



**Third edition of the  
Machine Learning @ INFN (ML\_INFN)  
advanced level hackathon**

# **Open Science Cloud**

Daniele Spiga INFN-Pg

**Bari 21-24 Novembre 2022**



# Agenda

Introduction to Open Science Cloud... a disclaimer

- My personal view

OSC at INFN (My personal view)

- Few examples
- A focus on ML\_INFN toolkits @ INFN-Cloud

Summary

# Disclaimer

What is Open Science Cloud? a buzzword ?!?

- **EOSC:** [European Open Science Cloud](#)
- **SOSC:** [School on Open Science Cloud](#)
- **GOSC:** [Global Open Science Cloud](#)
- ...

I'd read it as the sum of two "concepts":

- **Open Science:** An approach to the scientific process that focuses on spreading knowledge as soon as it is available using digital and collaborative technology (made of Expert groups, publications, news and events)
- **Cloud** computing: as a mean to use digital and collaborative technology making use of hosted services, such as data storage, servers, networking as well high level services and software over the internet

# From a research e-infrastructure perspective (my view)

Allow researcher to exploit “free” and open services to manage workflow, build pipeline, data processing and analysis and, of course, to share/to reuse technical solutions

- Allow researchers to focus on science

## Technical drivers:

- to enable users to create and provision infrastructure deployments, automatically and repeatedly, with almost zero effort.
- To Implement the *Infrastructure as Code* paradigm based on declarative approach: allows to describe “What” instead of “How”
  - Let the underlying system to deal with technicalities
- To promote (and support) **container-based solutions**
- To grant data sharing among users/infrastructures

# ...and from user perspective: few pillars

**At first order I think end users should handle just few pillars**

- What the user should/might see out of all of the underlying system?

**Software management:** a central role is played by container. A standard unit of software suitable to create **user tailored environment**, (share and port everywhere).

- Docker is an open source platform for building, deploying, and managing containerized applications: a **handy application encapsulation**.
- a de facto standard to manage runtime environments, and we use dockers everywhere [see later]
- Tip [Docker store](#)

**Infrastructure management:** in principle user might chose to know “nothing” about infrastructure (SaaS model and above).

- If a researcher need/want to customize its infrastructure, the system (the Cloud) should offer handles... **through templates** [see later]



# INFN-Cloud

An **internal effort** at the INFN level in order to manage a (large) fraction of the INFN resources, in order to decouple user needs from the availability of local and dedicated hardware: this applies both to data and compute

## Aims at providing solutions for a wide range of user/community needs:

- Computing **Resources optimization**
- **Reuse** of solutions
- Support R&D: **design your computing model**
- A platform for **training**
- Ease (democratize) the access to the computing capacity:
  - Think to **access specialized hw such as: accelerators** (GPU, FPGA...)

**Few  
highlights**

# A bit more in concrete

Services ready for use. Just instantiate your own

|   |   |  |
|---|---|--|
| Virtual machine<br>                      | Docker-compose<br>                         | Run docker<br>  |
| Elasticsearch and Kibana<br>             | Kubernetes cluster<br>                     | Spark + Jupyter cluster<br>                                       |
| HTCondor cluster<br>                      | Jupyter with persistence for Notebooks<br> | Computational environment for Machine Learning INFN (ML_INFN)<br> |
| Working Station for CYGNO experiment<br> | Sync&Share aaS<br>                         |  |

**Example of a production ready use case [more later]**

## Two major highlights

A ecosystem providing a **platform, toolkits, support and experience** to **develop and prototype ad hoc solutions** based on open source and industry standard (even cloud-native) technologies

- co-design and develop **customized solution for specific needs**

There are additional features and functionalities [check here](#)

# So if you'd like to raise the bar

Examples...

**Automation:** Building your pipeline

- Exploiting cloud-native services to build a “event based” system/workflow

**Workflow management:** interactive processing

- Exploiting parallelism, Implement Interactive analysis workflow

**Specialized hardware:**

- “pioneered” by INFN-Cloud & ML\_INFN joint venture (GPU access)

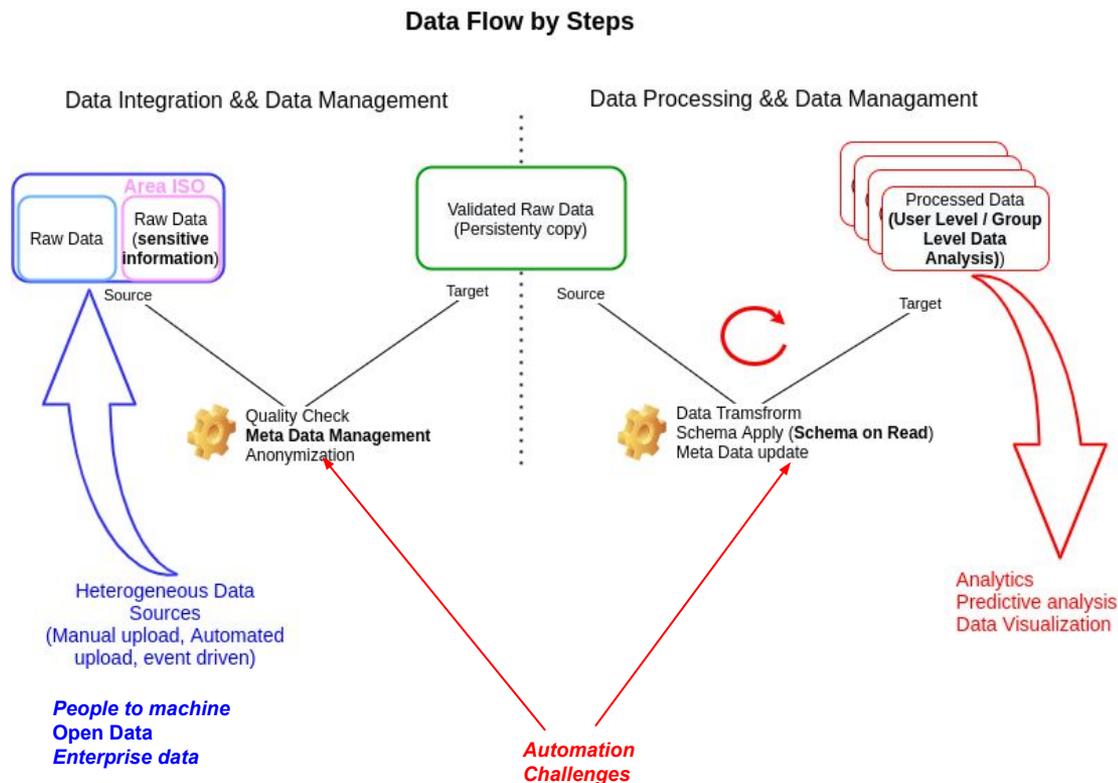
Several initiatives on going, two quite advanced, where this co-design is happening:

- CYGNO experiment, HERD Experiment (CSN2)

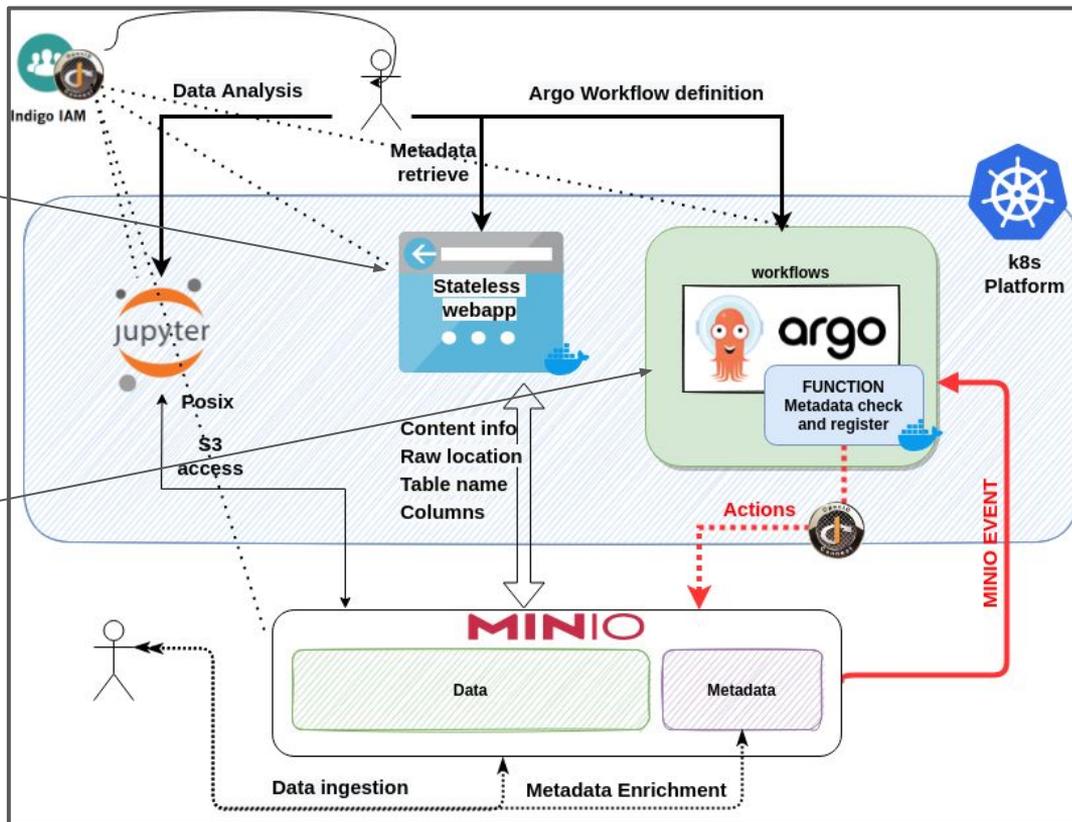
# Workflow automation (a real case at INFN)

## Some high level requirements:

- Structured and unstructured data archival
- Preserve data in its **original format**
- Enable automated data validation ( and organization )
- Enable automated data pre-processing, transformation...
- Easily find your data



# A cloud-native platform would look like



The only custom service (containerized)

**Argo Workflows** is an open source container-native workflow engine for orchestrating parallel jobs on Kubernetes.

Almost everything is implemented with **Industry Standard solutions**

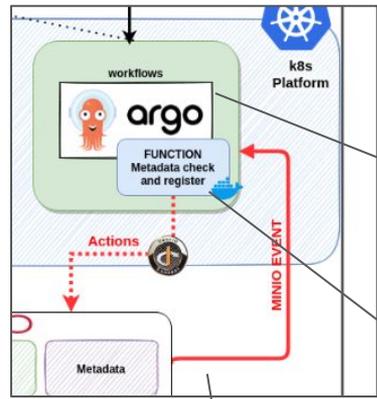
Container (docker) everywhere

Container Orchestration

- Self healing
- Automate
- Easy service deployment

Architecture of **The PLANET**  
[experiment@CSN5](mailto:experiment@CSN5)

# ... under the hood (declarative, container...)



Argo Sensor detect the event and trigger the validation

```
triggers:
- template:
  name: minio-workflow-trigger
  k8s:
    source:
      resource:
        apiVersion: argoproj.io/v1alpha1
        kind: Workflow
        metadata:
          generateName: artifact-workflow-2-
          namespace: argo-events
        spec:
          entrypoint: hook
          templates:
            - container:
                args:
                  - THIS_WILL_BE_REPLACED
                command:
                  hook
                env:
                  - name: ACCESSKEYID
                    value: admin-creds
                  - name: ENDPOINT
                    value: 'planet-store.cloud.cnaf.infn.it:9000'
                  - name: SECRETACCESSKEY
                    value: '223*sU#0!kss'
                image: 'dodasts/planet-demo-hook:v0'
                imagePullPolicy: Always
                name: hook
```



The EventSource: Each upload generate a Minio EVENT

```
spec:
  minio:
    example:
      endpoint: 'planet-store.cloud.cnaf.infn.it:9000'
      bucket:
        name: demo-raw
      accessKey:
        name: artifacts-minio
        key: accesskey
      secretKey:
        name: artifacts-minio
        key: secretkey
      events:
        - 's3:ObjectCreated:Put'
```



A customizable validation function is automatically executed:

**If metadata OK then:**

tell Minio to move data to the validated bucket

**else**

tell MinIO to move data to triage && notify

**fi**

Custom

Put here your code

# Managing workload: Transparent offloading

Distributing workload to parallelize data processing can be a complex task.

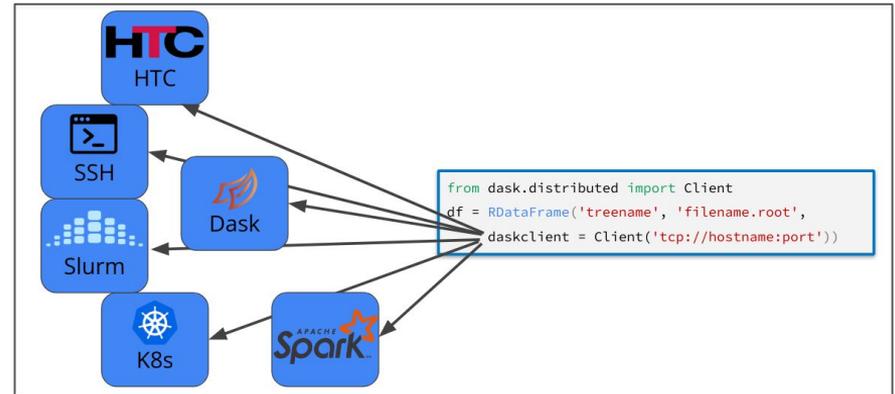
A dream would be to be able to access huge amount of computing capacity quickly and easily

- to process (huge amount of) data → **Interactive** or **Quasi interactive**
- reduce the time to insight: going interactive over huge amount of data

Simplifying: to being able to **scale up to a full workstation and transparently scale out to a cloud**

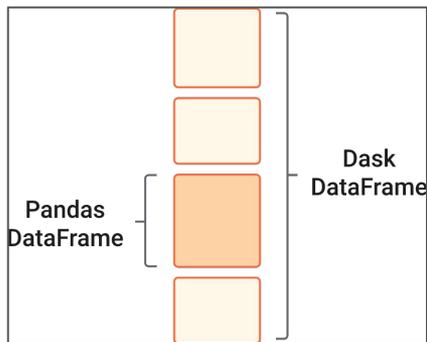
Several communities are exploiting the use high level frameworks capable of leverage distributed computing engine

- I.e. RDataFrame is getting traction @HEP



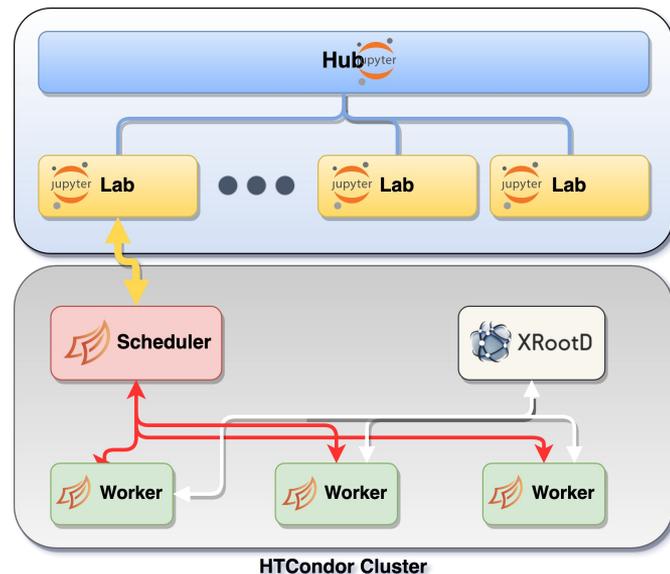
# An R&D @ INFN for HEP (CMS)

- **JupyterHub (JHub)** and **JupyterLab (JLab)** to manage the **user-facing part of the infrastructure**
  - **Comple abstraction**
- **DASK** to introduce the **scaling over a (highly) distributed system**
  - **Huge amount of resources, quickly and easily**
- **[XRootD is a bit HEP specific.. See it as a way to access any data anywhere]**



**Dask is not HEP.** It is a library that allow to scale the existing **Python and PyData ecosystem.**

- Looks and feels like the pandas API, but for parallel and distributed workflows.



R&D on interactive data analysis More details [here](#)

# Early results at CMS

Measured two different workflow distribution approaches

- Using VBS SSWW with a light lepton and an hadronic tau in final state
  - ported from legacy approach (nanoAOD-tools/plain PyROOT-based) to RDataFrame.
- Data processed about 2TB (Data + Monte Carlo)
- The comparison tests are performed
  - on the same nodes of the cluster
  - very same HTCondor infrastructure
  - A fair benchmark.



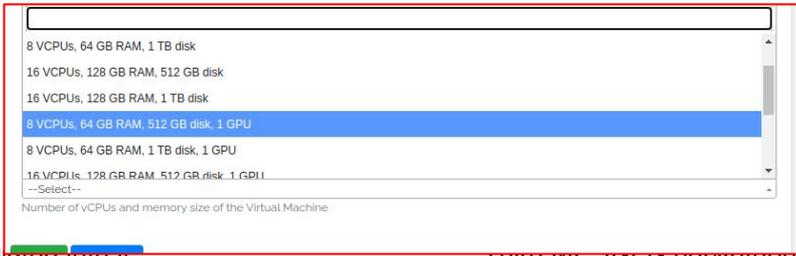
| Preselection                |           |           |
|-----------------------------|-----------|-----------|
|                             | Legacy    | RDF (O2)  |
| <b>Overall time</b>         | 3h 40min  | 25min     |
| <b>Overall rate</b>         | 693 Hz    | 7306 Hz   |
| <b>Event-loop rate</b>      | 721 Hz    | 8473 Hz   |
| <b>Overall network read</b> | 488 GB    | 371 GB    |
| <b>Average RSS per-node</b> | Ca. 13 GB | Ca. 17 GB |
| Postselection               |           |           |
|                             | Legacy    | RDF       |
| <b>Overall time</b>         | 0.25h     | 0.08h     |
| <b>Overall rate</b>         | 306 Hz    | 855 Hz    |
| <b>Event-loop rate</b>      | 412 Hz    | 1976 Hz   |
| <b>Overall network read</b> | 11 GB     | 10 GB     |
| <b>Average RSS per-node</b> | Ca. 1 GB  | Ca. 15 GB |



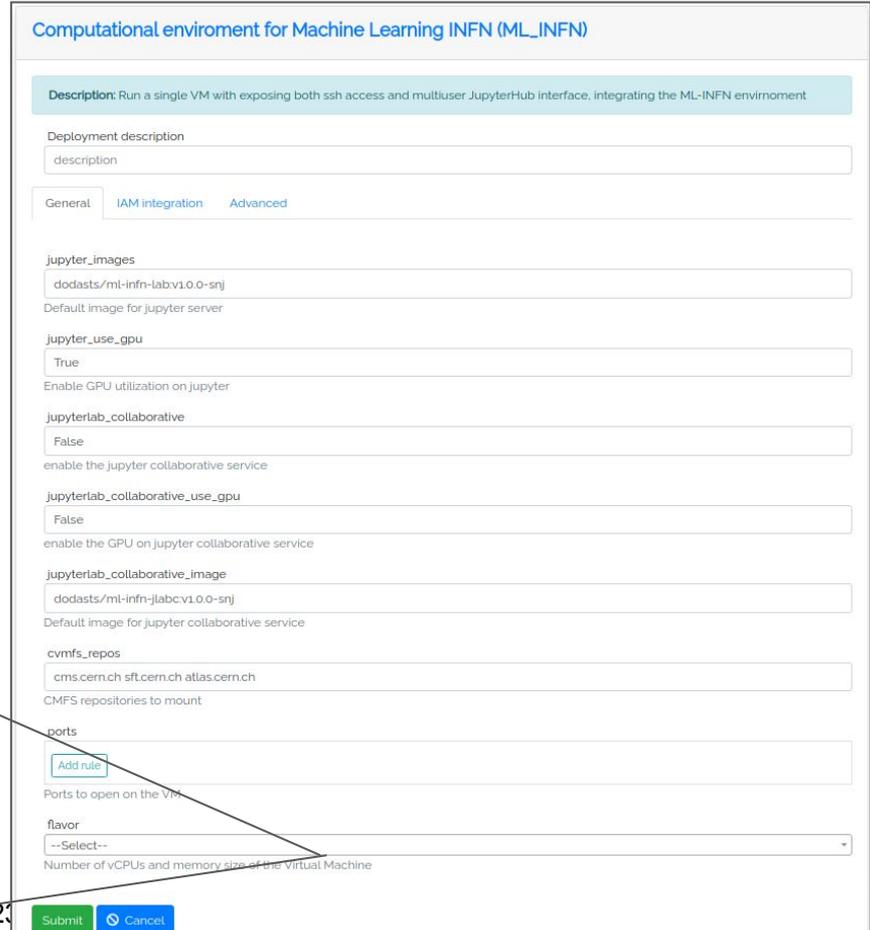
# ML-INFN @ INFN-Cloud

Build and promote the adoption of performant and tailored technological platforms because an effective platform for Machine Learning prototyping and developing might represent a technical challenge

- **access to hardware accelerators (GPUs)** and possibly a non-limiting **access to data** (i.e. training data).
- Handles to **create user tailored environment** is a key
- and finally groups need to **collaborate and share resources, data and code**.



8 VCPUs, 64 GB RAM, 1 TB disk  
 16 VCPUs, 128 GB RAM, 512 GB disk  
 16 VCPUs, 128 GB RAM, 1 TB disk  
**8 VCPUs, 64 GB RAM, 512 GB disk, 1 GPU**  
 8 VCPUs, 64 GB RAM, 1 TB disk, 1 GPU  
 16 VCPUs, 128 GB RAM, 512 GB disk, 1 GPU  
 --Select--  
 Number of vCPUs and memory size of the Virtual Machine



Computational environment for Machine Learning INFN (ML\_INFN)

**Description:** Run a single VM with exposing both ssh access and multiuser JupyterHub interface, integrating the ML-INFN environment

Deployment description  
 description

General | IAM integration | Advanced

jupyter\_images  
 dodasts/ml-infn-lab.v1.0.0-snj  
 Default image for jupyter server

jupyter\_use\_gpu  
 True  
 Enable GPU utilization on jupyter

jupyterlab\_collaborative  
 False  
 enable the jupyter collaborative service

jupyterlab\_collaborative\_use\_gpu  
 False  
 enable the GPU on jupyter collaborative service

jupyterlab\_collaborative\_image  
 dodasts/ml-infn-jlab.v1.0.0-snj  
 Default image for jupyter collaborative service

cvmfs\_repos  
 cms.cern.ch sft.cern.ch atlas.cern.ch  
 CMFS repositories to mount

ports  
 Add rule

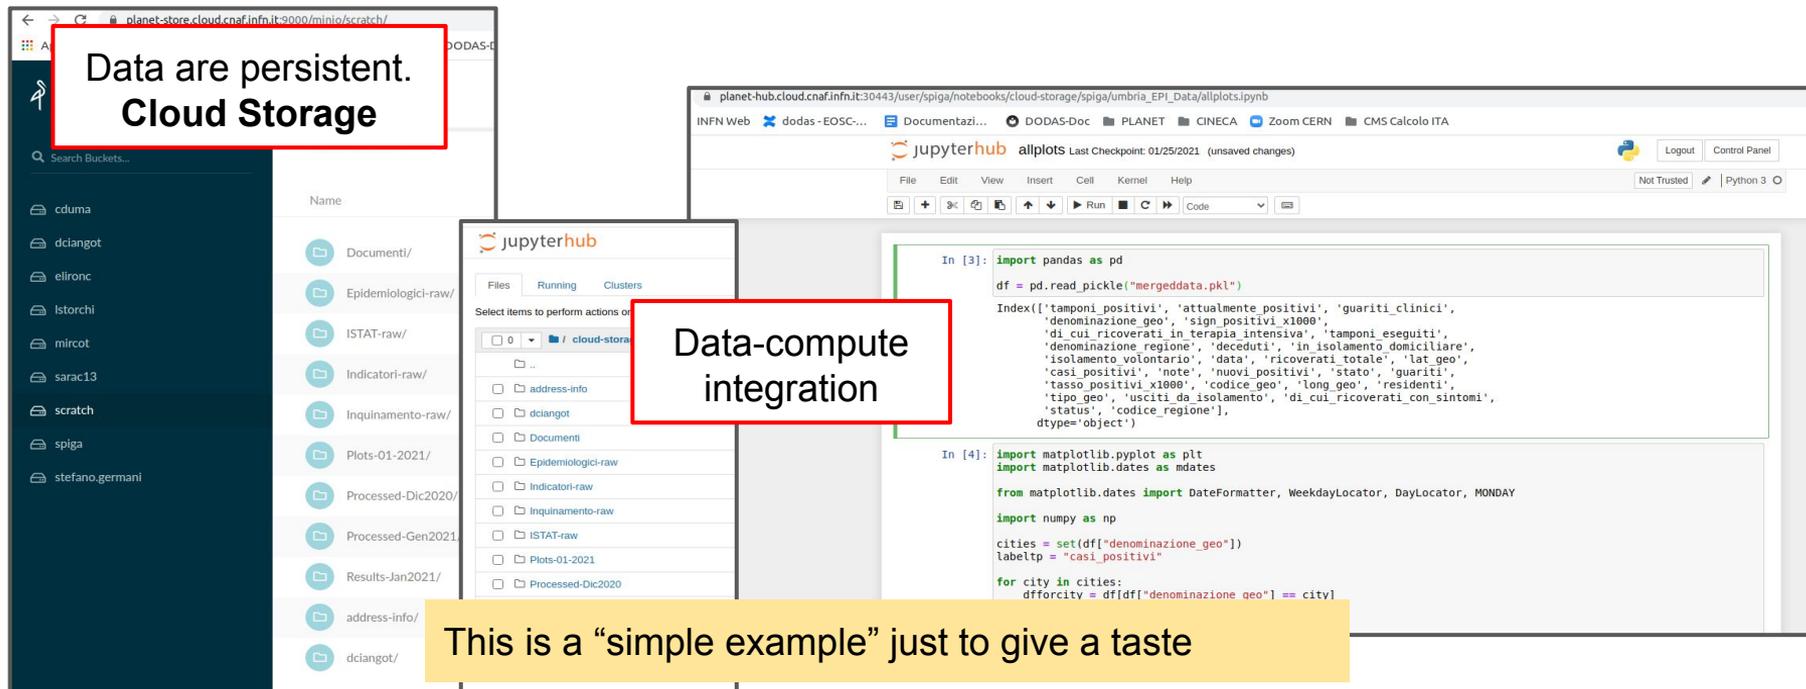
Ports to open on the VM

flavor  
 --Select--  
 Number of vCPUs and memory size of the Virtual Machine

Submit Cancel

# Data and compute: connecting (a few) dots

The ultimate goal is to seamlessly integrating data and compute (**The INFN DataLake**).



**Data are persistent. Cloud Storage**

**Data-compute integration**

```

In [3]: import pandas as pd
df = pd.read_pickle("mergeddata.pkl")
Index(['tamponi positivi', 'attualmente positivi', 'guariti clinici',
       'denominazione_geo', 'sign_positivi_x1000',
       'di_cui_ricoverati_in_terapia_intensiva', 'tamponi eseguiti',
       'denominazione regione', 'deceduti', 'in_isolamento_domiciliare',
       'isolamento_volontario', 'data', 'ricoverati_totale', 'lat_geo',
       'casi_positivi', 'note', 'nuovi_positivi', 'stato', 'guariti',
       'tasso_positivi_x1000', 'codice_geo', 'long_geo', 'residenti',
       'tipo_geo', 'usciti_da_isolamento', 'di_cui_ricoverati_con_sintomi',
       'status', 'codice_regione'],
      dtype='object')

In [4]: import matplotlib.pyplot as plt
import matplotlib.dates as mdates

from matplotlib.dates import DateFormatter, WeekdayLocator, DayLocator, MONDAY
import numpy as np

cities = set(df["denominazione_geo"])
labeltp = "casi_positivi"

for city in cities:
    dfforcity = df[df["denominazione_geo"] == city]
  
```

This is a "simple example" just to give a taste

# Summary

Open Science Cloud: presented a personal point of view and INFN perspectives

- At INFN we are build and growing an ecosystem for these technical matters

An overview of tools and solutions in the scope of the INFN-Cloud Portfolio

- Where to start and where to possibly contribute (idea and solutions)

Few R&D examples made with the aim to stimulate questions, curiosity and...  
possibly synergies

Mentioned R&D developments are there thanks to the work of many people

- Credits: INFN-Cloud Team; CMS Italy computing team; PLANET Team; ML\_INFN Team