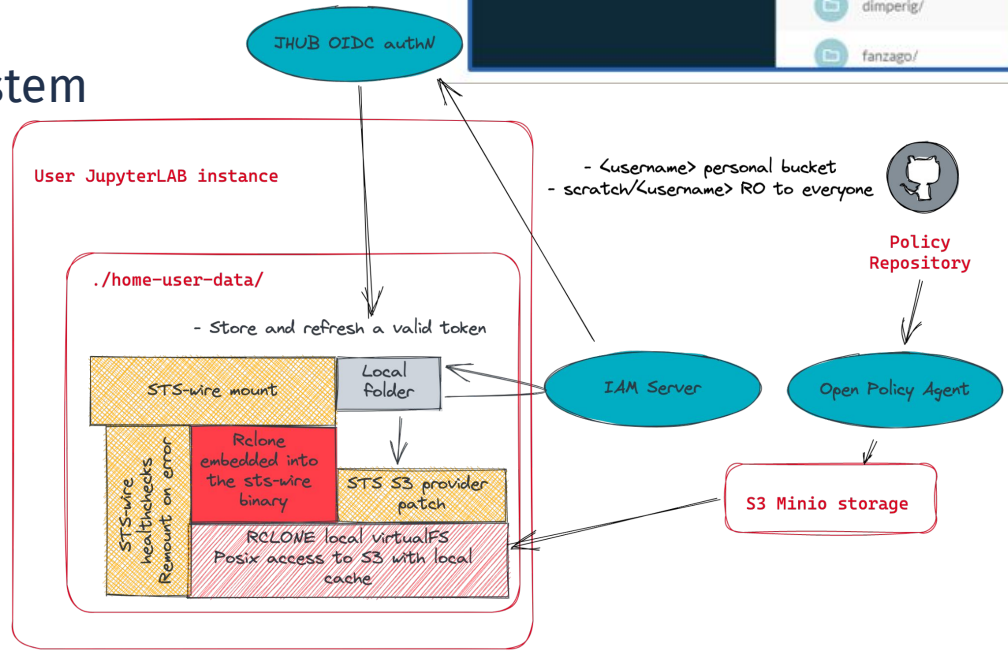

All in all, this is what you get

You login with your credentials, select your docker image and...

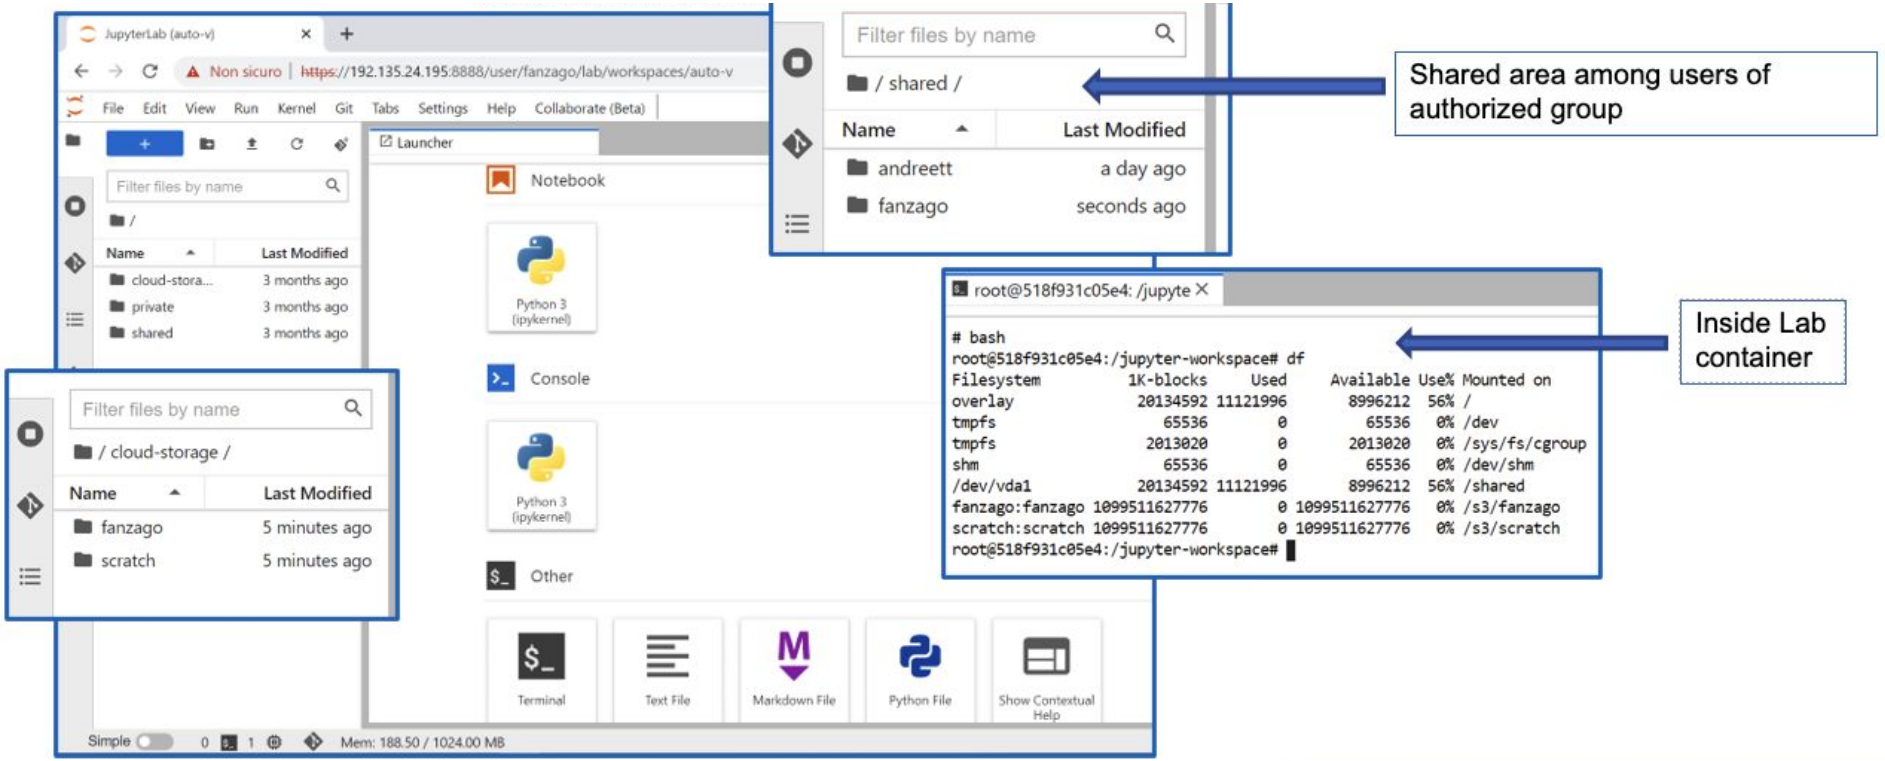

Not only on-demand...

Also a a SaaS managed service available for every INFN Cloud user.

Hosted on the INFN Cloud Backbone (HA/failover on the two sites)

Enabling a quick evaluation for new use cases

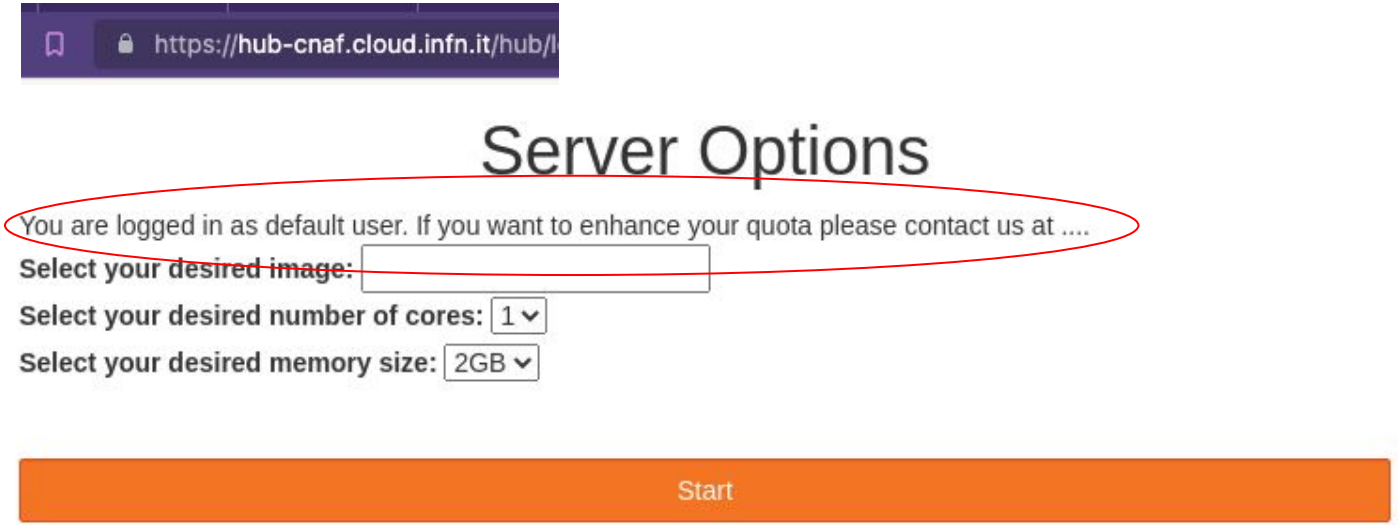

Conclusions

- INFN Cloud focus is on enabling users to perform their analysis work in a intuitive and easy to use fashion
- If you want to do data-science like experience with JupyterLab you can now:
  - Instantiate your own JupyterHub instance to share with your collaborators
  - Get access to the SaaS-level, for a ready to use environment
- At all the implementation you will be able to store and access data on the same object storage backend

More is yet to come:

- Swift to Ceph migration for S3-compatible object storage backend
- Managing data at different sites
  - data-lake approach, RUCIO-FTS under study
- Exploiting ML specific setups for the whole pipeline
  - A first functional setup already available under ML-INFN project activity
  - Preprocessing, training, inference…
    - putting it all together in a cloud-native way is the main challenge now