V irtual **R** esearch **E** nvironment

Towards a comprehensive analysis platform

Elena Gazzarrini, Enrique Garcia, Domenic Gosein

European Union's Horizon 2020 programme Grant Agreement 824064 and 101017536

Agenda

Motivation

Science use case

Components

Demo

Motivation

Data volumes growing not only at LHC

The LHC at CERN was the first large scientific experiment to generate and manage multi PBs of data per year.

Technologies to manage and process data initially developed at CERN are being adopted by other collaborations, as new generation of detectors, antennas and telescopes are producing and processing large data volumes as well.

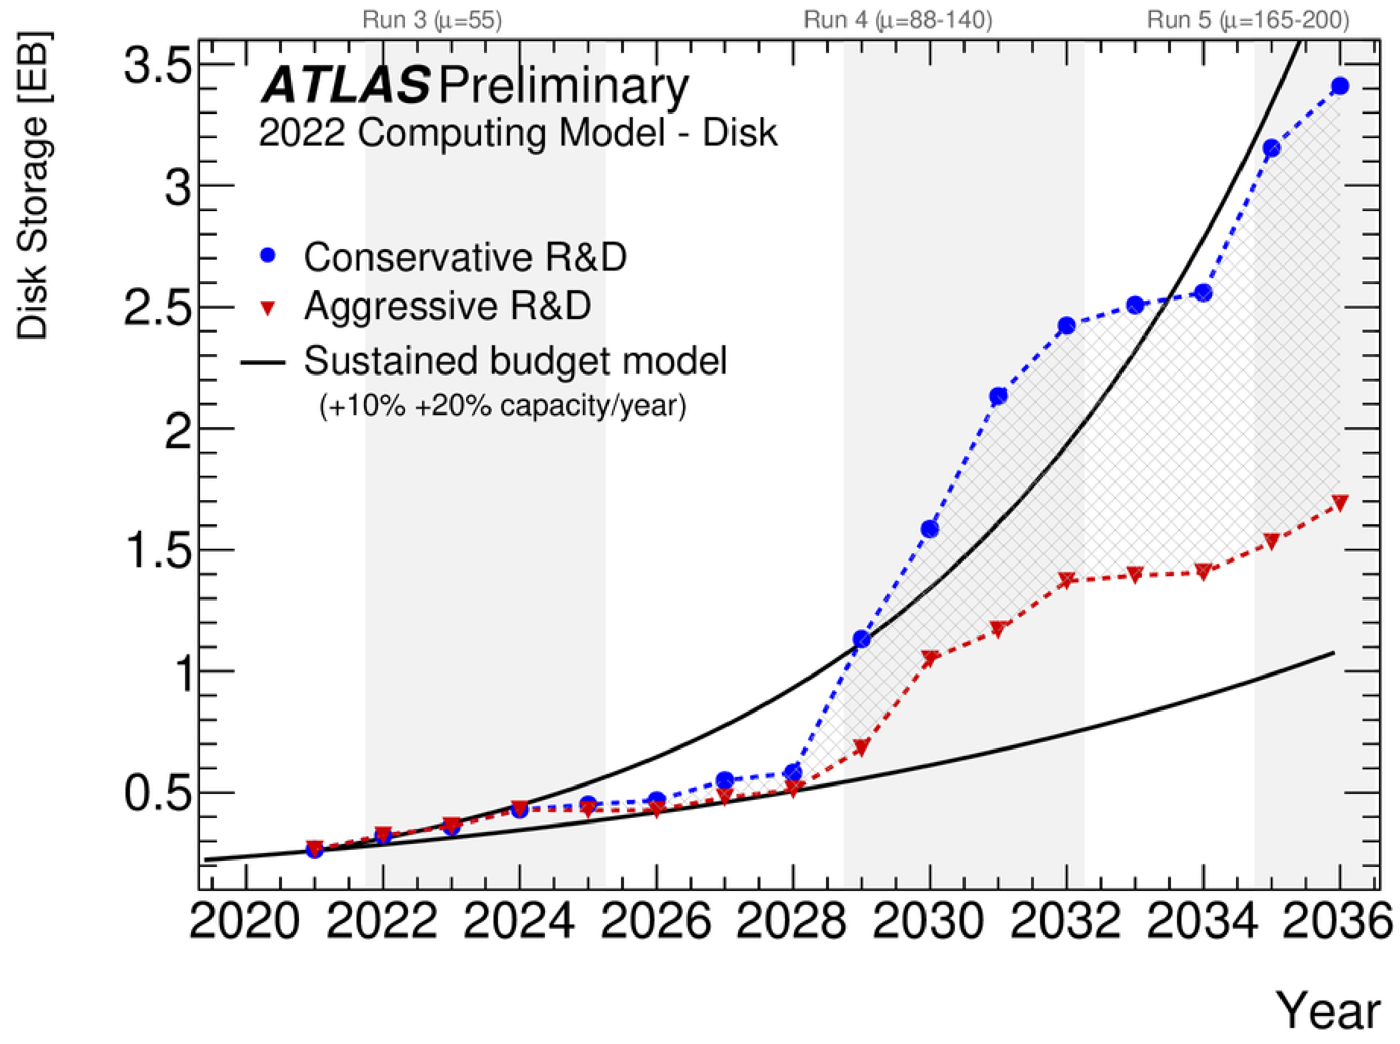

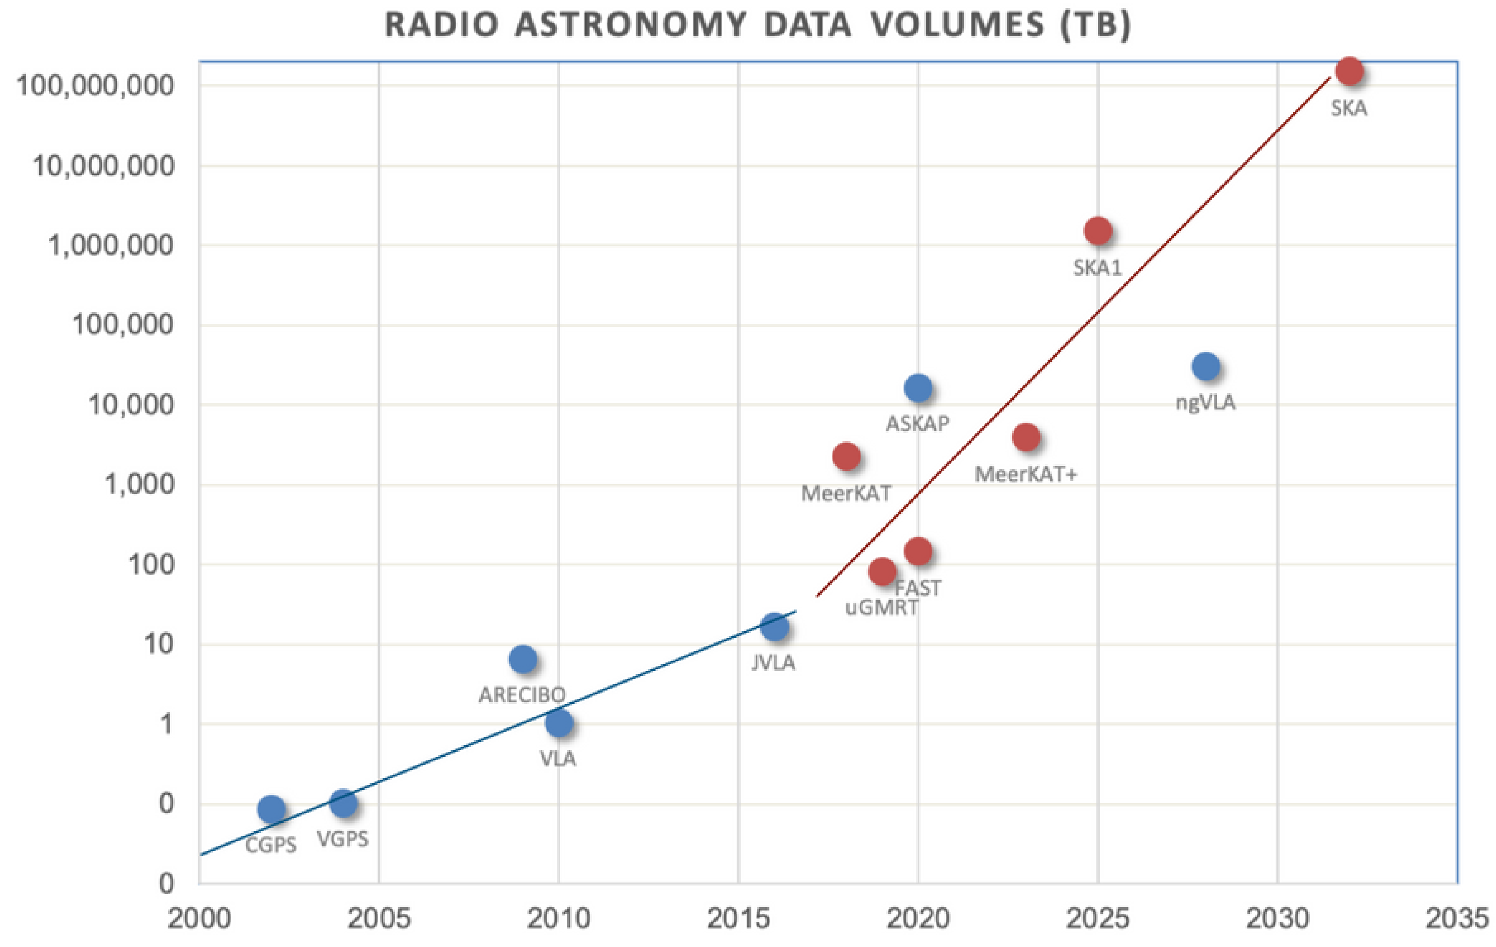

Taylor R. et al. Big Data Research Infrastructure Collaboration Toward the SKA (BRICSKA). doi: 10.1590/0001-3765202120201027. PMID: 34076205.

See plenary talk by Rosie Bolton

The challenge

A common infrastructure across Research Infrastructures would foster:

- economy of scale
- collaboration across domains
- scientific reuse
- sustainability

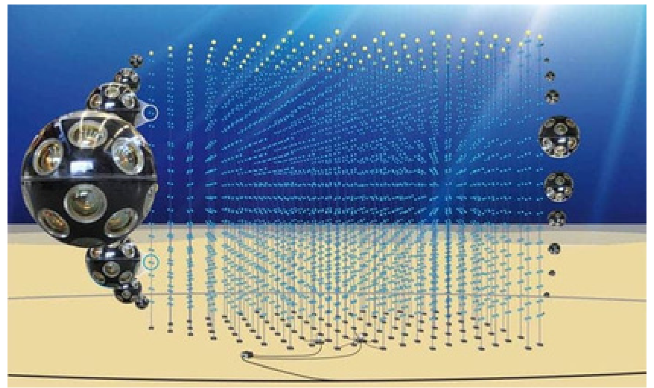
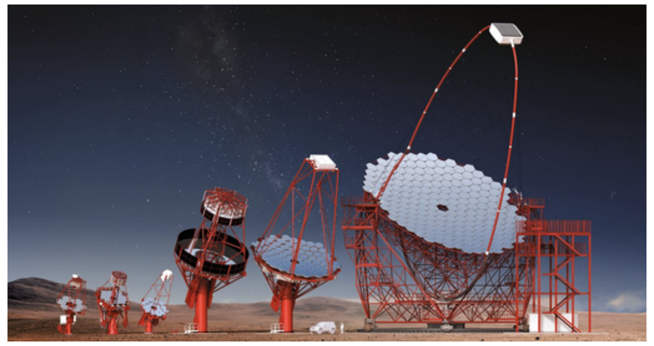
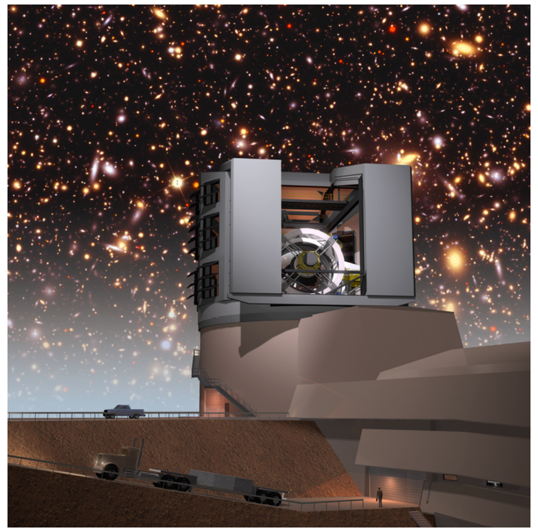


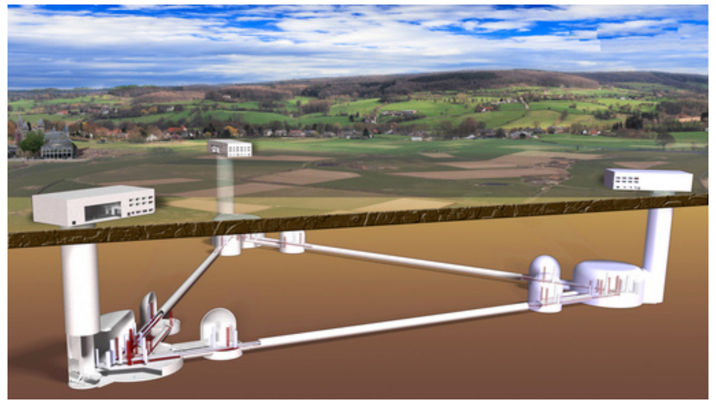

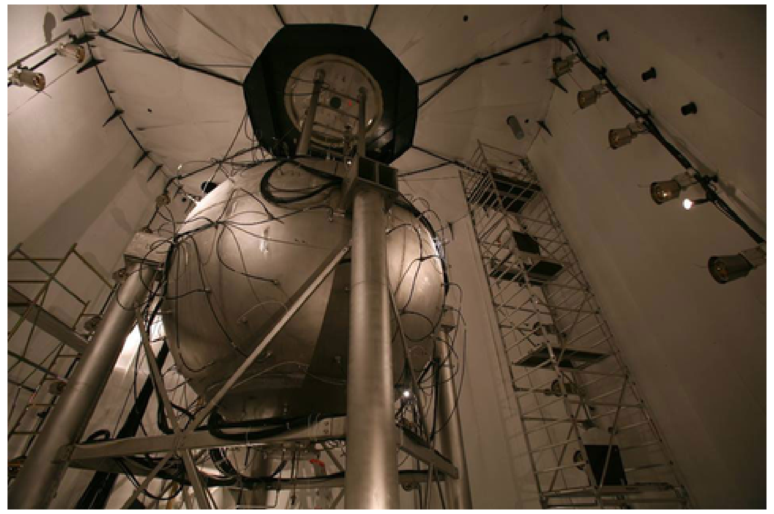


Have you failed to reproduce a result?

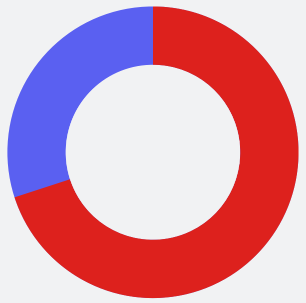

70%

researchers tried and failed to reproduce others' results

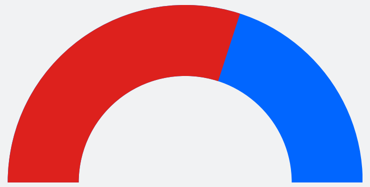

> 50%

researchers failed to reproduce own results

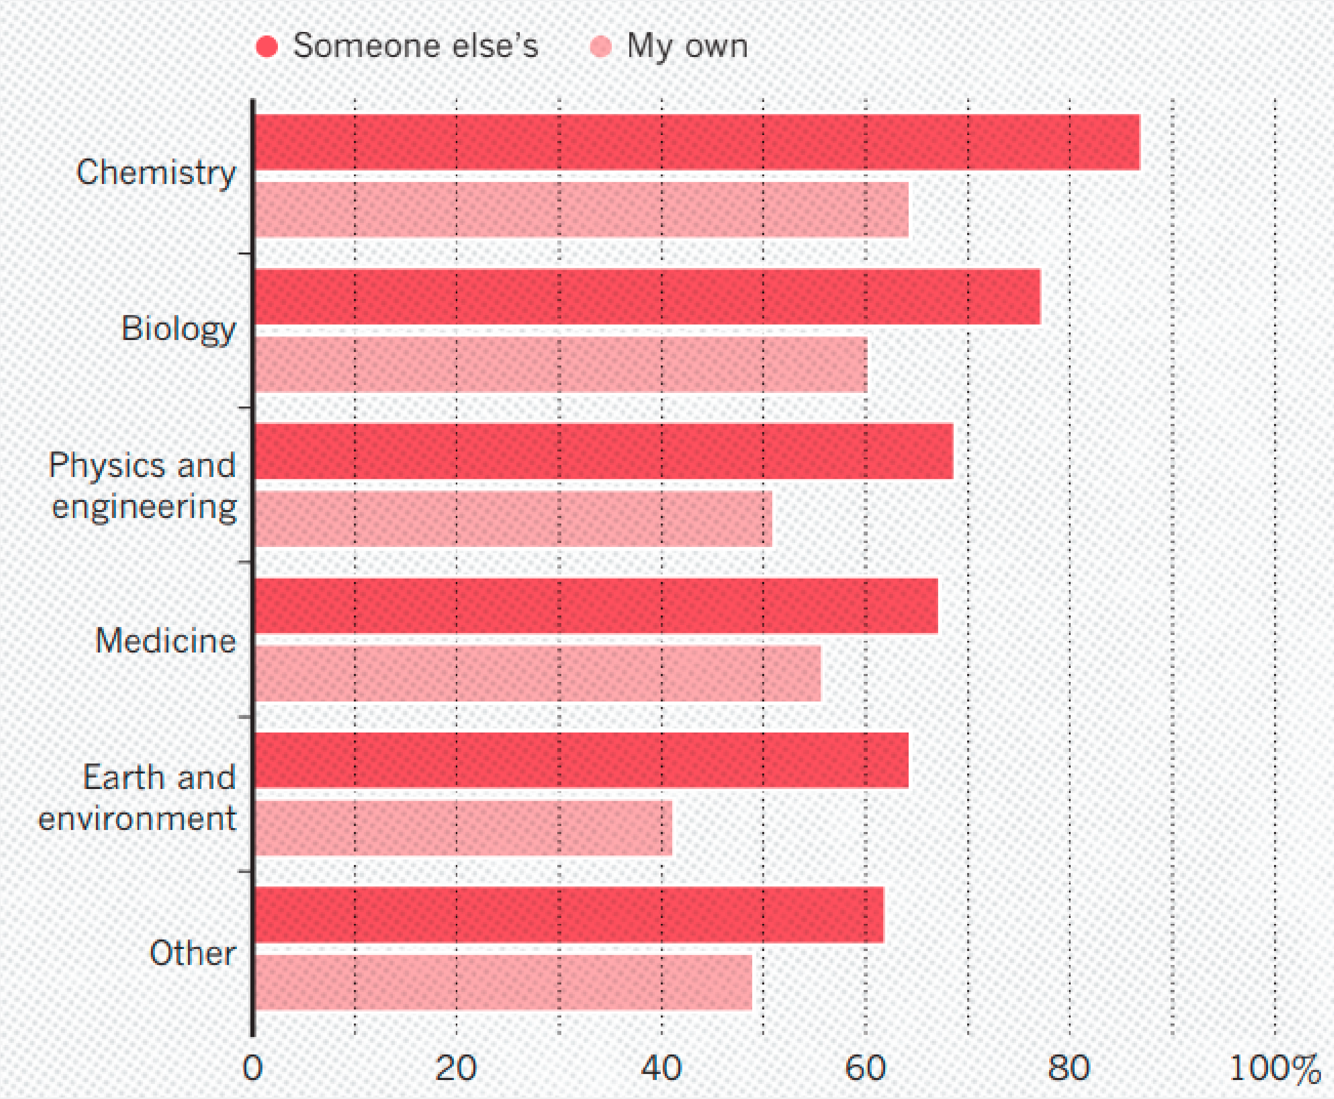

Baker, M. 1,500 scientists lift the lid on reproducibility. NatBure 533, 452–454 (2016).

EU collaborations

EU-funded projects promote cross-fertilisation across Research Infrastructures and scientific domains to find common, consistent and useful solutions to challenges of

- Federated Data Management and Transfer Services
- Distributed Data Processing
- Software Sustainability
- Analysis Preservation and Reusability

… all in one common Analysis Platform!

See yesterday's talk by G. Lamanna

CERN EOSC Future ESCAPE European Science Cluster of Astronomy & Particle physics ESFRI research infrastructures European Commission

The Virtual Research Environment

The VRE is an **open source** analysis platform where researchers have access to all the digital content needed to **develop, share and reproduce an end-to-end scientific result** in compliance with **FAIR** (findable, accessible, interoperable, reproducible) principles.

- Researchers
- Community Software & Analysis development
- Data & Workflow Management
- IT administrators

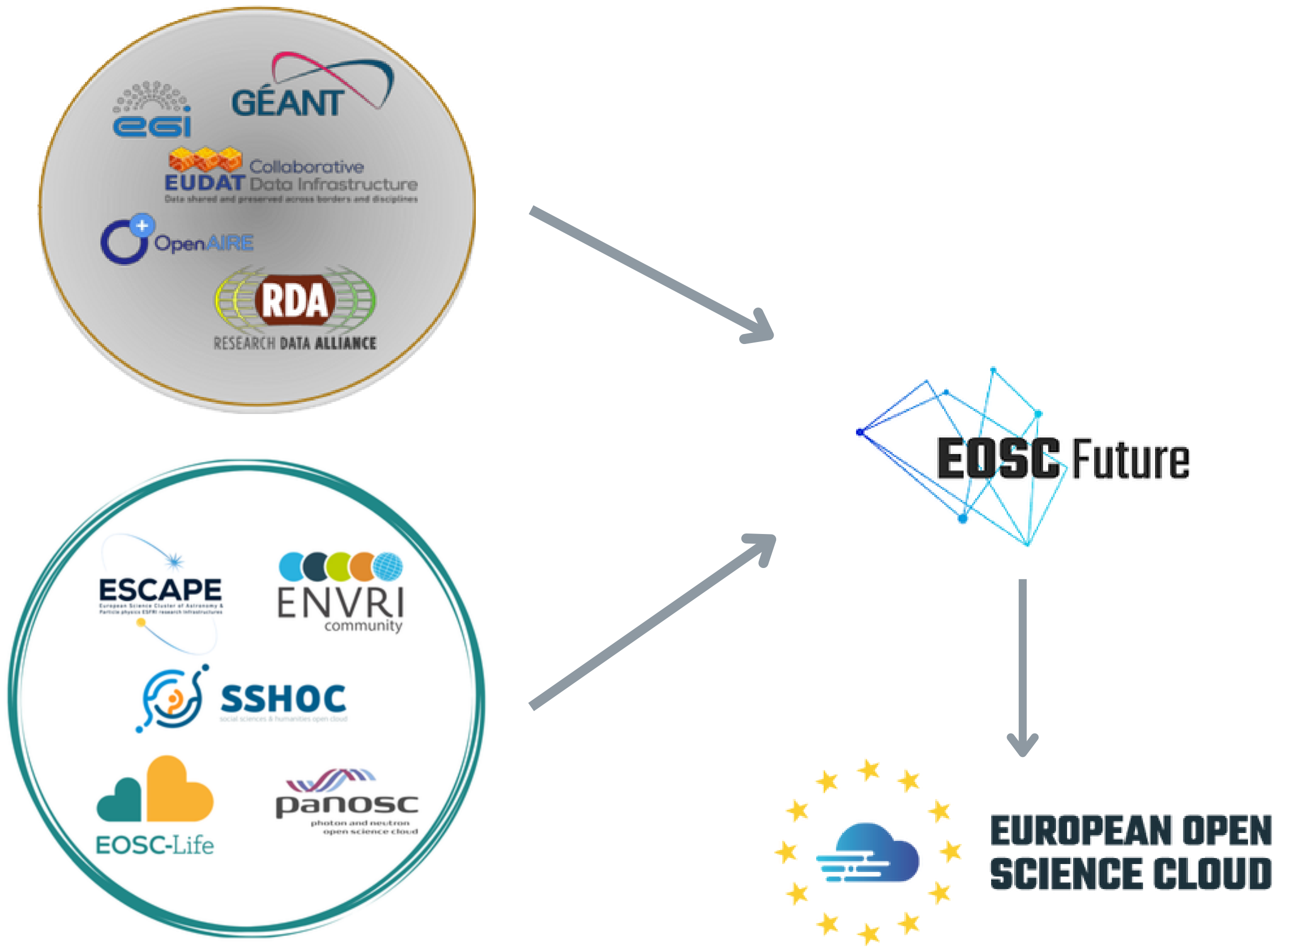

Timeline

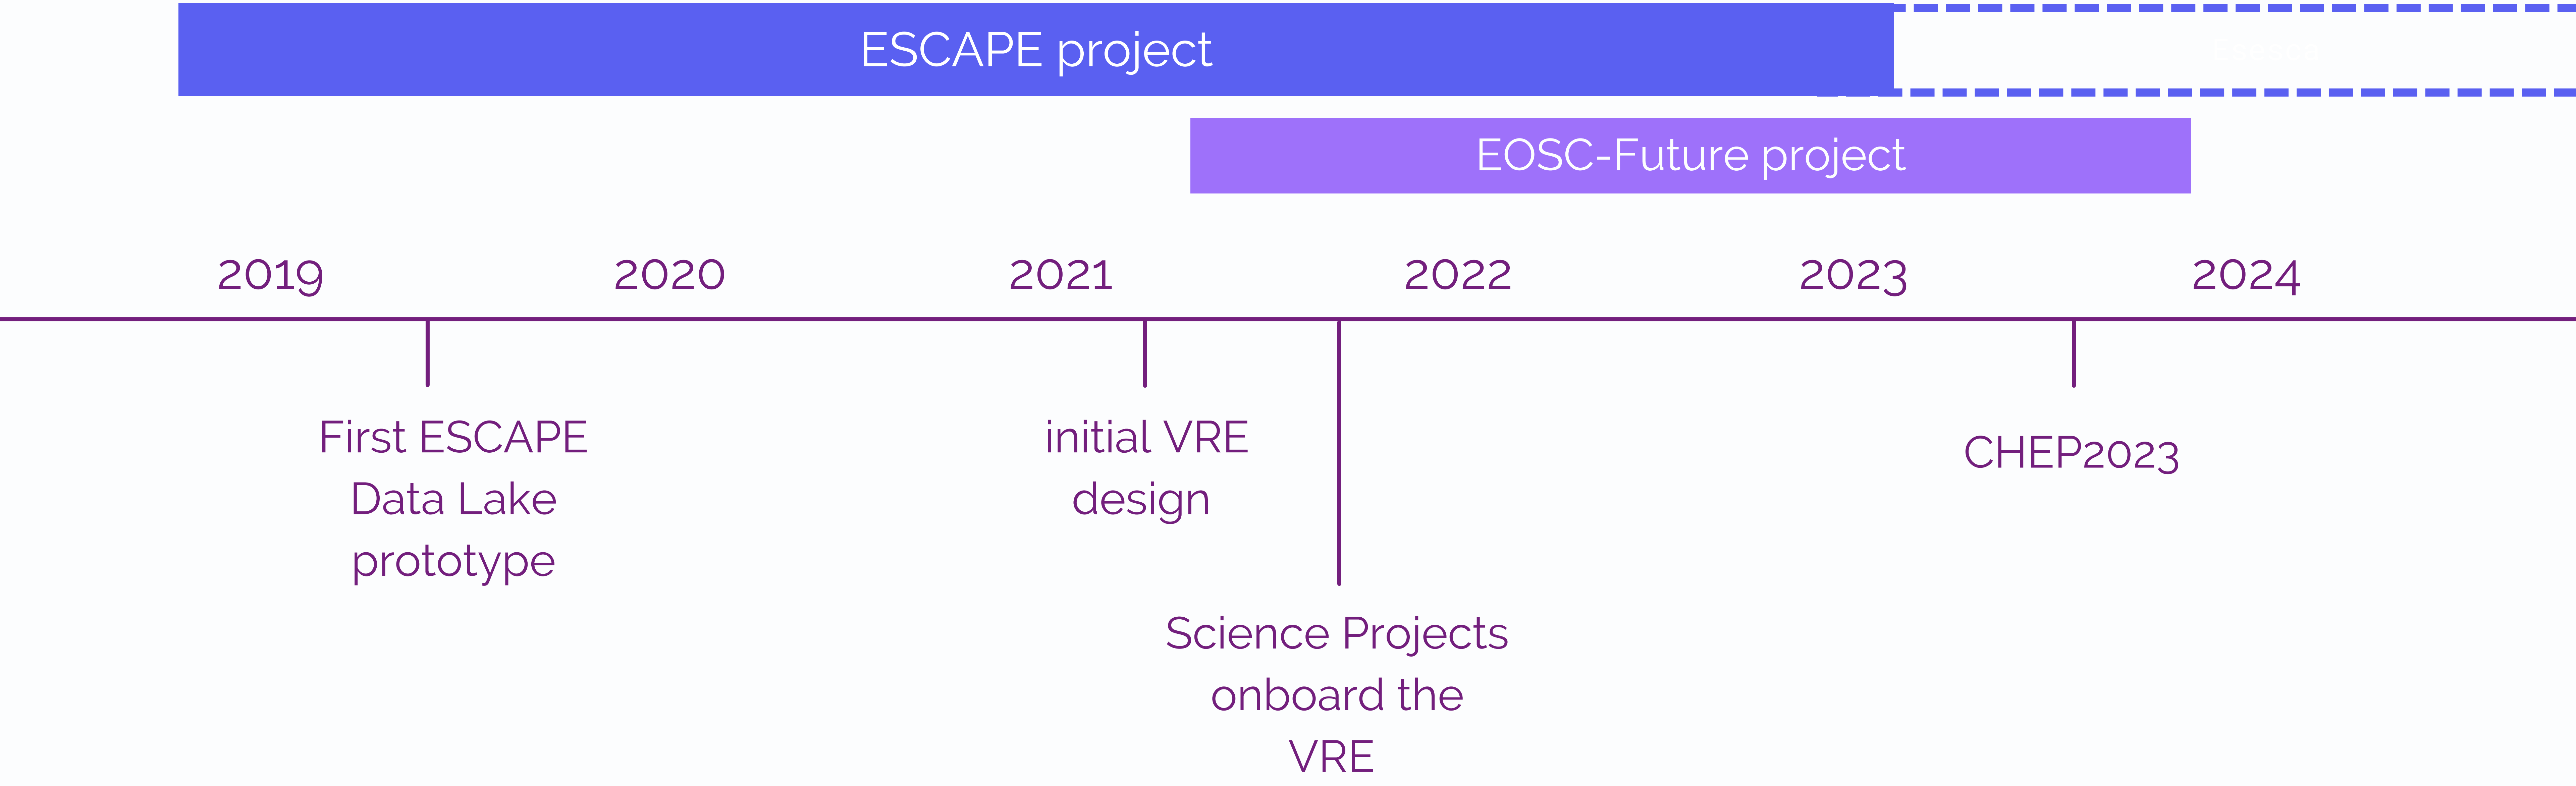

ESCAPE project
EOSC-Future project
2019
2020
2021
2022
2023
2024
First ESCAPE Data Lake prototype
initial VRE design
Science Projects onboard the VRE
CHEP2023

Science Use Case

Analysis workflow

reduce data formats for analysis

with preserved, versioned software on computing resources

combine resultls, compare across workflows and share

Data collection

Storage

Data processing

Analysis

Results interpretation & preservation

Experimental data generated at site

in data centres

Context: EOSC-Future

EOSC-Future Science Projects demonstrate

- multi-domain science integration across the **ESCAPE** project
- unification of services under **one Proof of Concept (PoC) analysis platform,** the **VRE**
- **interdisciplinary open science** example from bottom-up effort as a science driver for other communities

- Neutrino Observations
- High Energy Astronomy
- Low Frequency Telescopes
- Particle Physics
- Gravitational Waves

→

DARK MATTER

EXTREME UNIVERSE

→

Virtual Research Environment

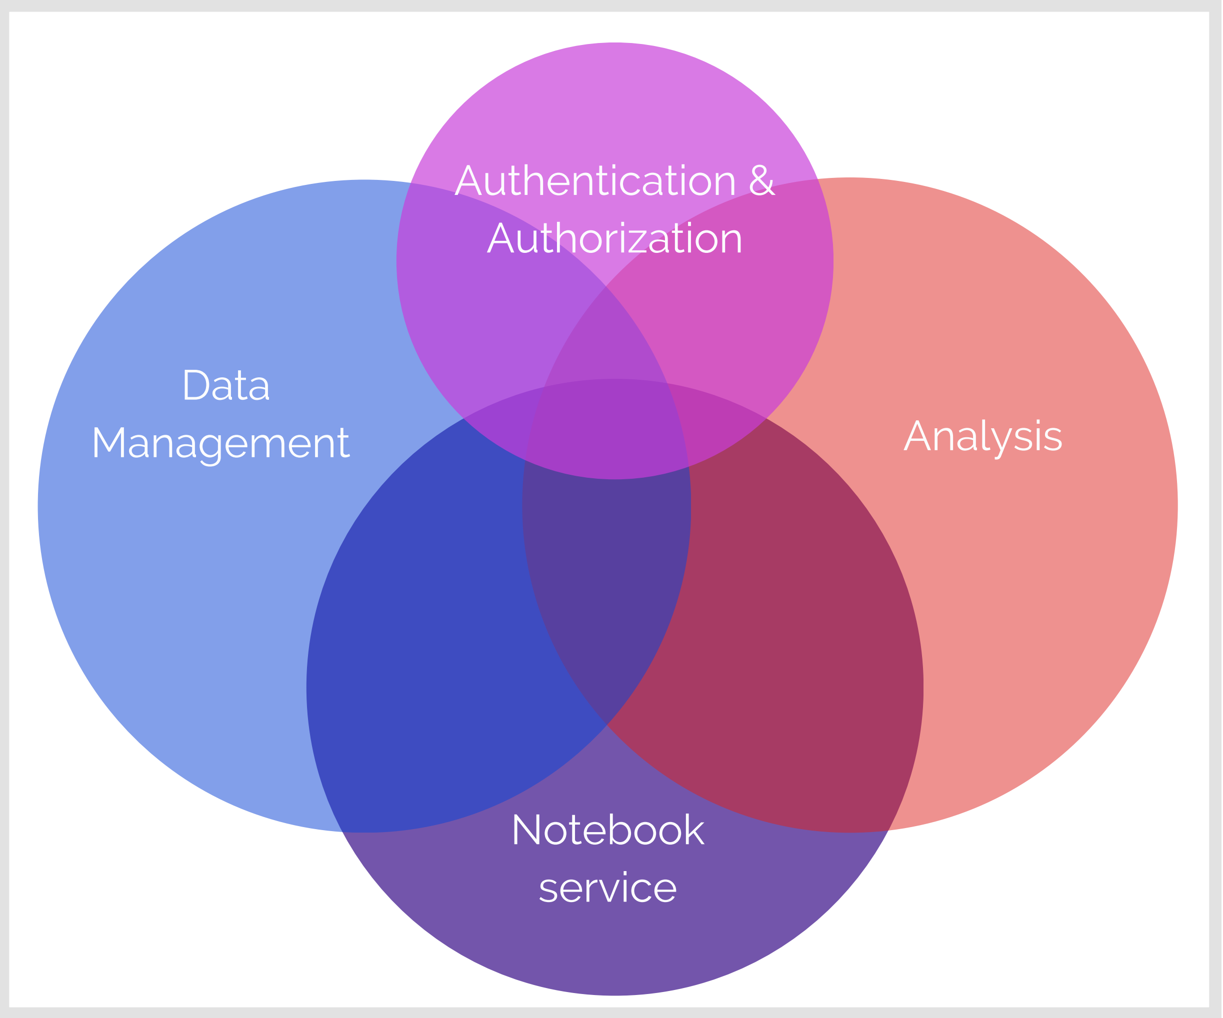

Dark Matter searches

Both HEP and astrophysics and researching limits of Dark matter, and preserving analysis workflows would enable to output **combined plots**

See yesterday's talk by J. Little

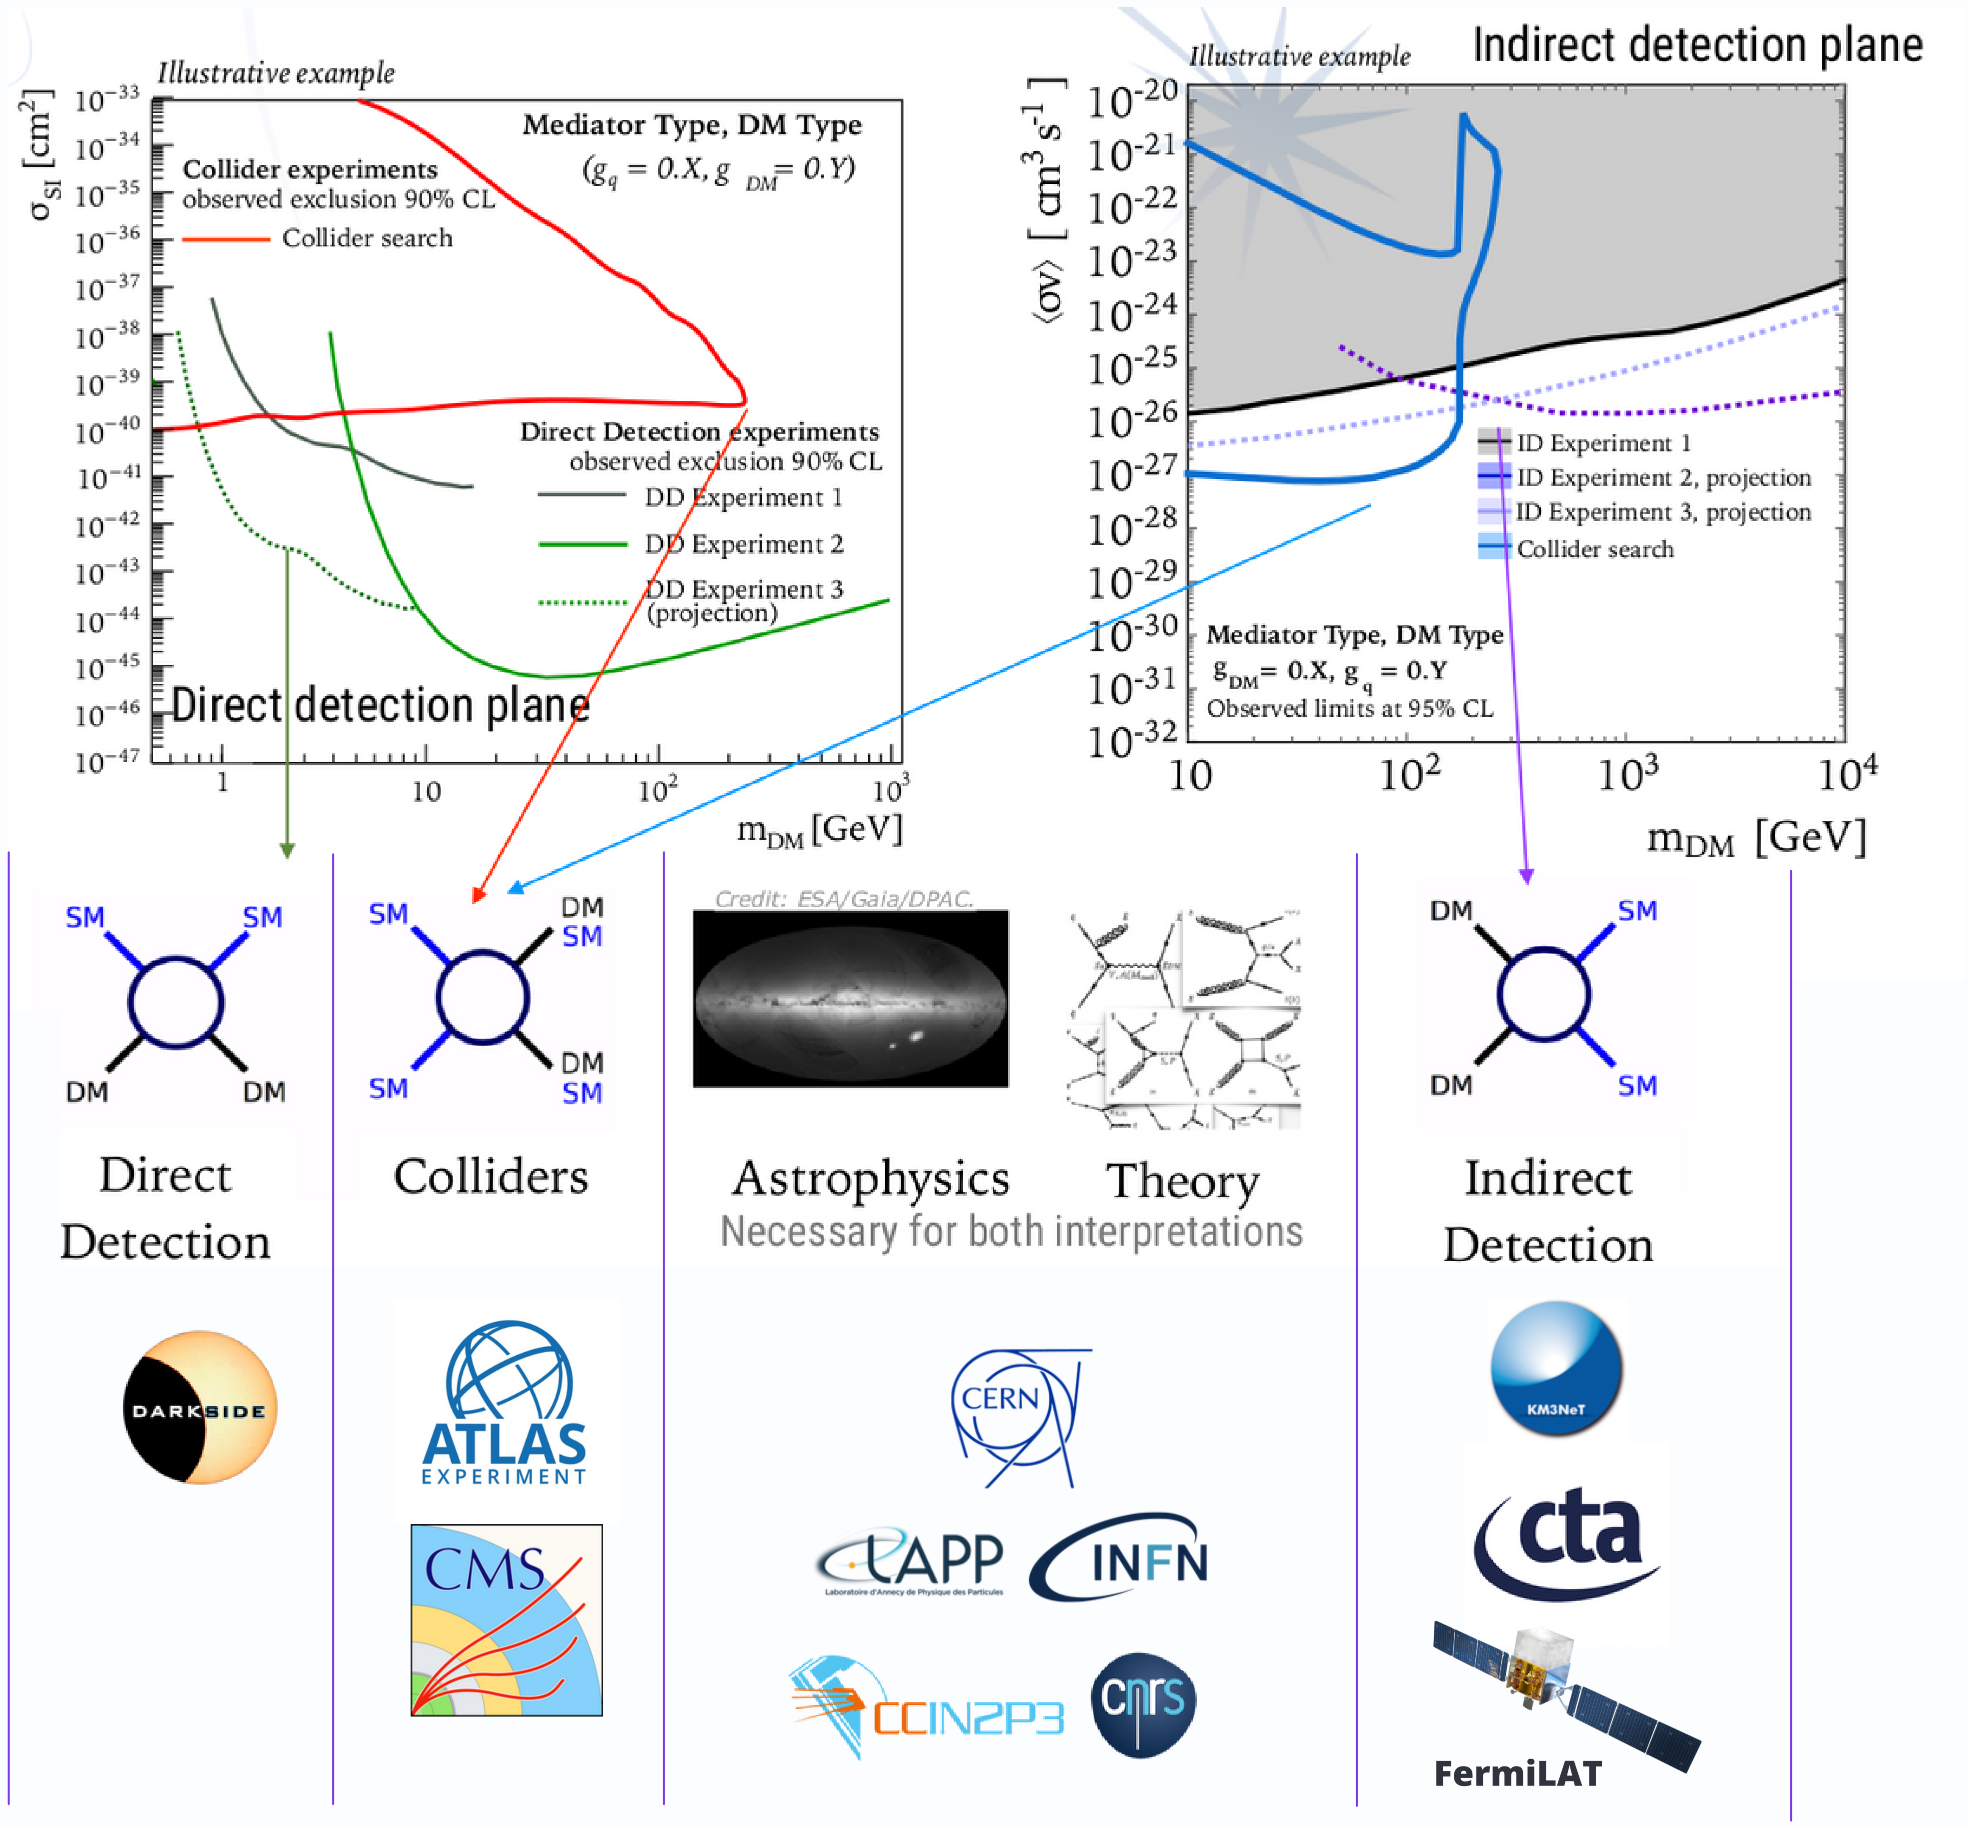

Extreme Universe searches

Multi-messenger astronomy: EM radiation, GW, neutrinos, cosmic rays are created by different astrophysical processes, and thus reveal different information about their sources)

INPUT DATA	Binary Neutron Star Merger	Active Galactic Nuclei	Core-Collapse Supernovae
STUDY	• GW • Fast Radio Bursts • Broadband follow ups	multi wavelength observations	• Neutrinos • GW
EXPERIMENT	VIRGO LOFAR	FermiLAT	VIRGO KM3NeT

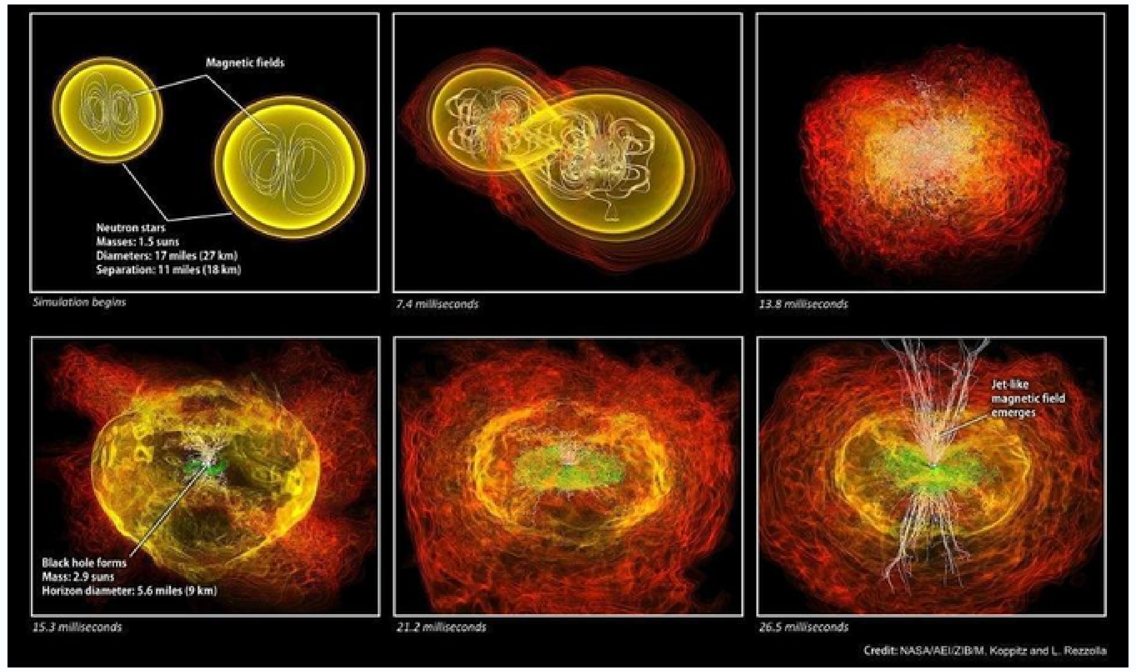

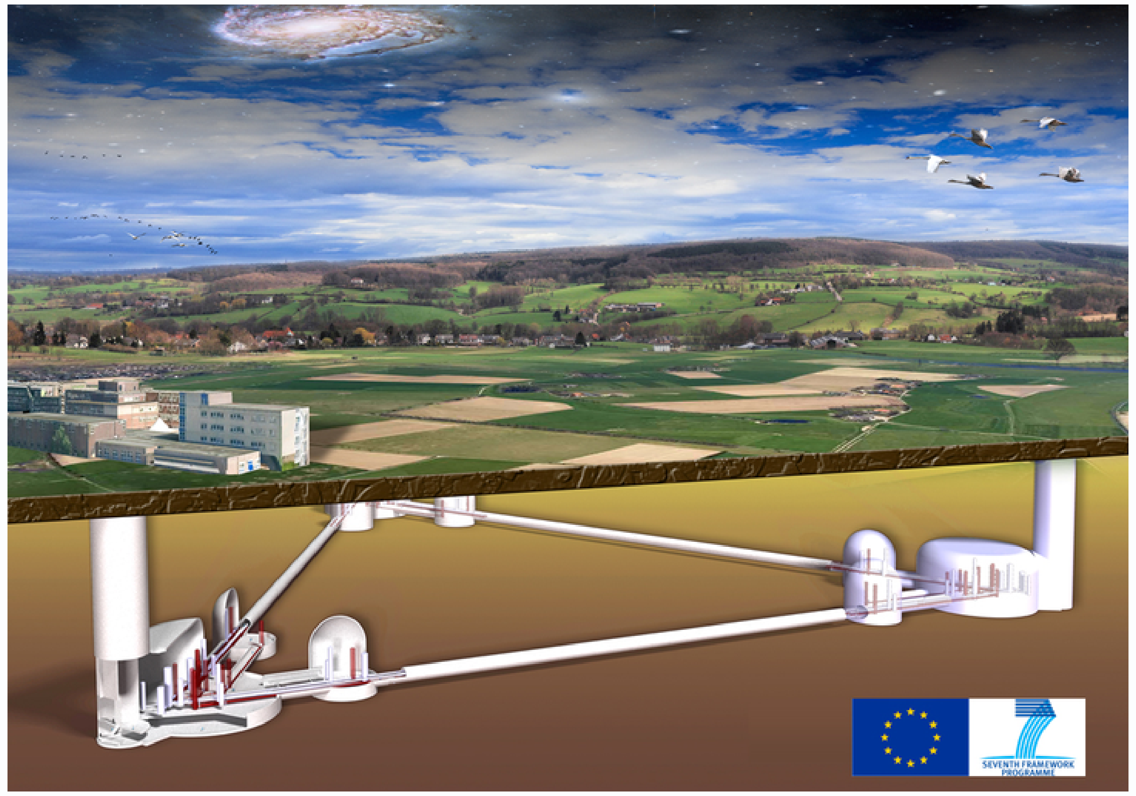

VRE Components

The building blocks

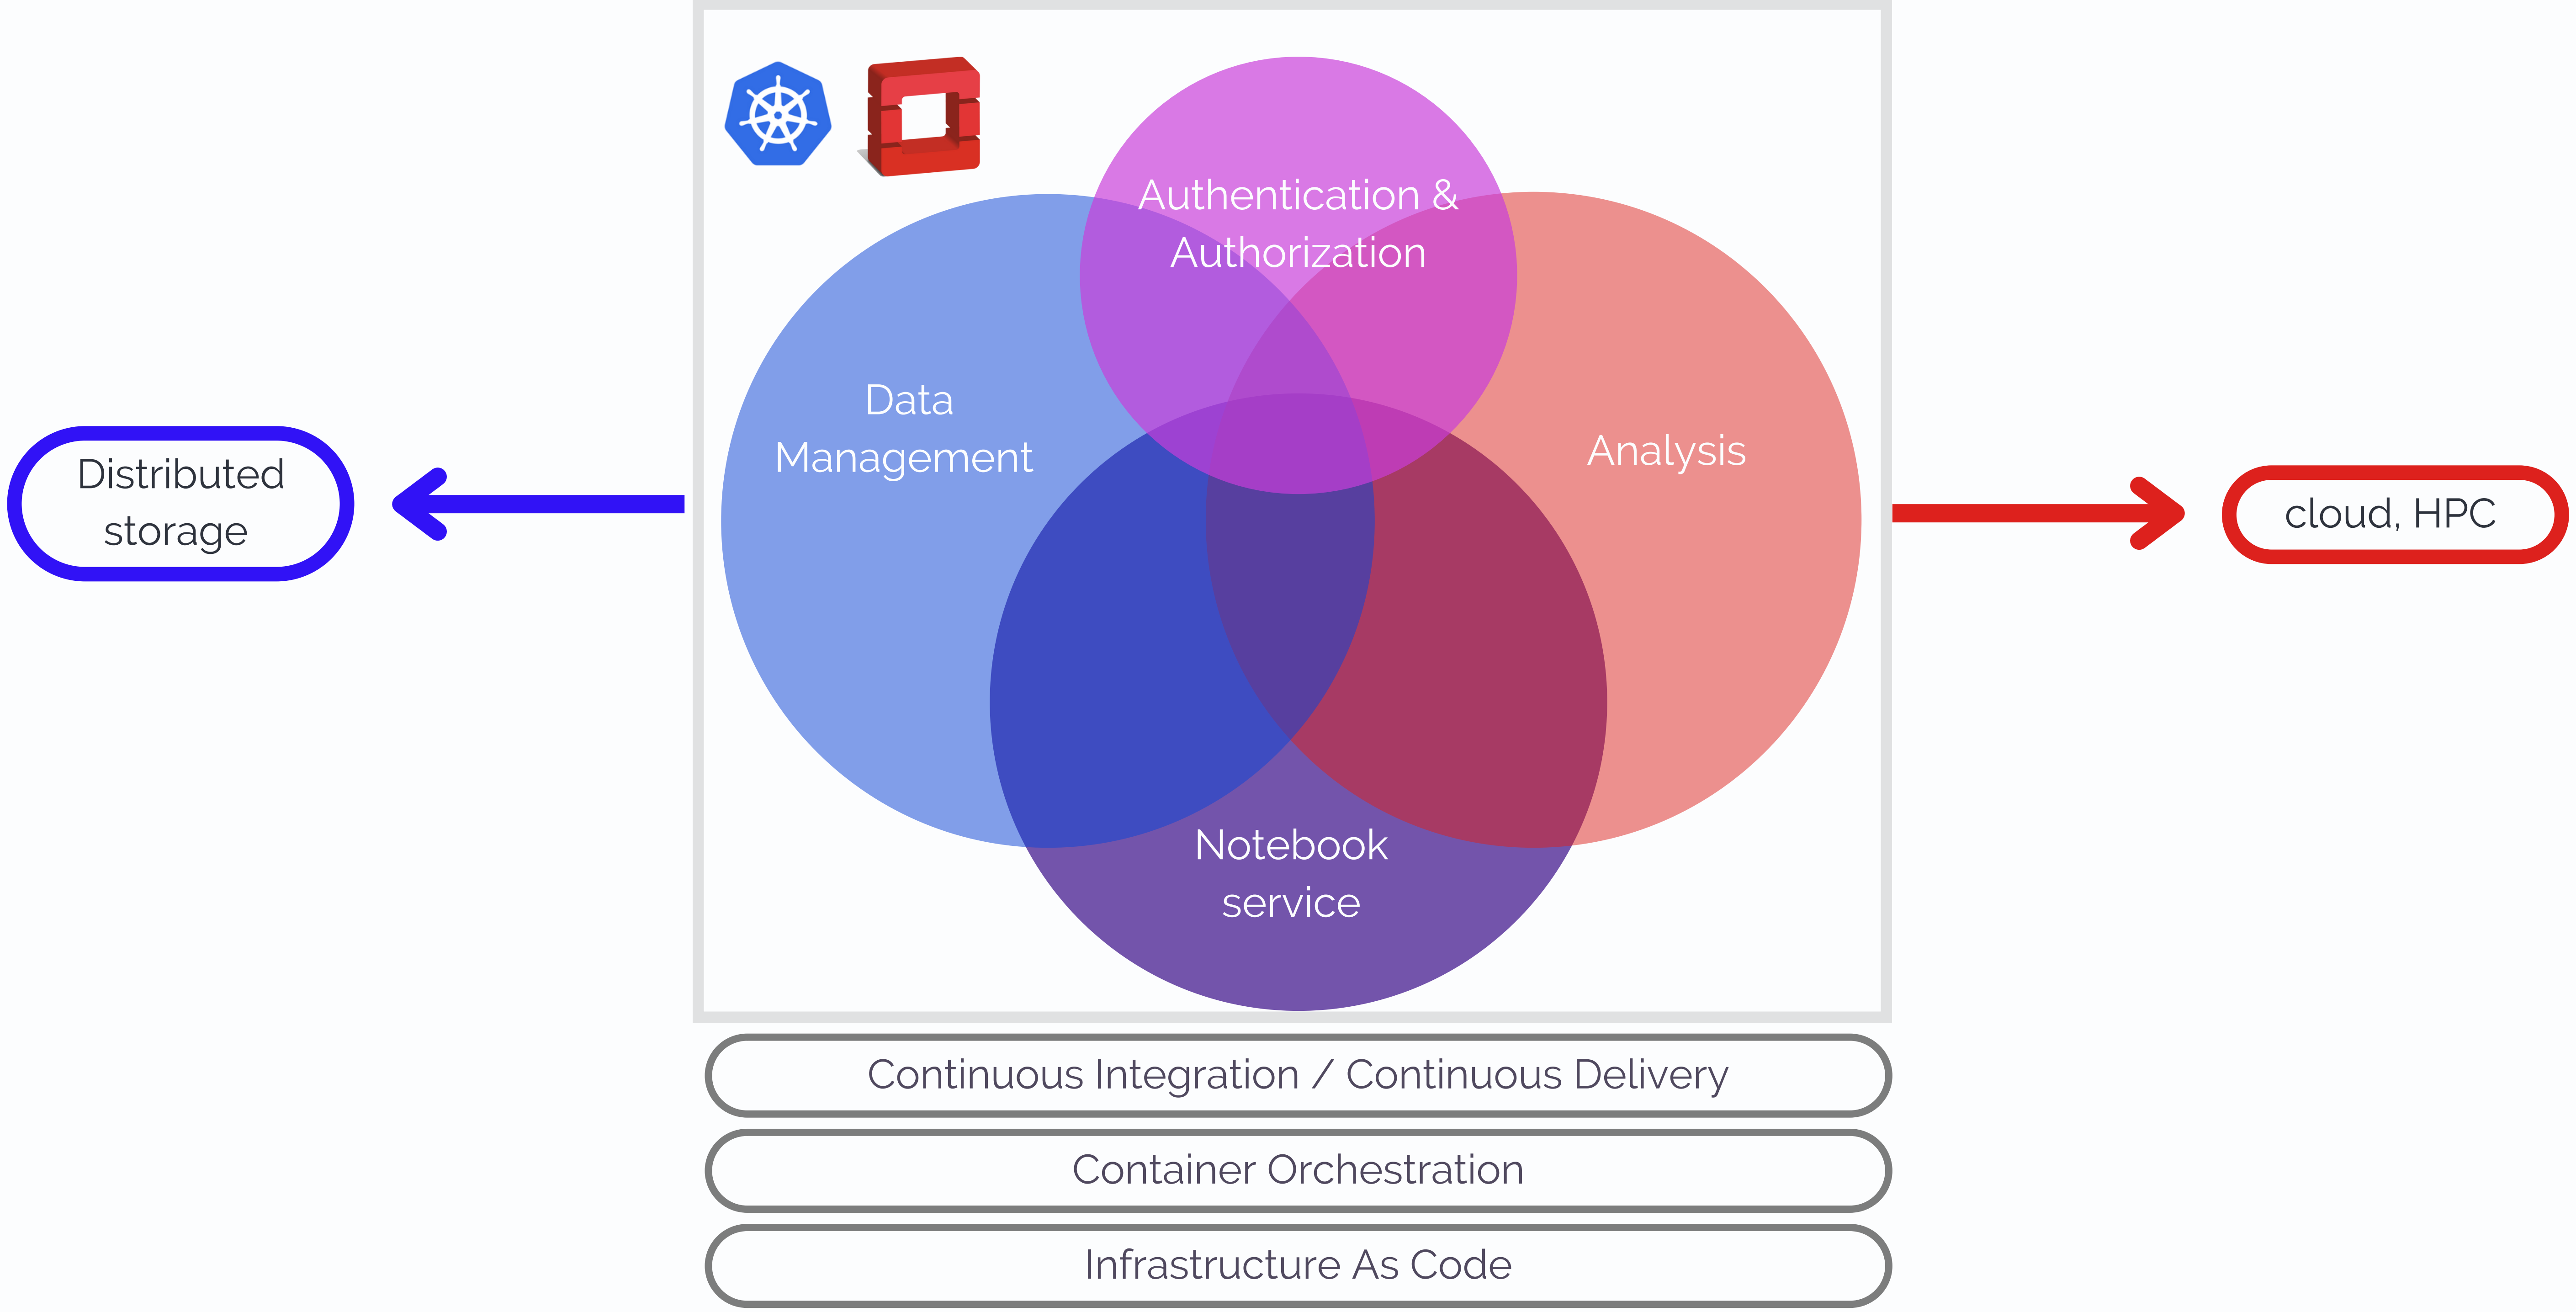

Authentication & Authorisation

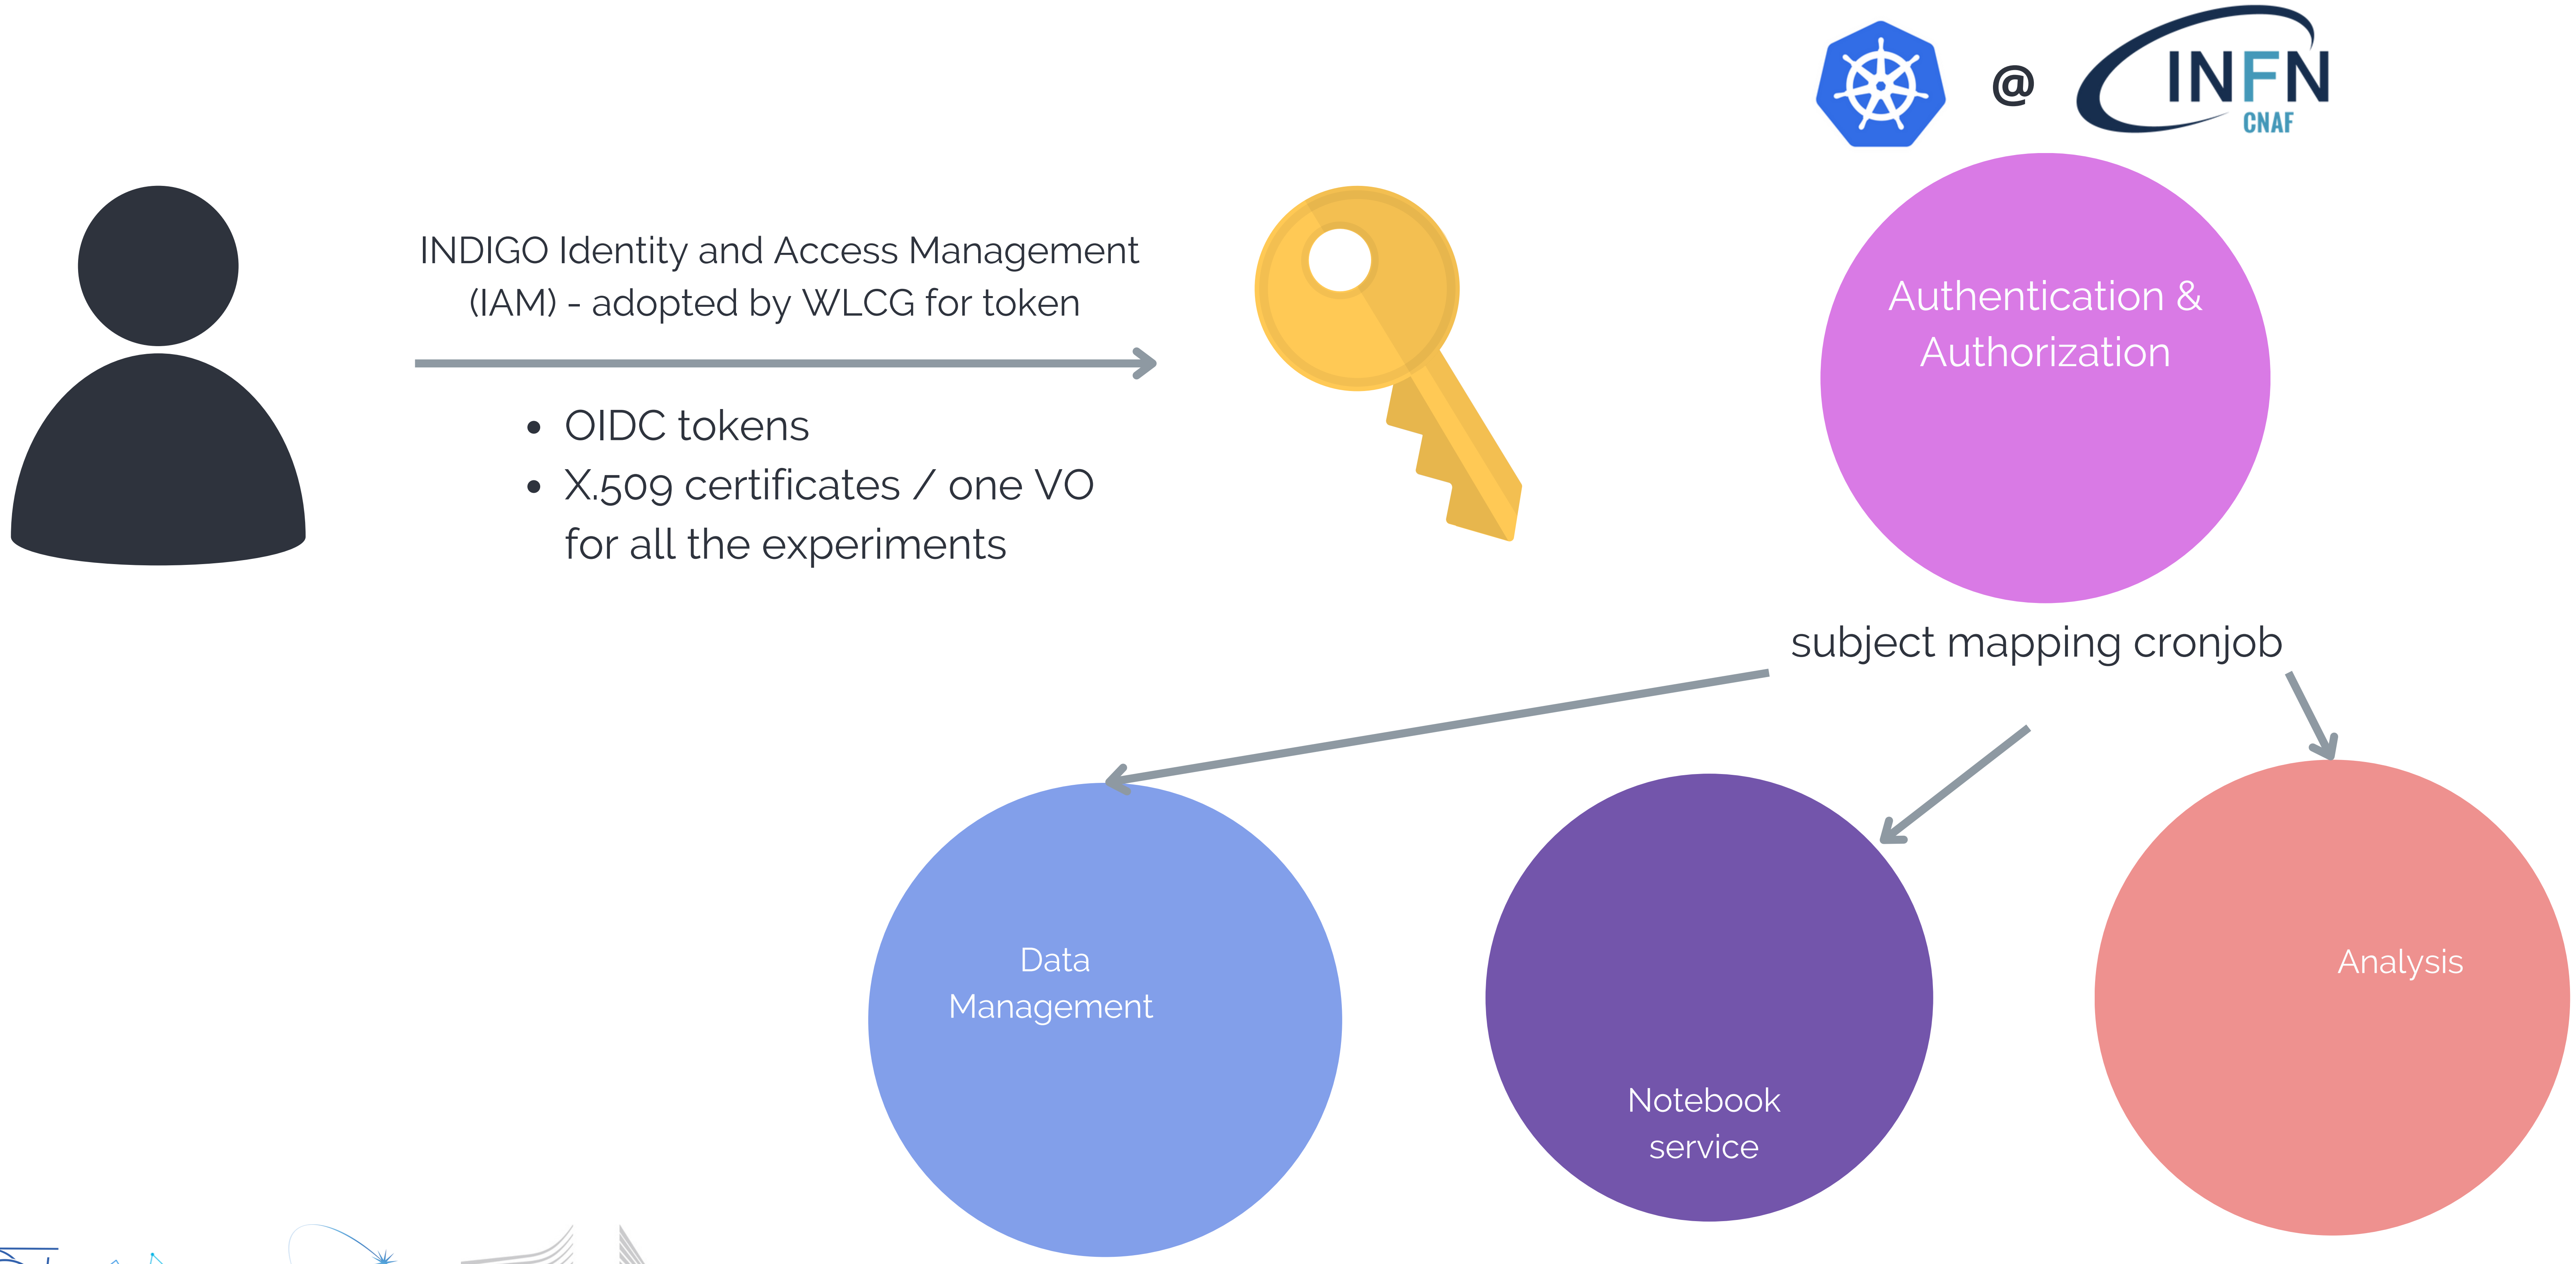

Data Management

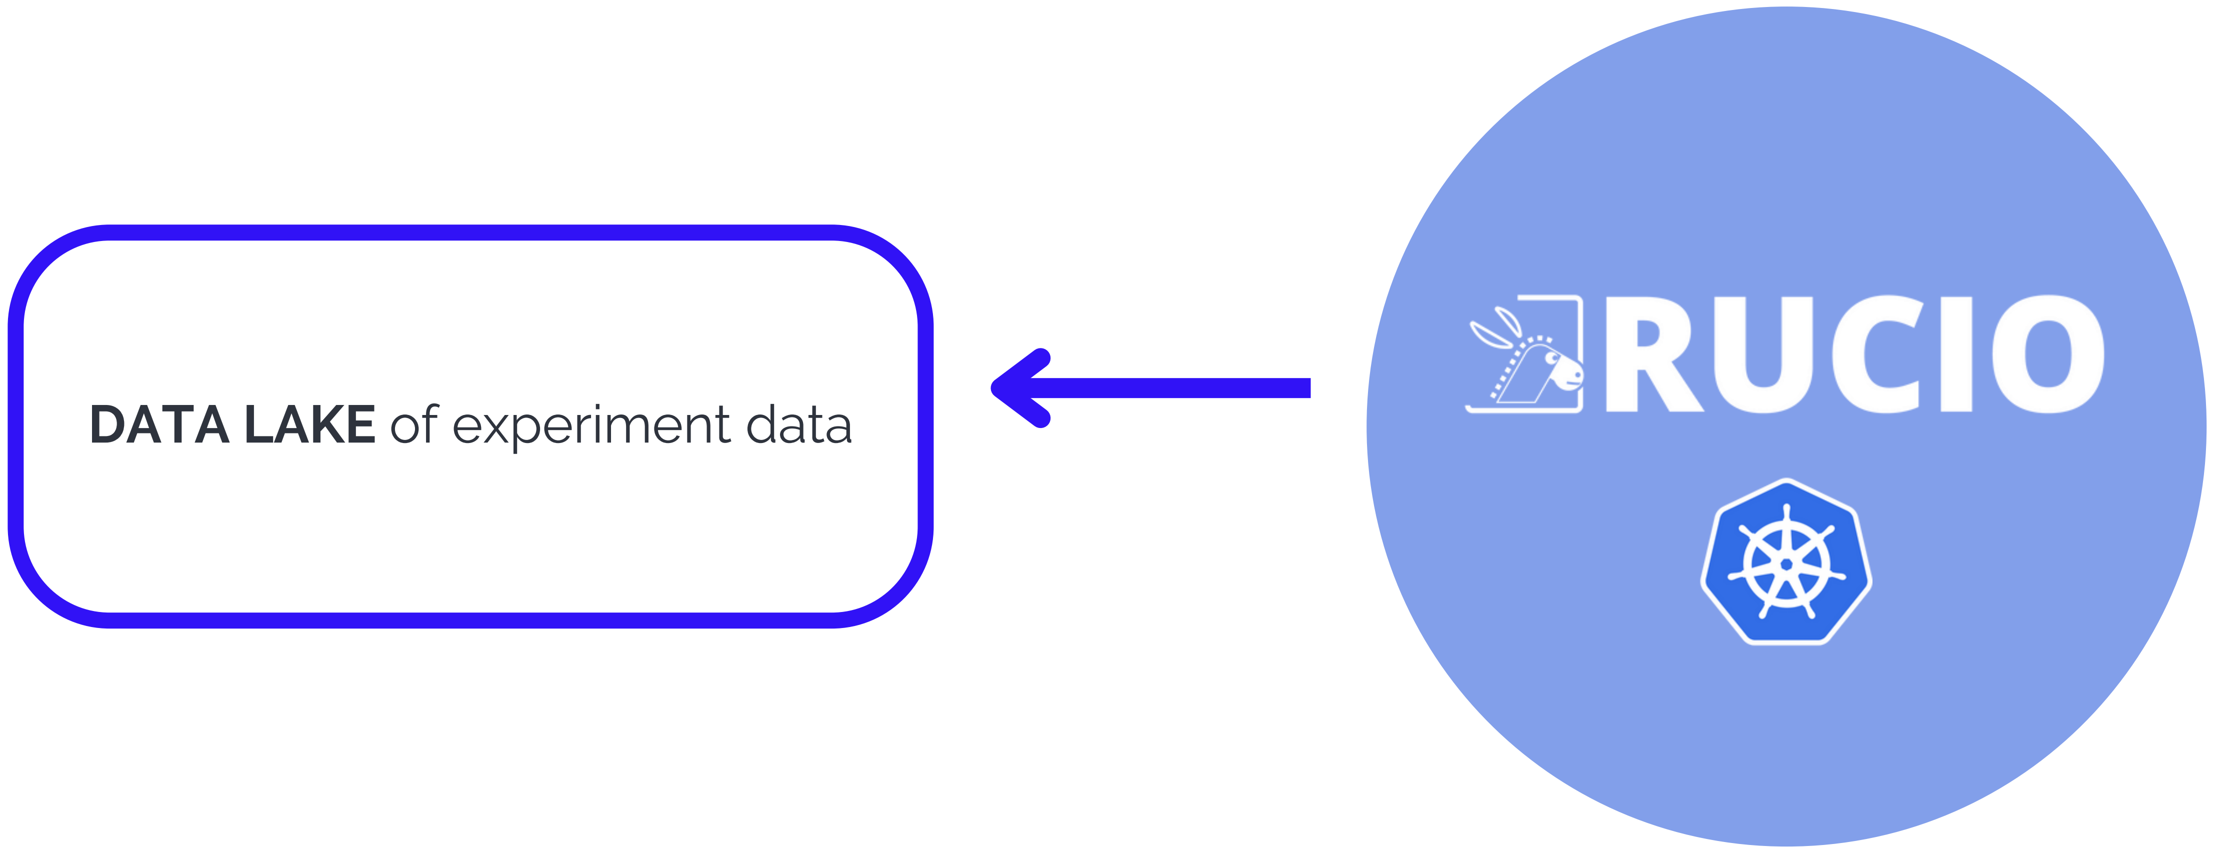

Rucio is an **open-source data management and orchestration** project initially developed by the ATLAS experiment to manage large volumes of data. It is now used by various CERN and non-CERN communities.

The **Data Lake** is a **policy-driven, reliable, distributed data infrastructure** able to deliver data **on-demand at low latency** to all types of processing facilities. It ensures **data security, quality and access.** The storage elements are managed by partner institutions.

Rucio instance

Data Lake

- EOS
- DPM
- dCache
- StoRM
- XRootD
- S3

(see Mario Lassnig talk)

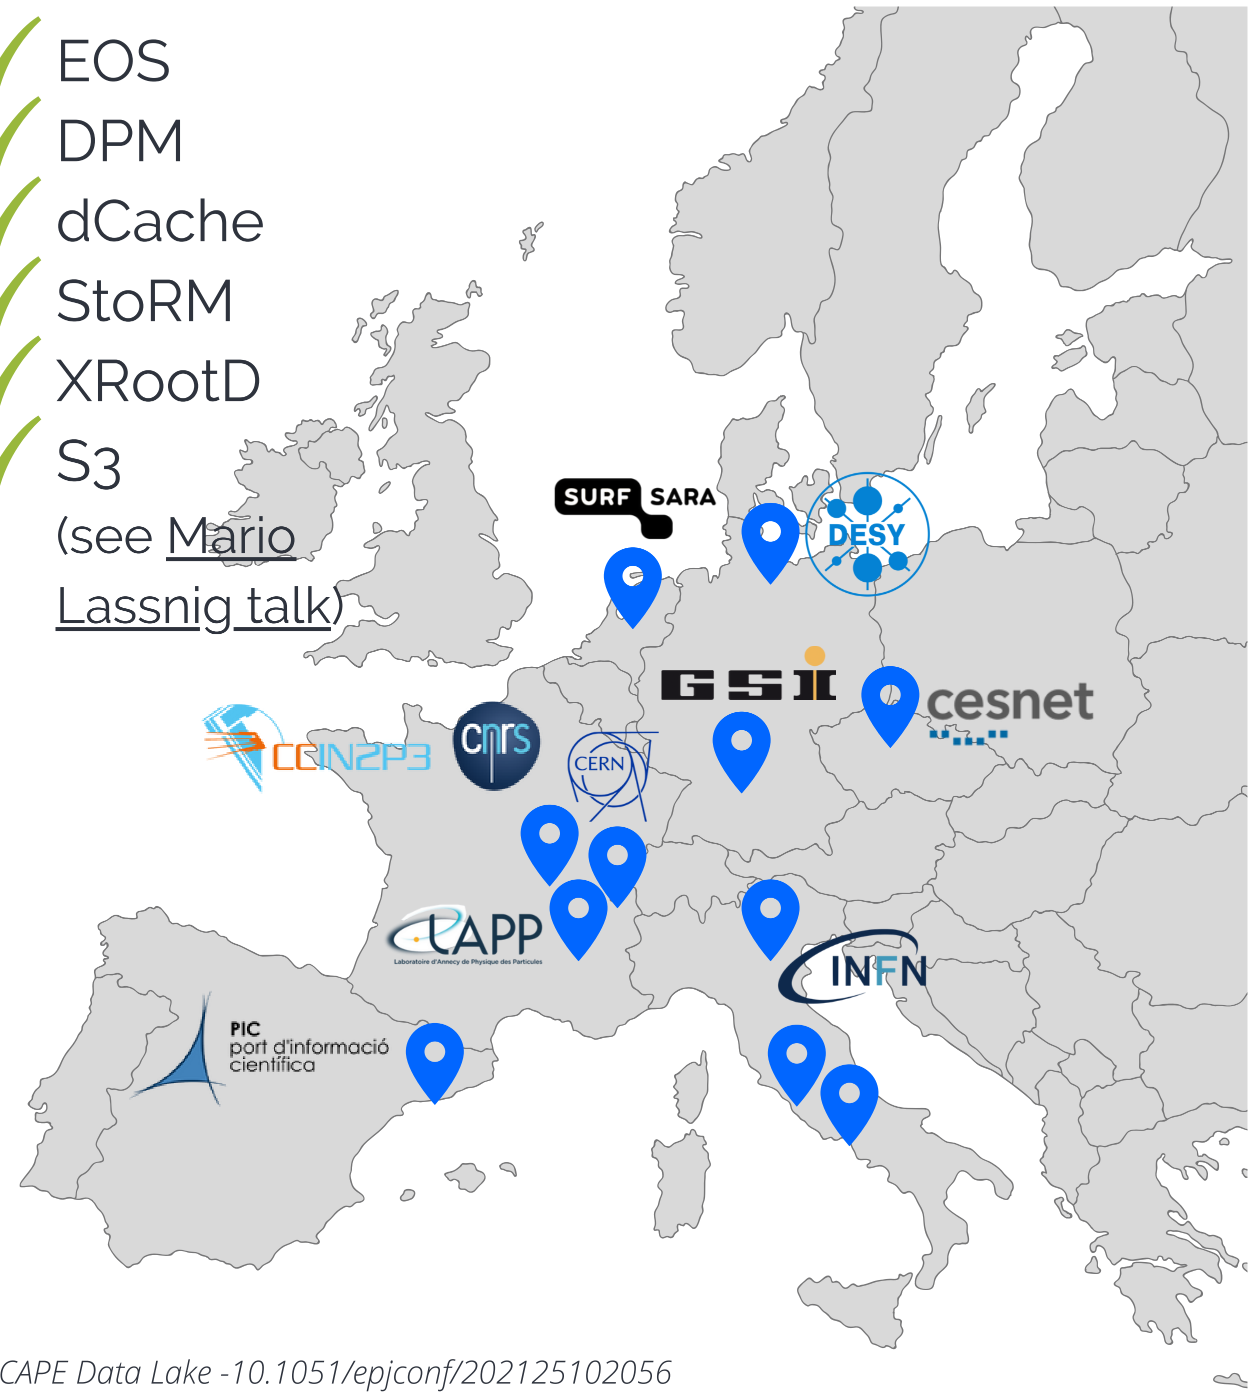

ESCAPE Data Lake -10.1051/epjconf/202125102056

Central Relational Database

FTS3and GFAL2

RUCIO

Daemons for data access & replications through gridFTP, HTTP(S)/webDAV, XRoot protocols

RUCIO

Main & authorisation servers

Notebook Service

To facilitate interactive analysis.

jupyter

interface to run preliminary analysis

docker

containerised environments on public repositories

CERN Virtual Machine FS (CVMFS) installed

zenodo

client libraries and software installed to interact with underlying services

ceph

CephFS volumes provided as shared, temporary storage solution

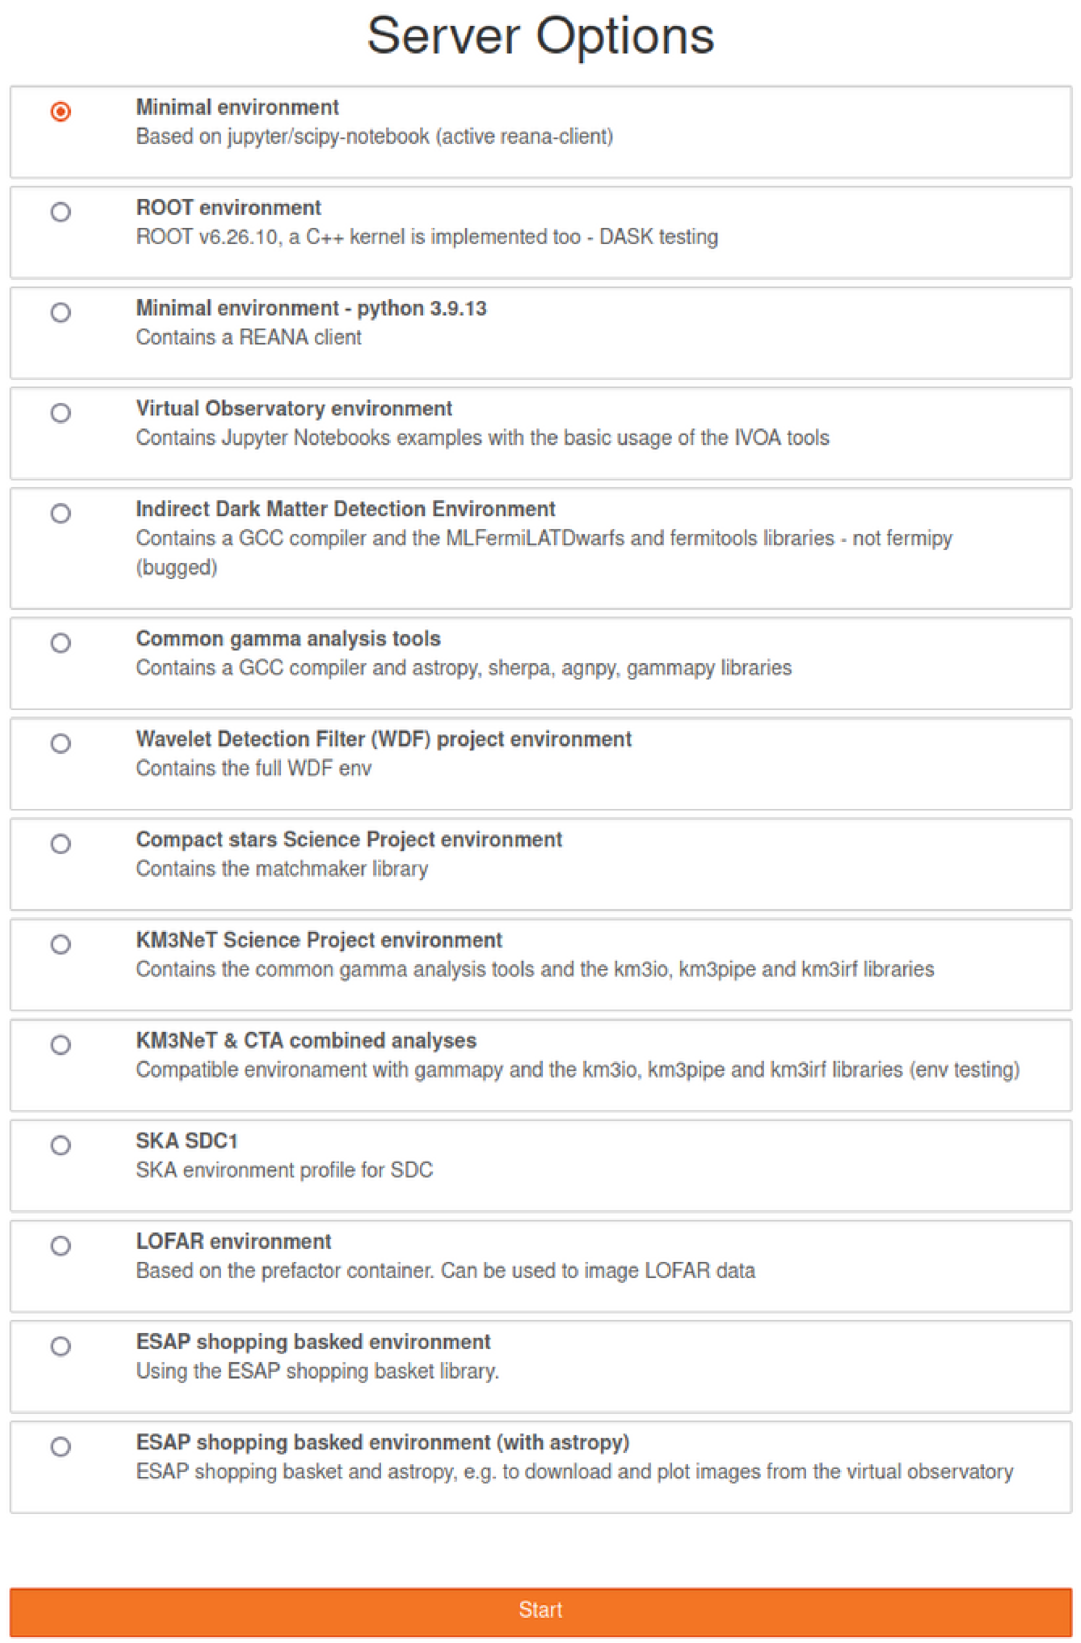

CERN EOSC Future ESCAPE European Commission

Data into the notebook

The **Jupyterhub Rucio extension** hides the complexity of the Data Lake and allows users to

- browse experiments' data catalogue
- authenticate with OIDC tokens to the Rucio infrastructure
- replicate data into the notebook
- import the data into the notebook by assigning a parameter to it
- run preliminary analysis to prototype code

RUCIO → jupyter

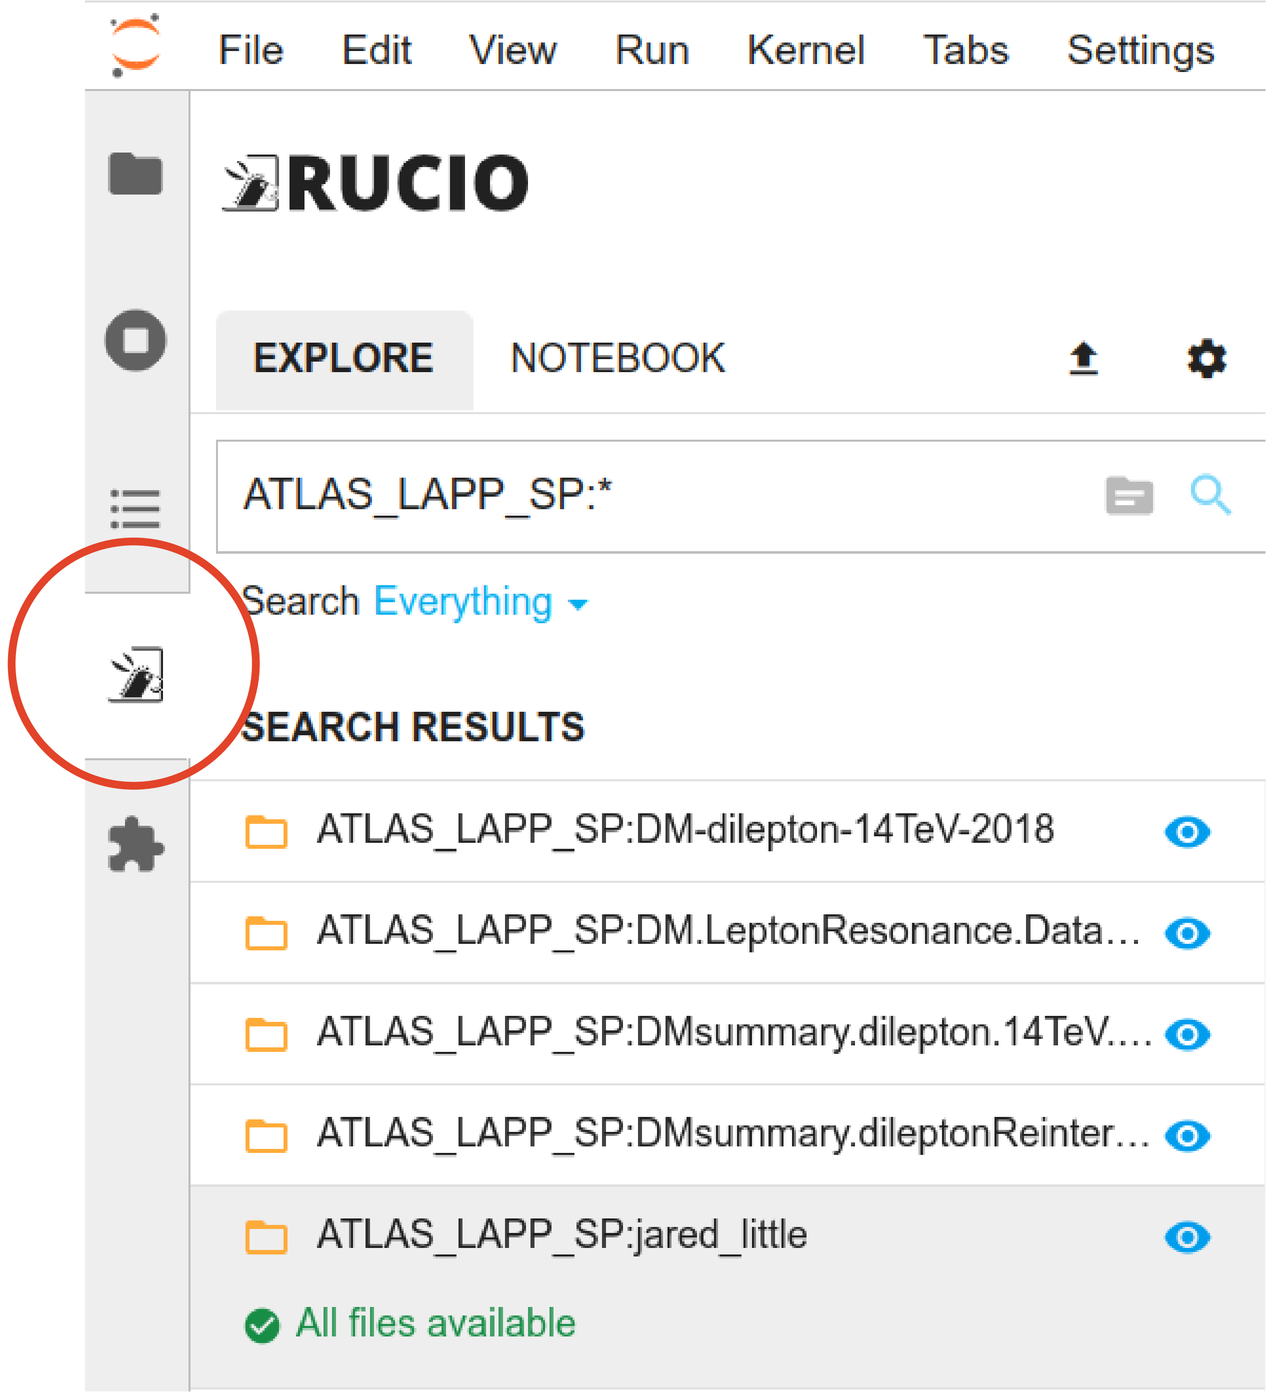

Data into the notebook

Data gets replicated through Rucio daemons from any storage element to an EOS storage element of half a Petabyte FUSE mounted on the Jupyterhub node.

The computation is limited to the CPU capacity of the node.

How do we SCALE OUT?

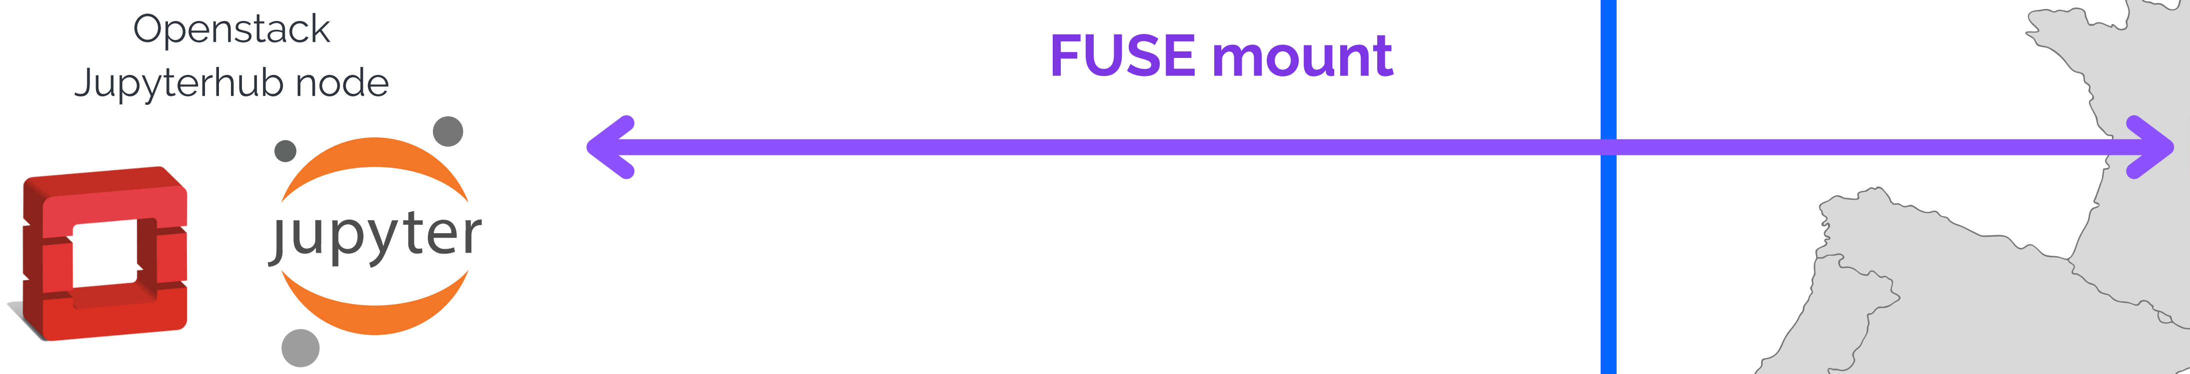

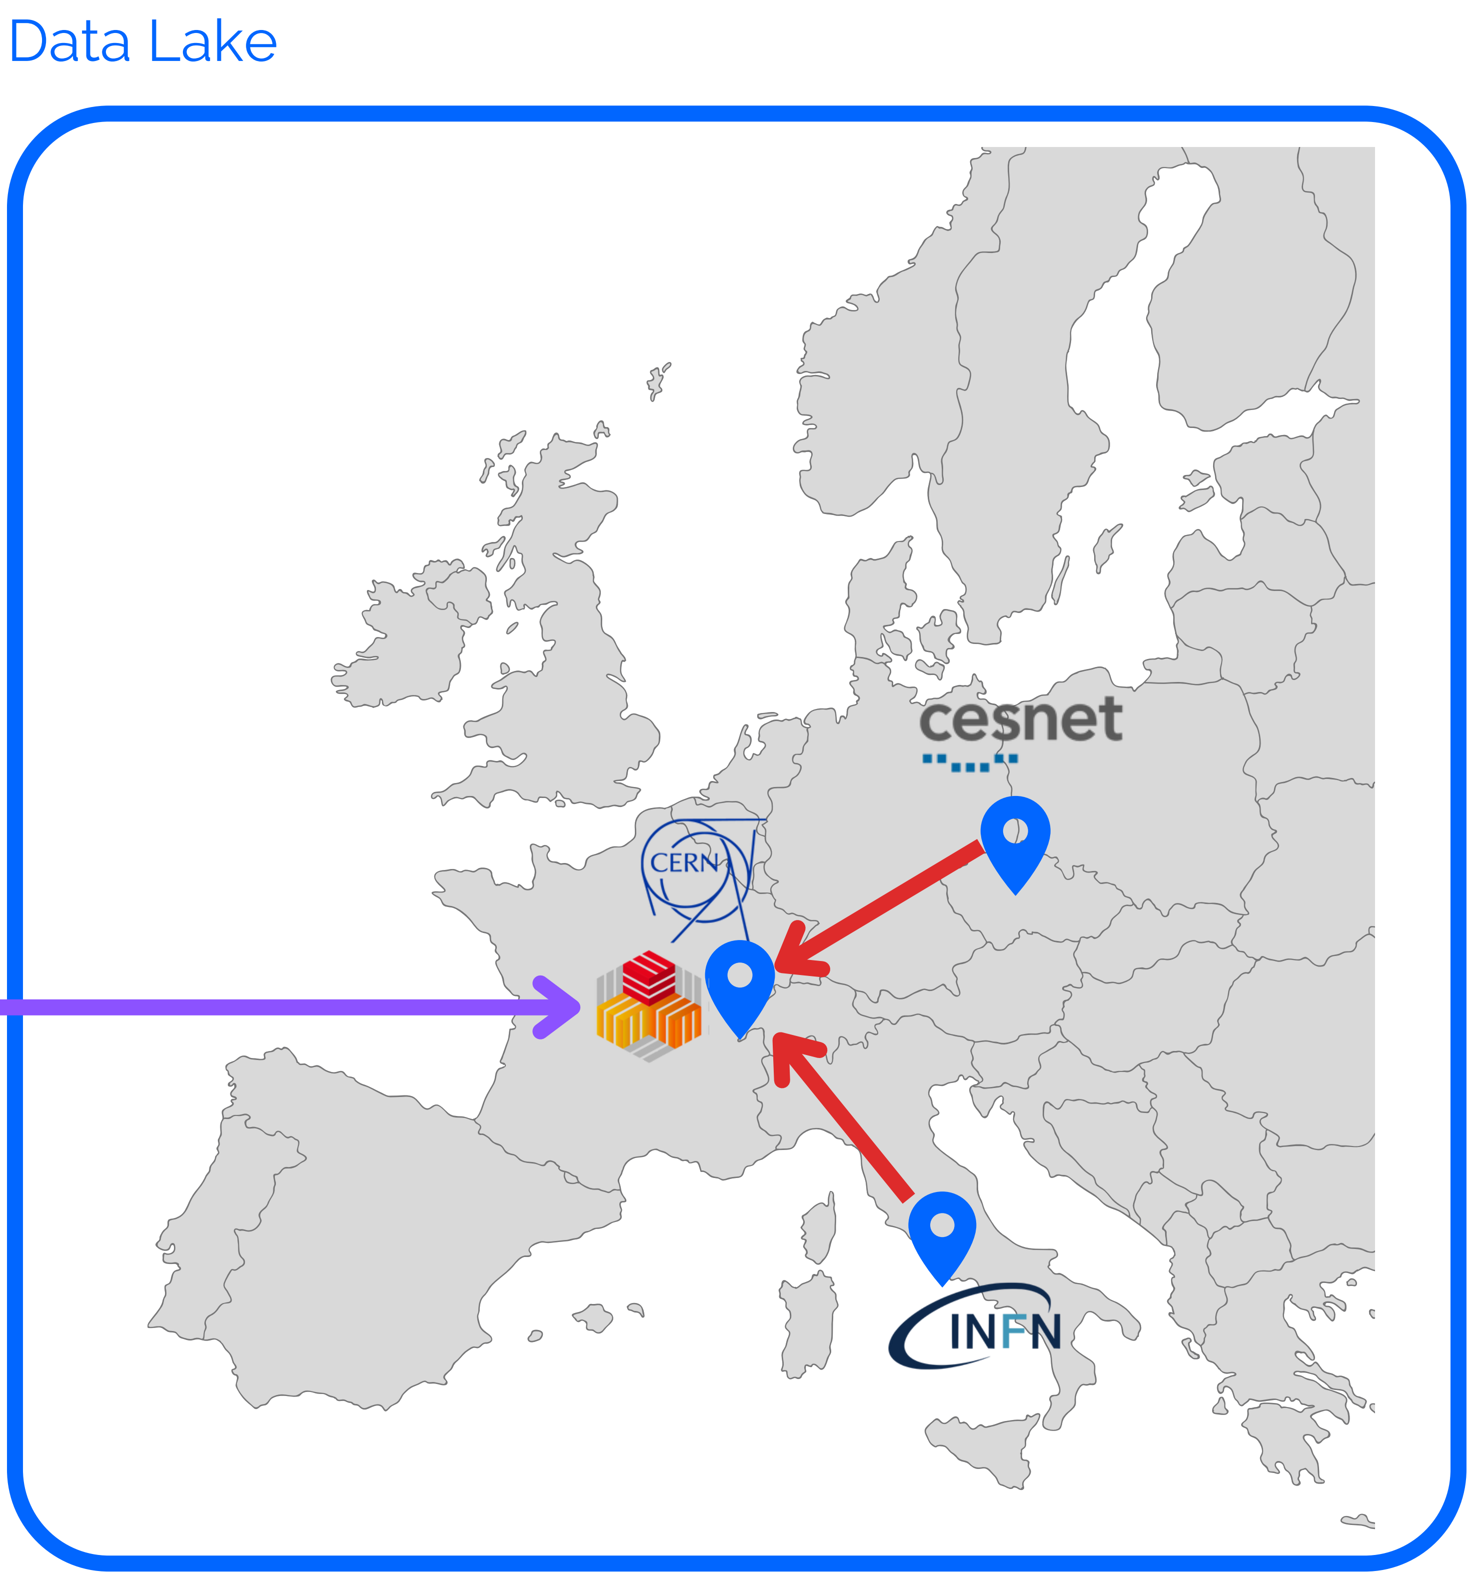

Computing

- **Distribute** the analysis
  - **resource managers** (Kubernetes, HTCondor (High Throughput Computing (HTC)) and Slurm (High Performance Computing (HPC))
  - **work schedulers** (Dask, Reana, Spark)

- **Preserve** the analysis for reuse
  - work schedulers (Reana)

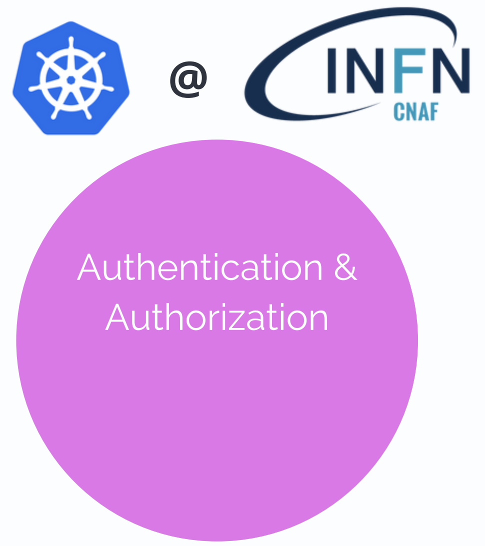

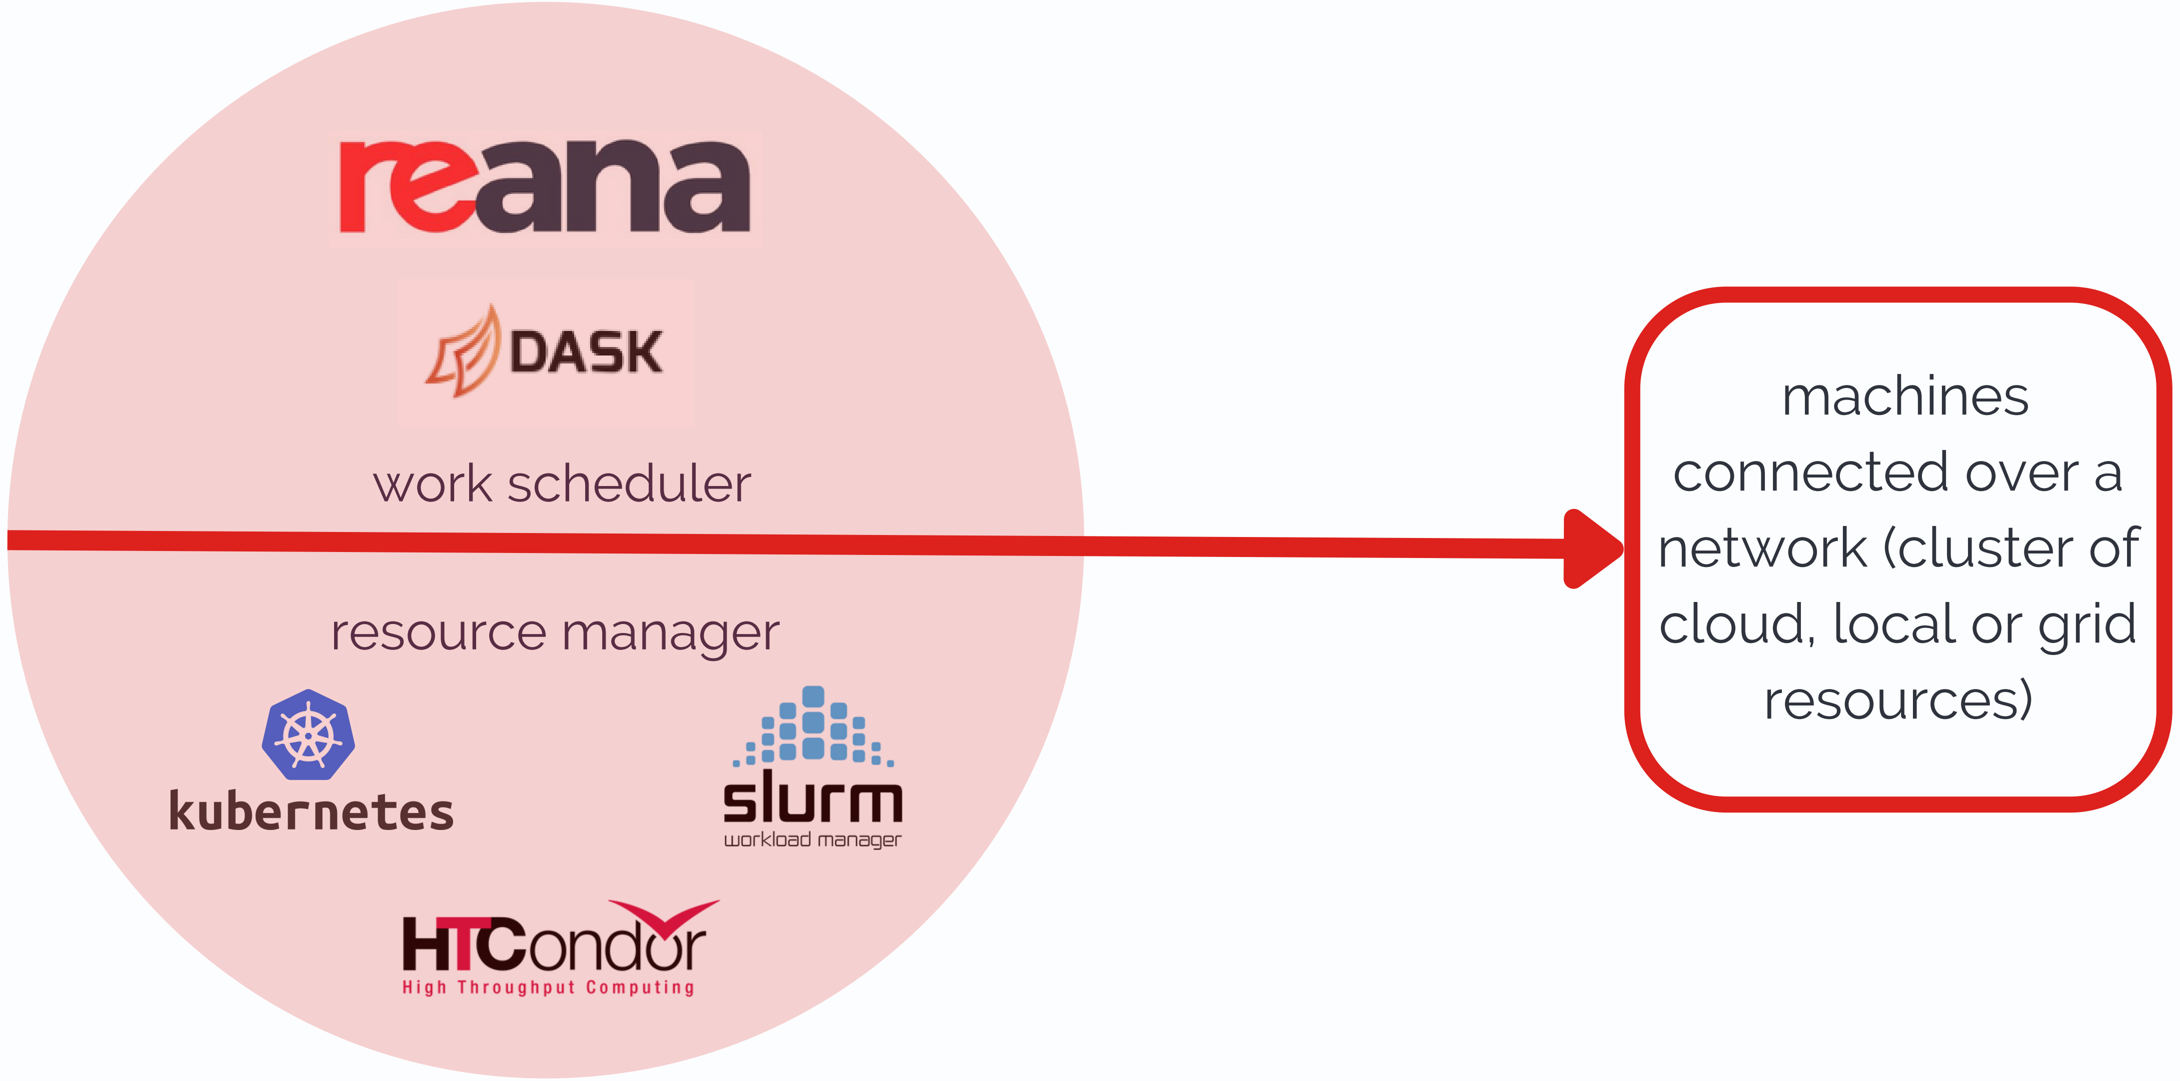

Analysis preservation and distribution

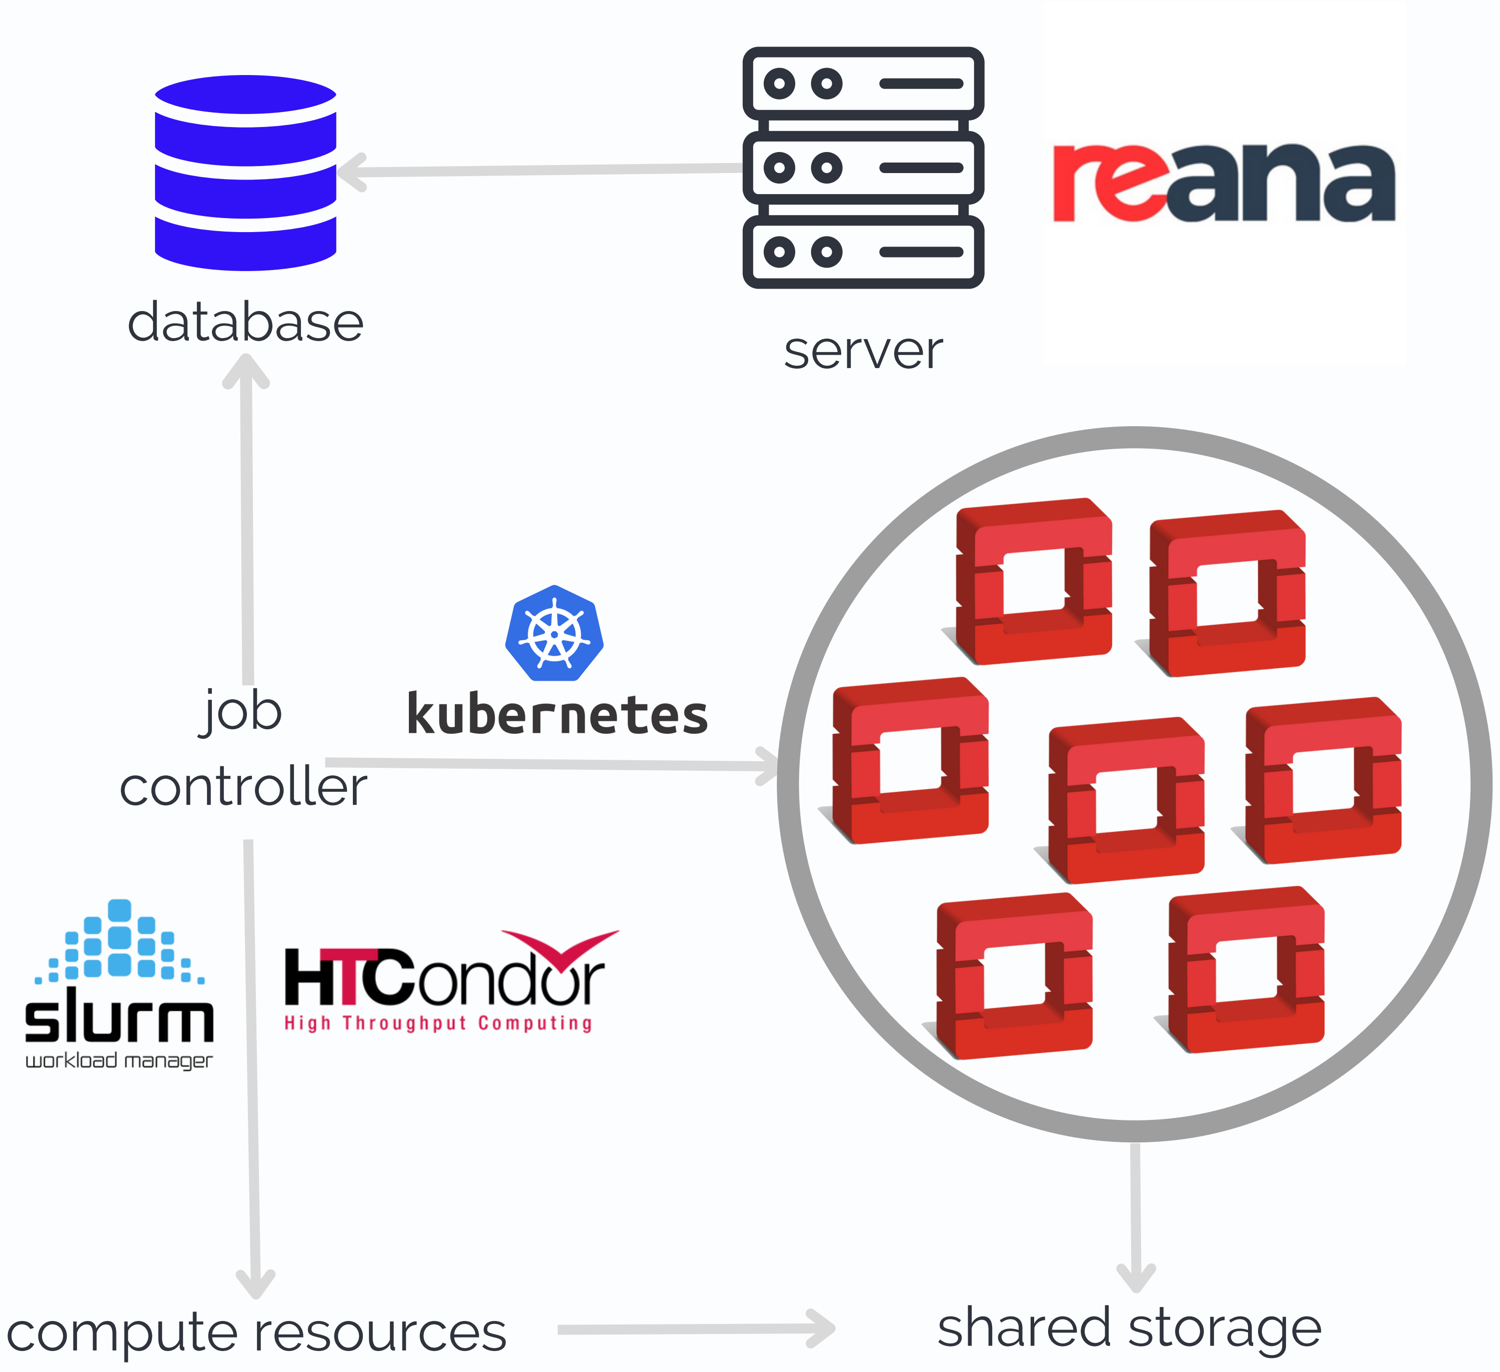

Reana is a reproducible analysis project developed at CERN, to make the preservation of heavier analyses seamless.

- Easily installed via Helm
- Intuitive declarative programming approach (reana.yaml file) with:
  - input data
  - environment
  - code
  - computational steps
- Isolates each step with different containers
- Supports workflow engines
  - CWL
  - Snakemake
  - Yadage --> workflow concatenation (output becomes input)

Non-local analysis preservation

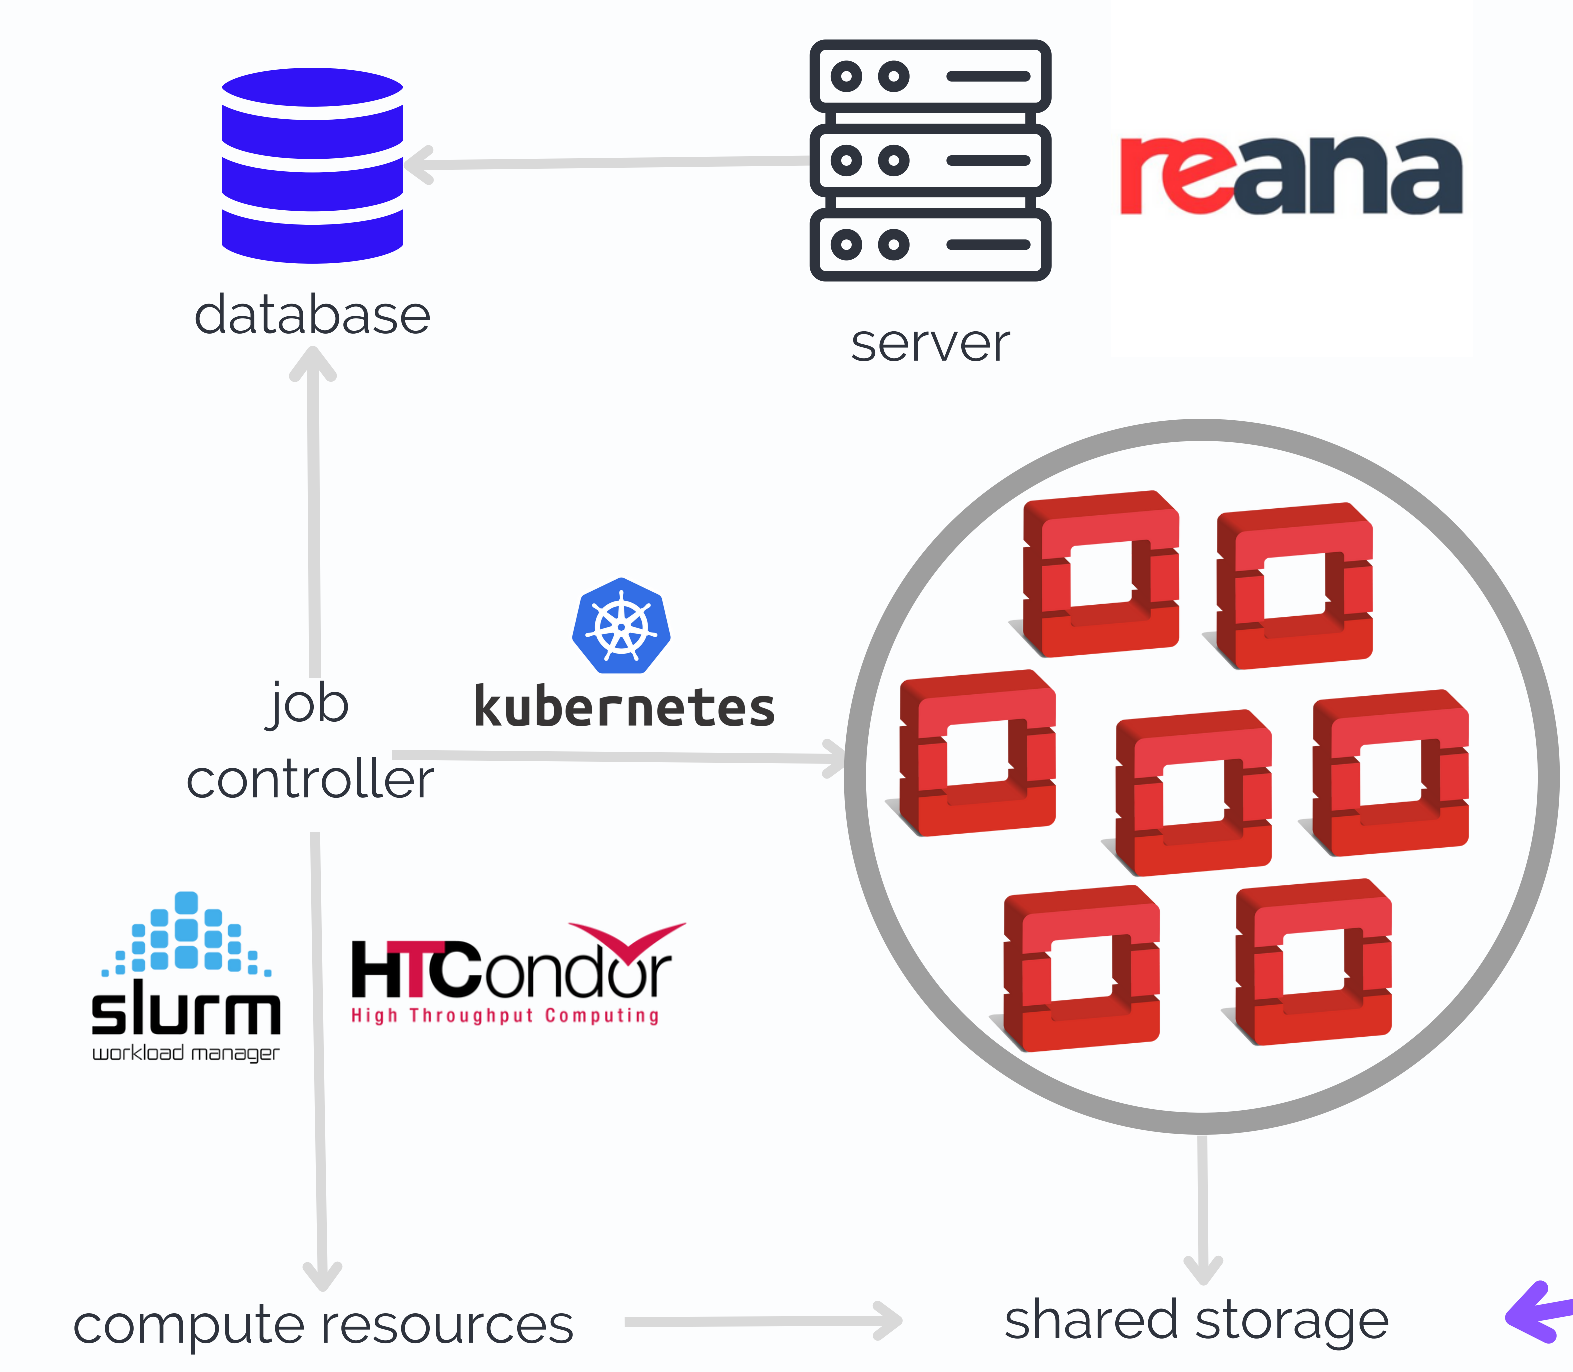

From the Reana client:

- authenticate via IAM to Rucio with a side-car container
- get data from distributed storage

→ the analysis can be reproduced fully and **independently from local storage**

Data Lake

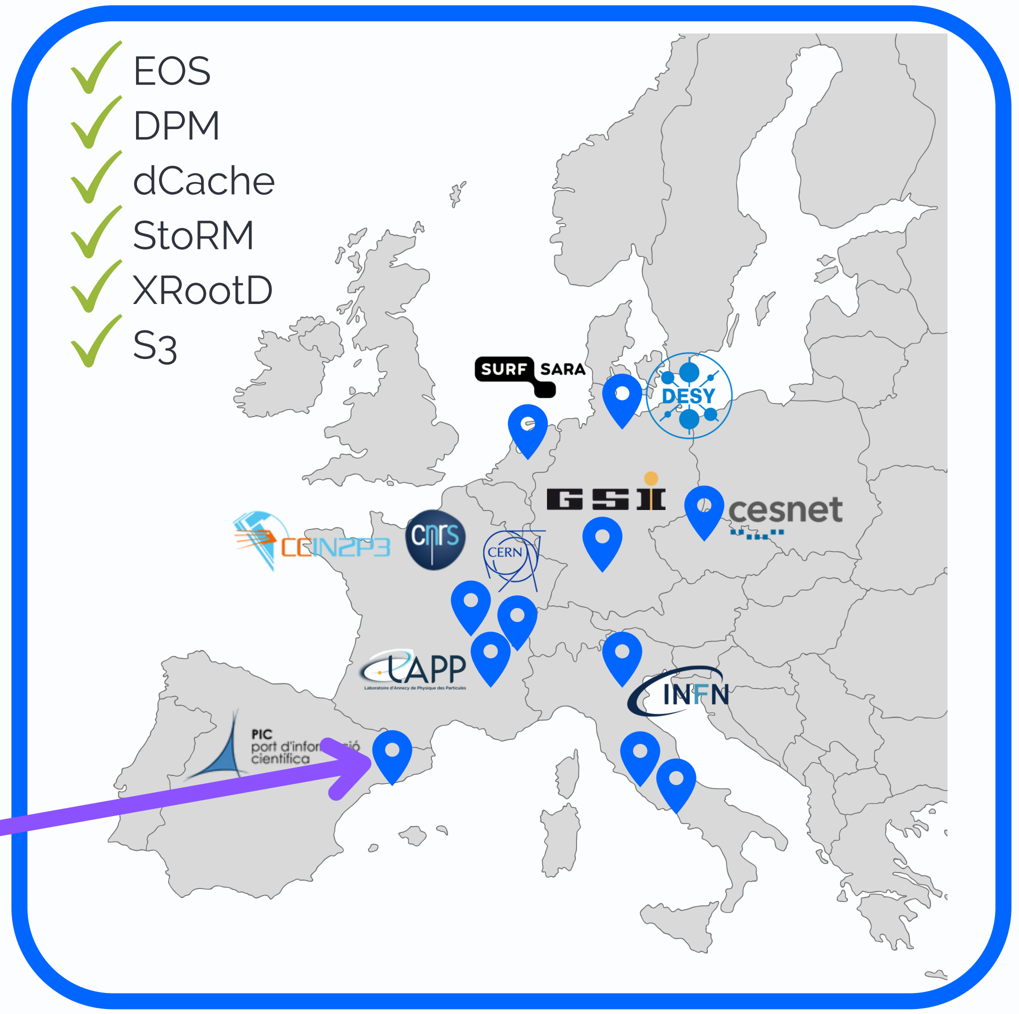

Workflow distribution with Dask

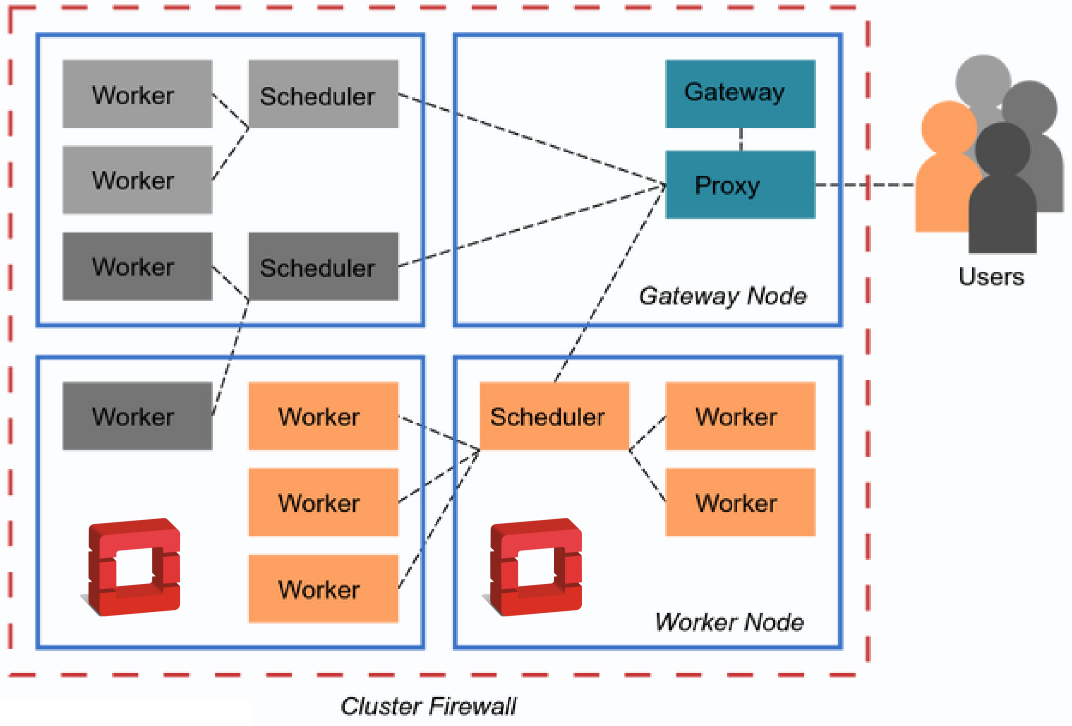

Daskhub helm chart: Dask Gateway + Jupyterhub

- multi-user, configurable usage profiles
- gateway to distribute access to all cloud nodes of the VRE
- code needs to be adapted
- dashboards of work progress

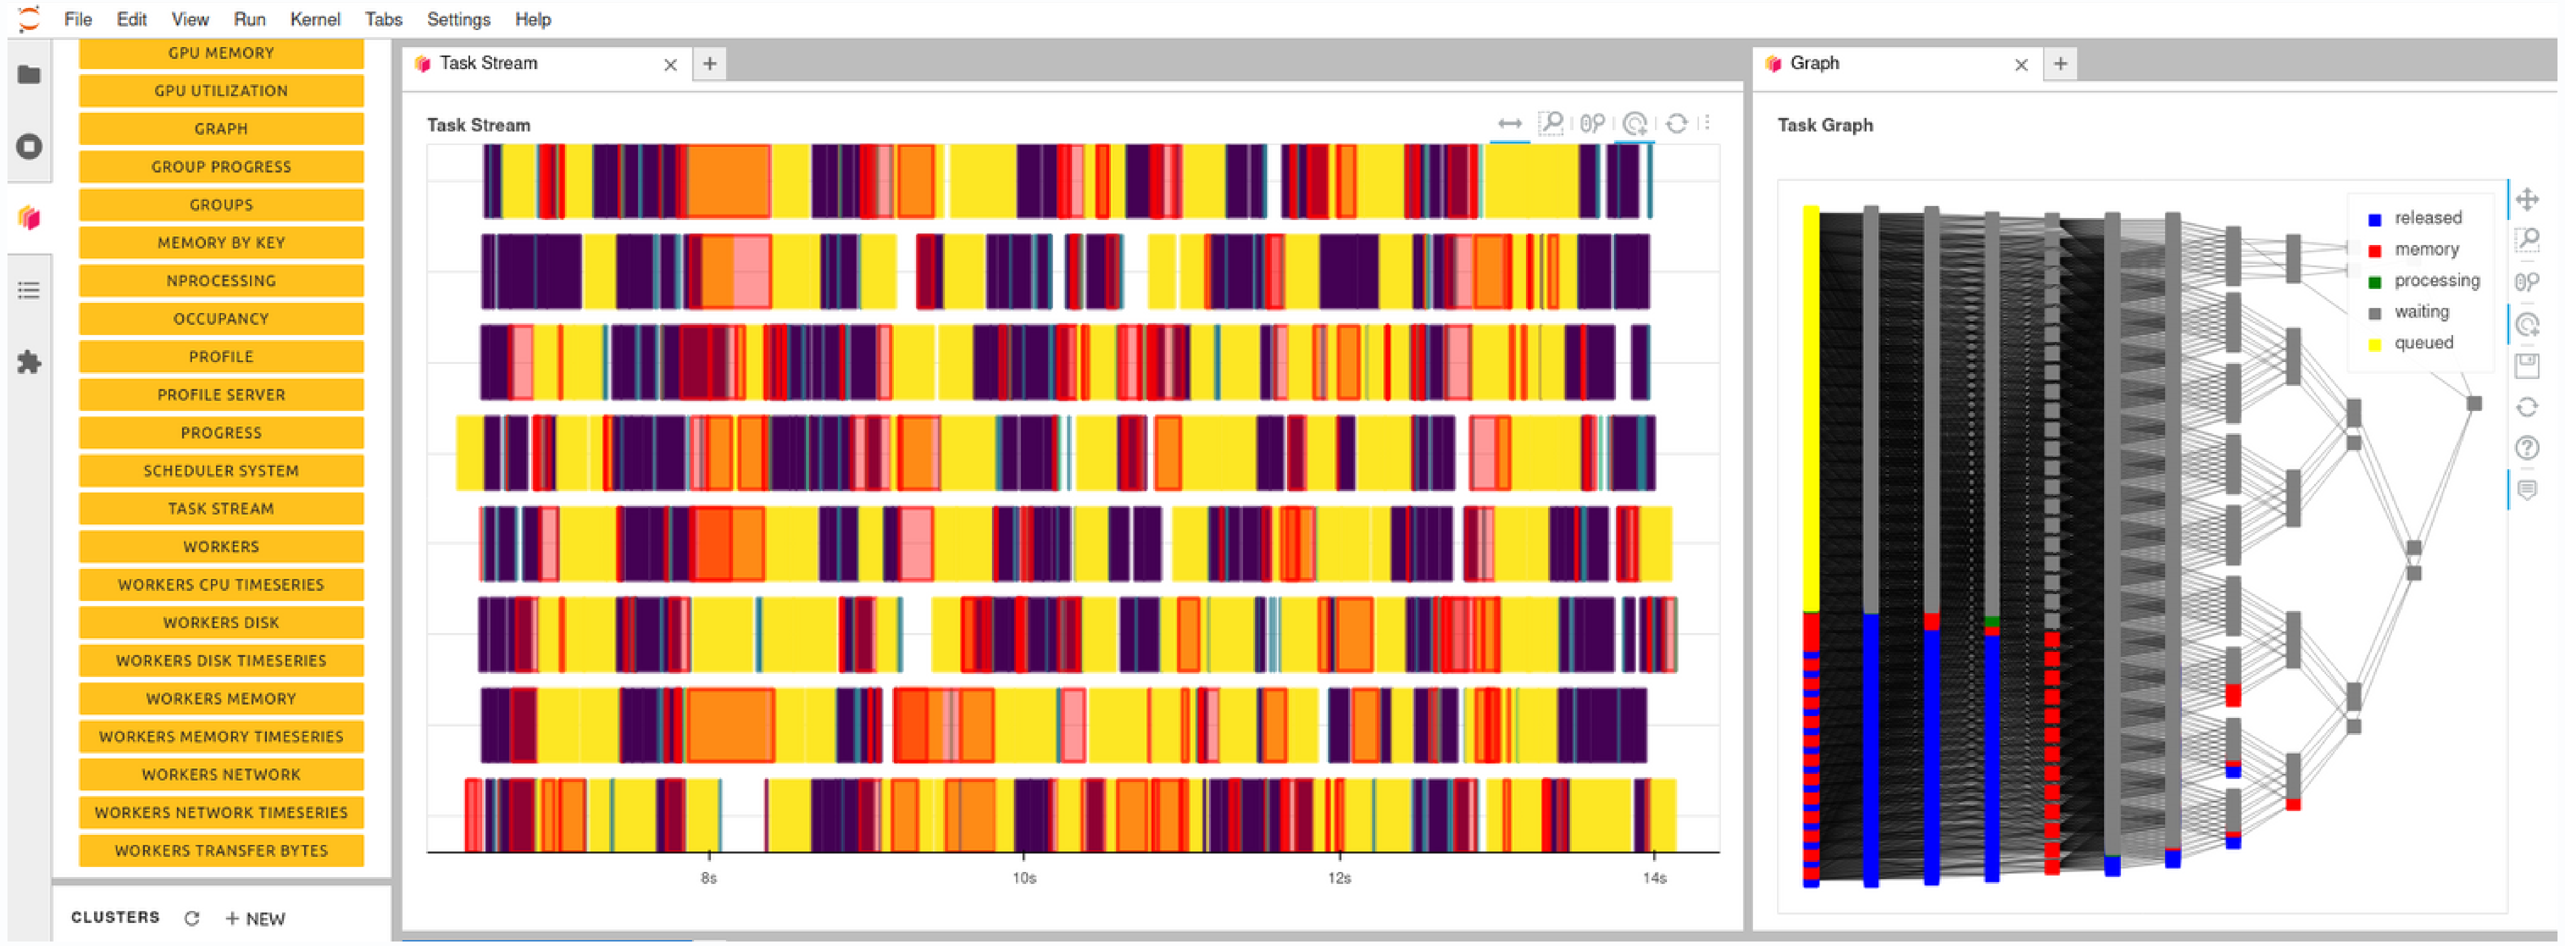

jupyter + DASK

Monitoring, testing, dashboards, on-boarding

- Continuous **monitoring and testing** of transfers between Rucio Storage Elements (RSEs) is in place on Grafana dashboards hosted at CERN.

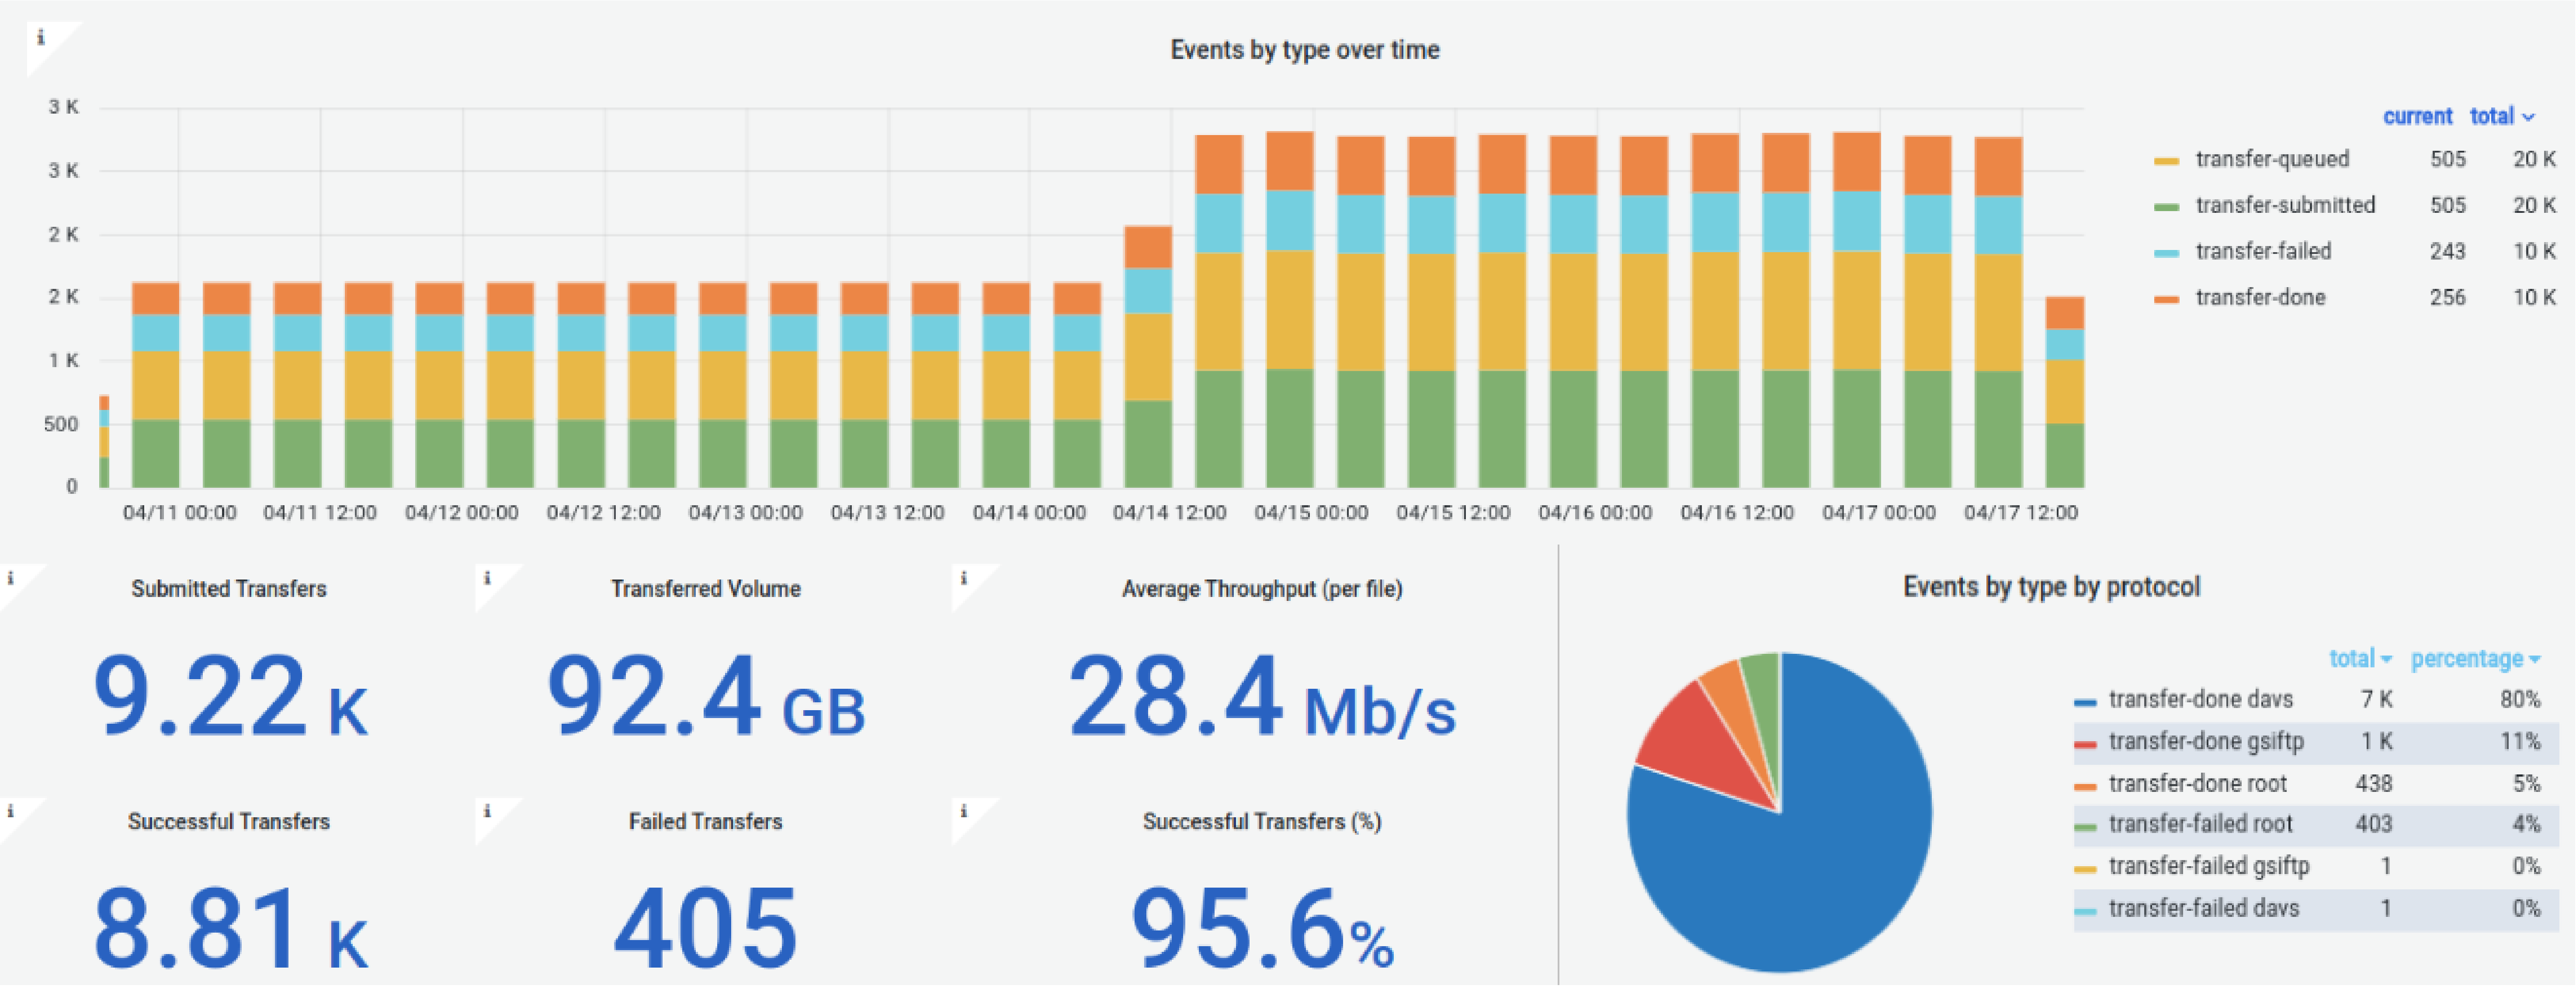

Monitoring, testing, dashboards, on-boarding

- Continuous **monitoring and testing** of transfers between Rucio Storage Elements (RSEs) is in place on Grafana dashboards hosted at CERN.
- **Rucio and Reana UI** interfaces deployed with K8s allow to explore and debug failed transfers and workflows.

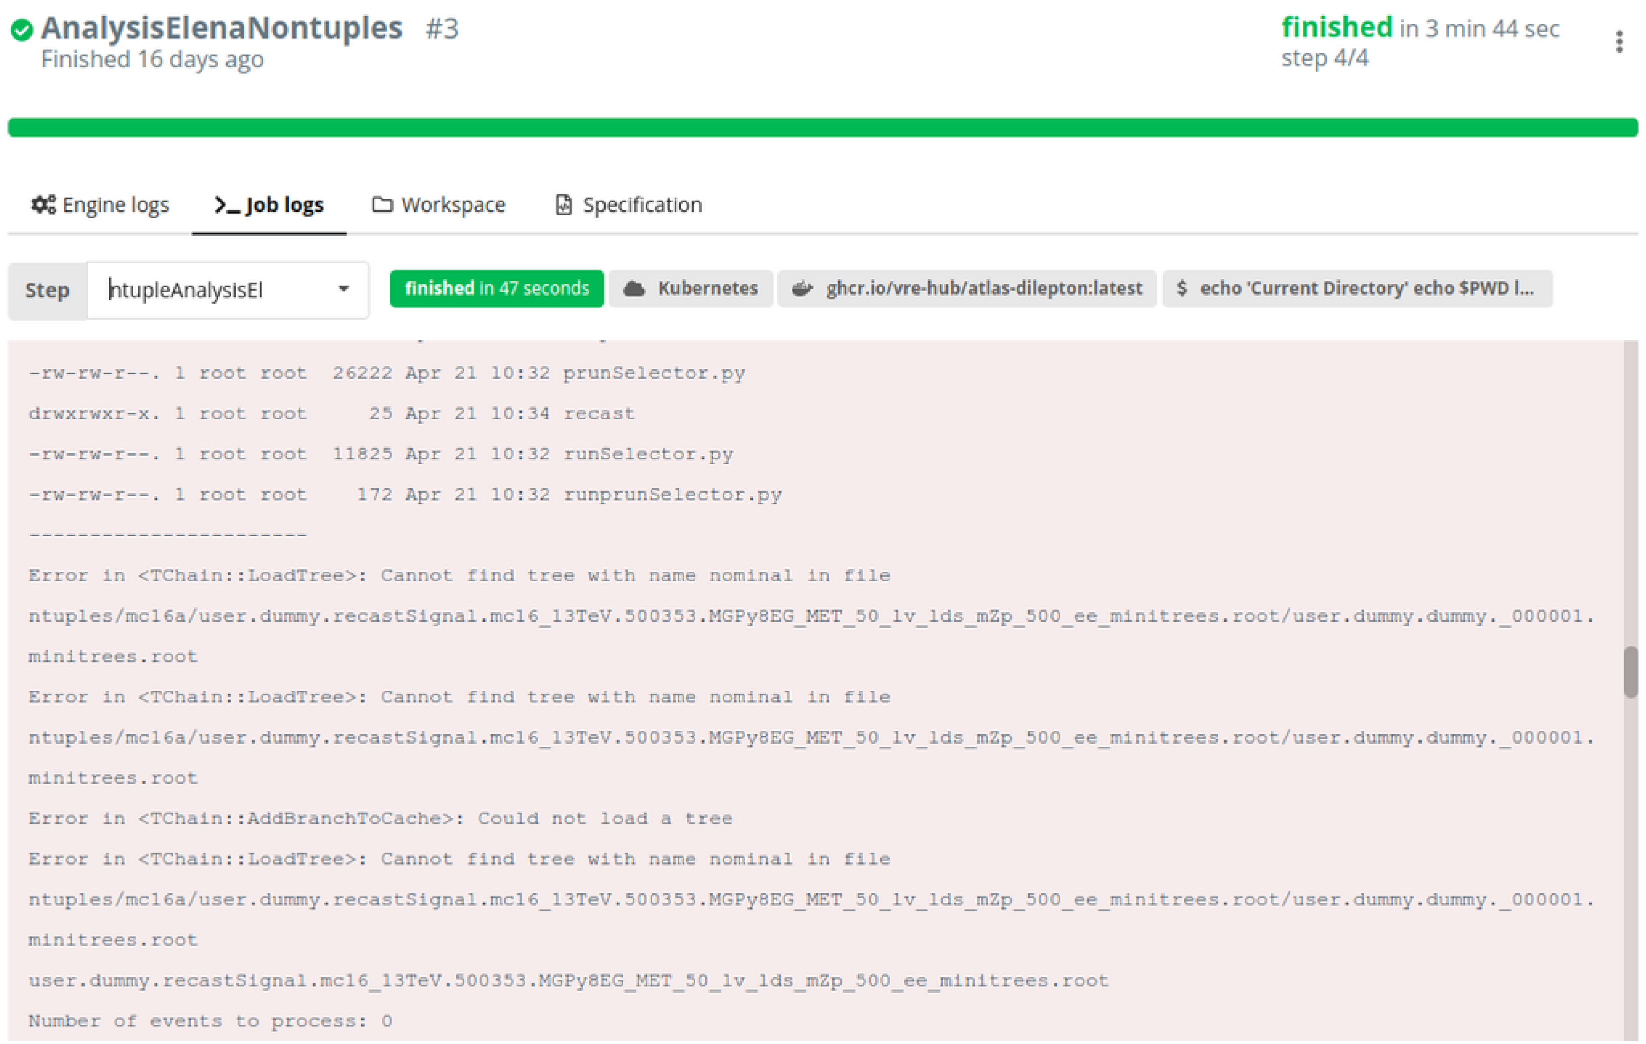

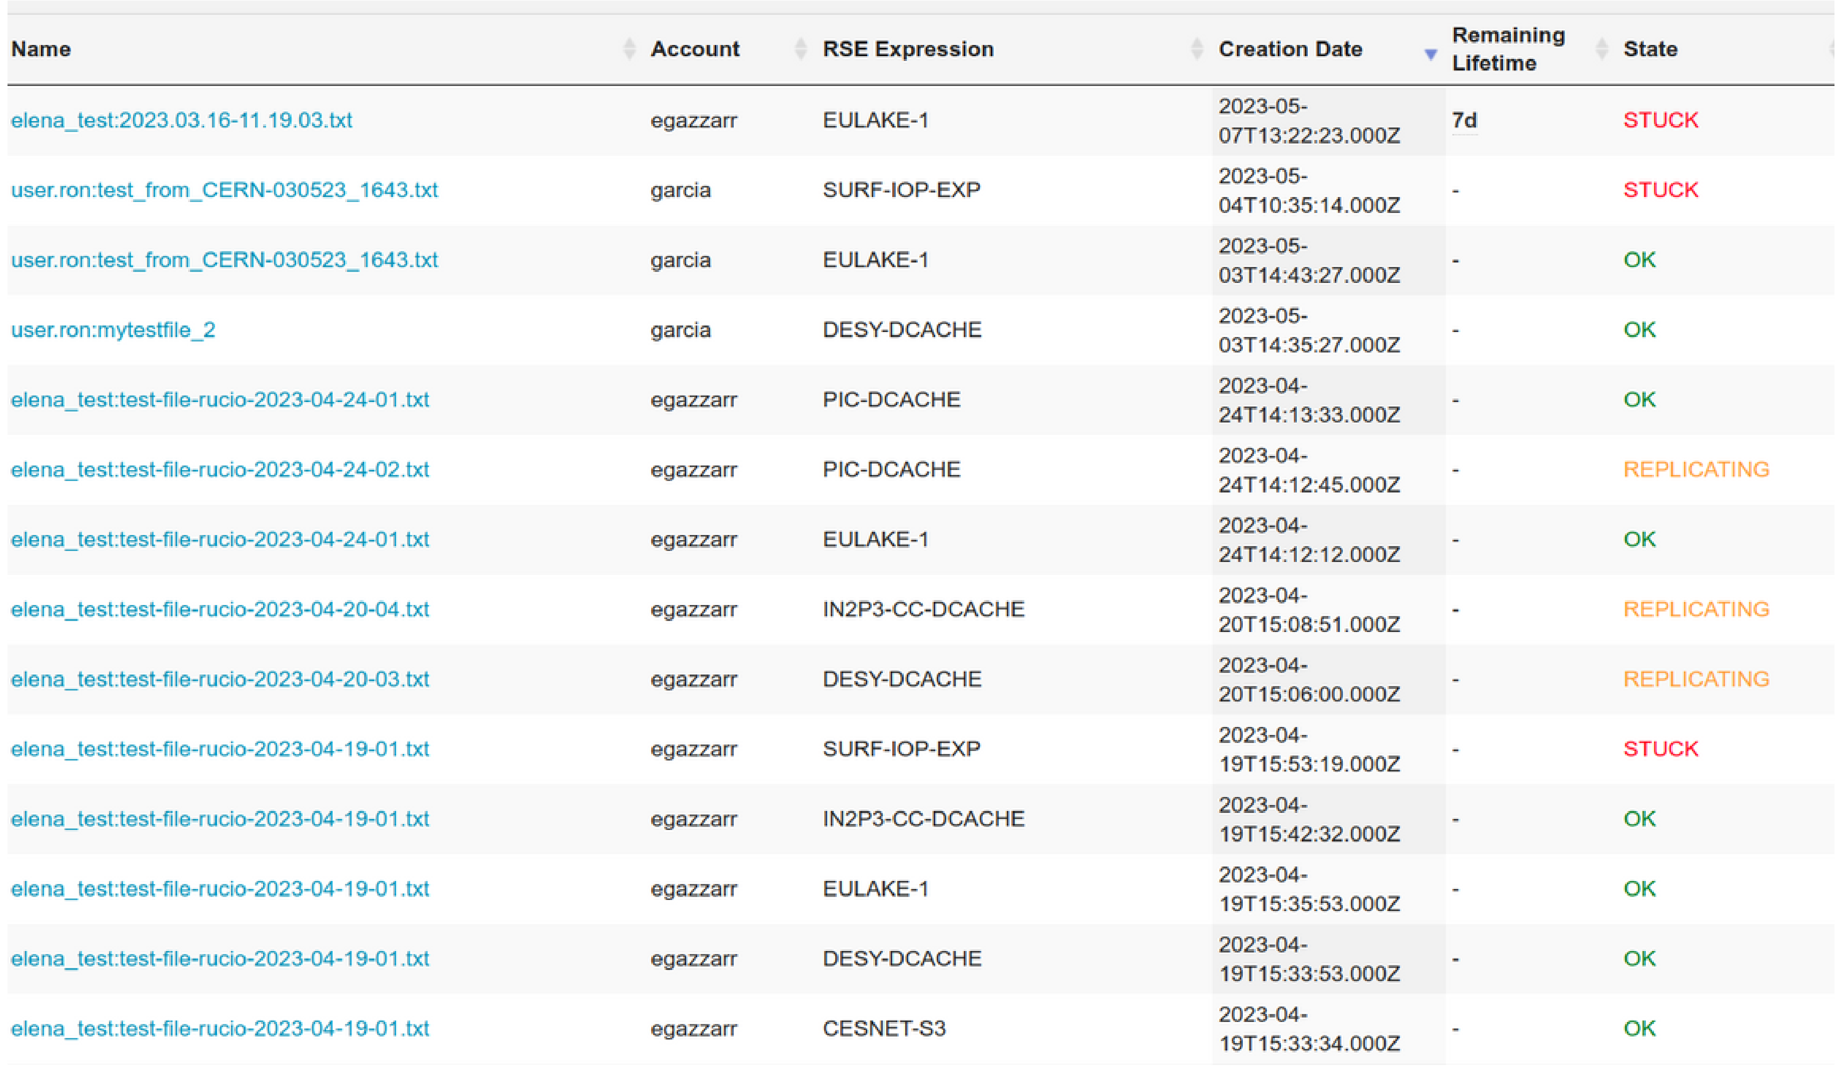

Name	Account	RSE Expression	Creation Date	Remaining Lifetime	State
elena_test:2023.03.16-11.19.03.txt	egazzarr	EULAKE-1	2023-05-07T13:22:23.000Z	7d	STUCK
user.ron:test_from_CERN-030523_1643.txt	garcia	SURF-IOP-EXP	2023-05-04T10:35:14.000Z	-	STUCK
user.ron:test_from_CERN-030523_1643.txt	garcia	EULAKE-1	2023-05-03T14:43:27.000Z	-	OK
user.ron:mytestfile_2	garcia	DESY-DCACHE	2023-05-03T14:35:27.000Z	-	OK
elena_test:test-file-rucio-2023-04-24-01.txt	egazzarr	PIC-DCACHE	2023-04-24T14:13:33.000Z	-	OK
elena_test:test-file-rucio-2023-04-24-02.txt	egazzarr	PIC-DCACHE	2023-04-24T14:12:45.000Z	-	REPLICATING
elena_test:test-file-rucio-2023-04-24-01.txt	egazzarr	EULAKE-1	2023-04-24T14:12:12.000Z	-	OK
elena_test:test-file-rucio-2023-04-20-04.txt	egazzarr	IN2P3-CC-DCACHE	2023-04-20T15:08:51.000Z	-	REPLICATING
elena_test:test-file-rucio-2023-04-20-03.txt	egazzarr	DESY-DCACHE	2023-04-20T15:06:00.000Z	-	REPLICATING
elena_test:test-file-rucio-2023-04-19-01.txt	egazzarr	SURF-IOP-EXP	2023-04-19T15:53:19.000Z	-	STUCK
elena_test:test-file-rucio-2023-04-19-01.txt	egazzarr	IN2P3-CC-DCACHE	2023-04-19T15:42:32.000Z	-	OK
elena_test:test-file-rucio-2023-04-19-01.txt	egazzarr	EULAKE-1	2023-04-19T15:35:53.000Z	-	OK
elena_test:test-file-rucio-2023-04-19-01.txt	egazzarr	DESY-DCACHE	2023-04-19T15:33:53.000Z	-	OK
elena_test:test-file-rucio-2023-04-19-01.txt	egazzarr	CESNET-S3	2023-04-19T15:33:34.000Z	-	OK

Monitoring, testing, dashboards, on-boarding

- Continuous **monitoring and testing** of transfers between Rucio Storage Elements (RSEs) is in place on Grafana dashboards hosted at CERN.
- **Rucio and Reana UI** interfaces deployed with K8s allow to explore and debug failed transfers and workflows.
- **Documentation** is hosted on Github pages and is made easy for both users and system administrators who would like to get inspired by the VRE model.

The VRE

A comprehensive analysis platform to serve the particle physics and astrophysics community.

View My GitHub Profile

The Virtual Research Environment

The Virtual Research Environment is an analysis platform developed at **CERN** serving the needs of scientific communities involved in European Projects. Its scope is to facilitate the development of **end-to-end physics workflows**, providing researchers with access to an **infrastructure** and to the digital content necessary to produce and preserve a scientific result in compliance with **FAIR** principles. The platform's development is aimed at demonstrating how sciences spanning from High Energy Physics to Astrophysics could benefit from the usage of common technologies, initially born to satisfy CERN's **exabyte-scale data** management needs.

The Virtual Research Environment's main components are:

1. A federated and reliable **Authentication and Authorization** layer
2. A **federated distributed storage** solution (the Data Lake), providing functionalities for data injection and replication through a Data Management framework (Rucio)
3. A **computing** cluster supplying the processing power to run full analyses with Reana, a re-analysis software
4. An enhanced **notebook interface** with containerised environments to hide the infrastructure's complexity from the user.

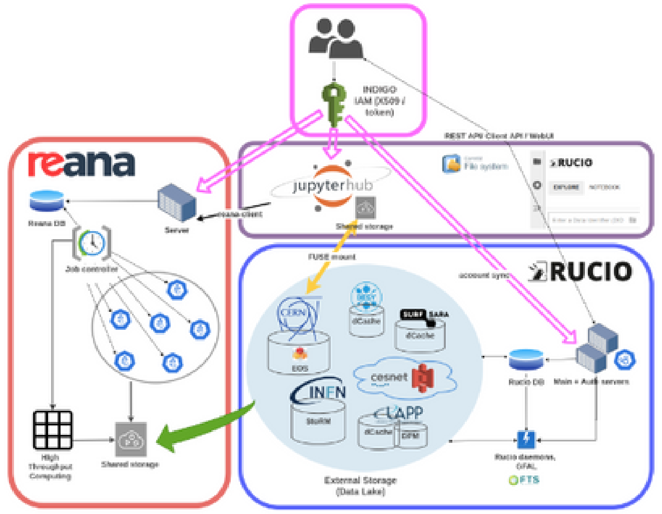

The deployment of the Virtual Research Environment is open-source and modular, in order to make it easily reproducible by partner institutions; it is publicly accessible and kept up to date by taking advantage of state of the art IT-infrastructure technologies.

The Science Projects which are using the VRE are described here.

If you are a scientist or a new user curious to use the above resources, please refer to the following documentation:

1. AAI
2. Rucio Data Lake
3. Reana cluster
4. Notebook service

Hosted on GitHub Pages — Theme by orderedlist

Deployment

VRE public Github repository hosts

- cloud deployment of the infrastructure components with Helm, Flux, Terraform and K8s
- Science Projects software to produce the environments for the Jupyterhub instance
- scientific code to be shared
- reana.yaml files to reproduce the analysis
- forums and discussions

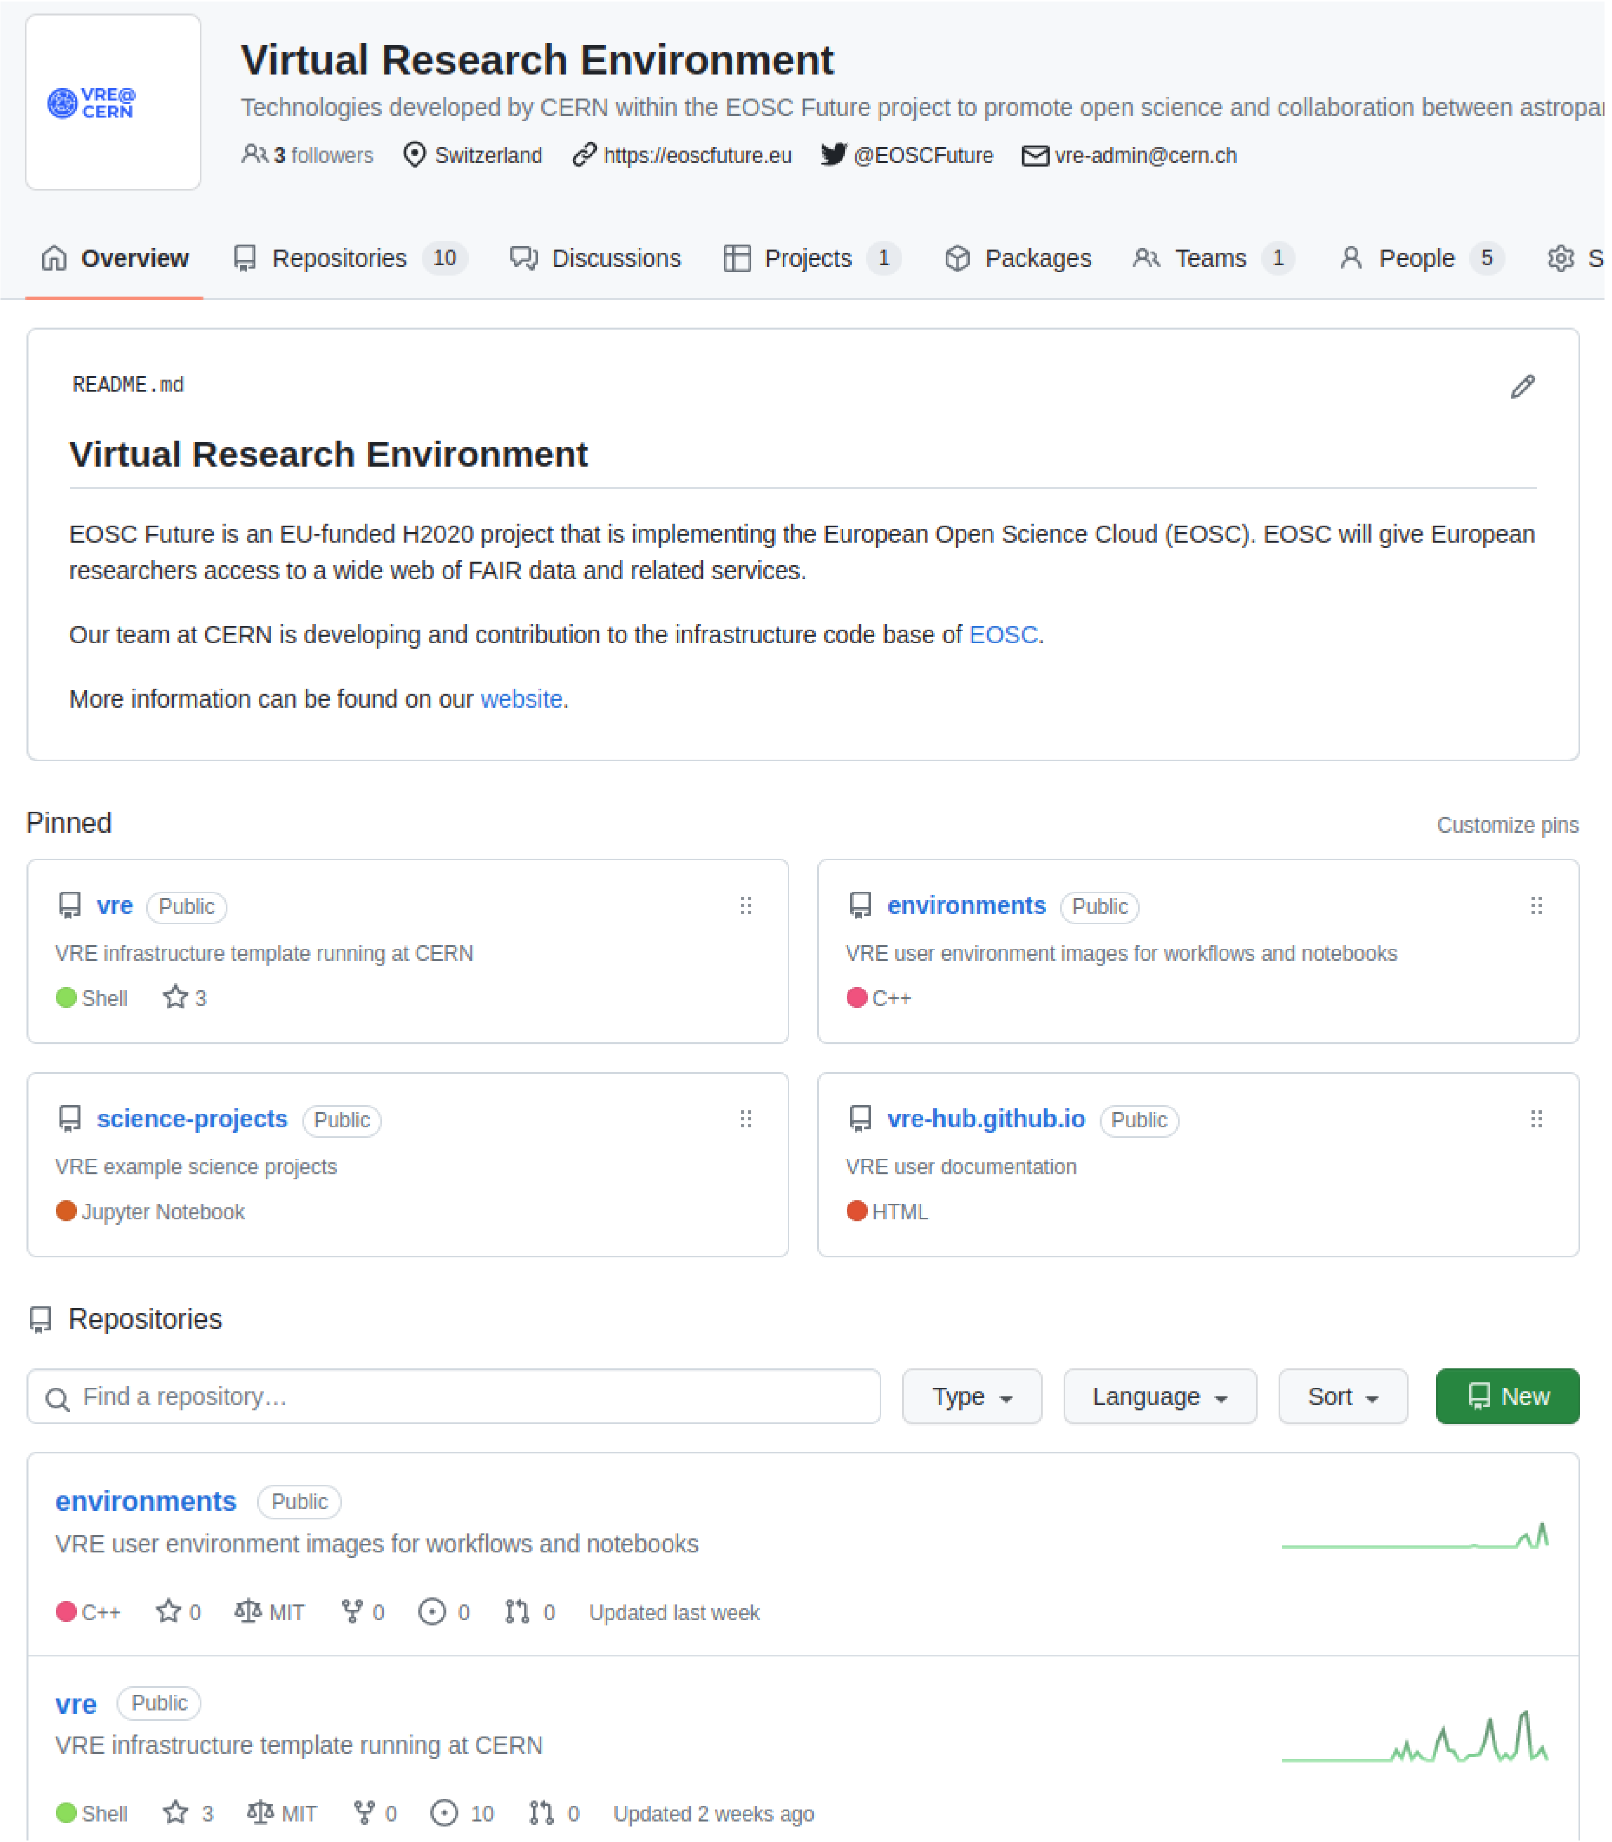

Demo

DM@LHC with ATLAS

1. **Dark Matter Reinterpretation**: setting limits on High-Luminosity LHC contraints on $Z' \rightarrow \chi\chi$ (Z' mediated Dark Matter models).
2. The **dilepton** inclusive search (right) concluded in 2019
   a. objective: projecting limits to 14 TeV and computing the fiducial cross-sections in **lower mass regions.**

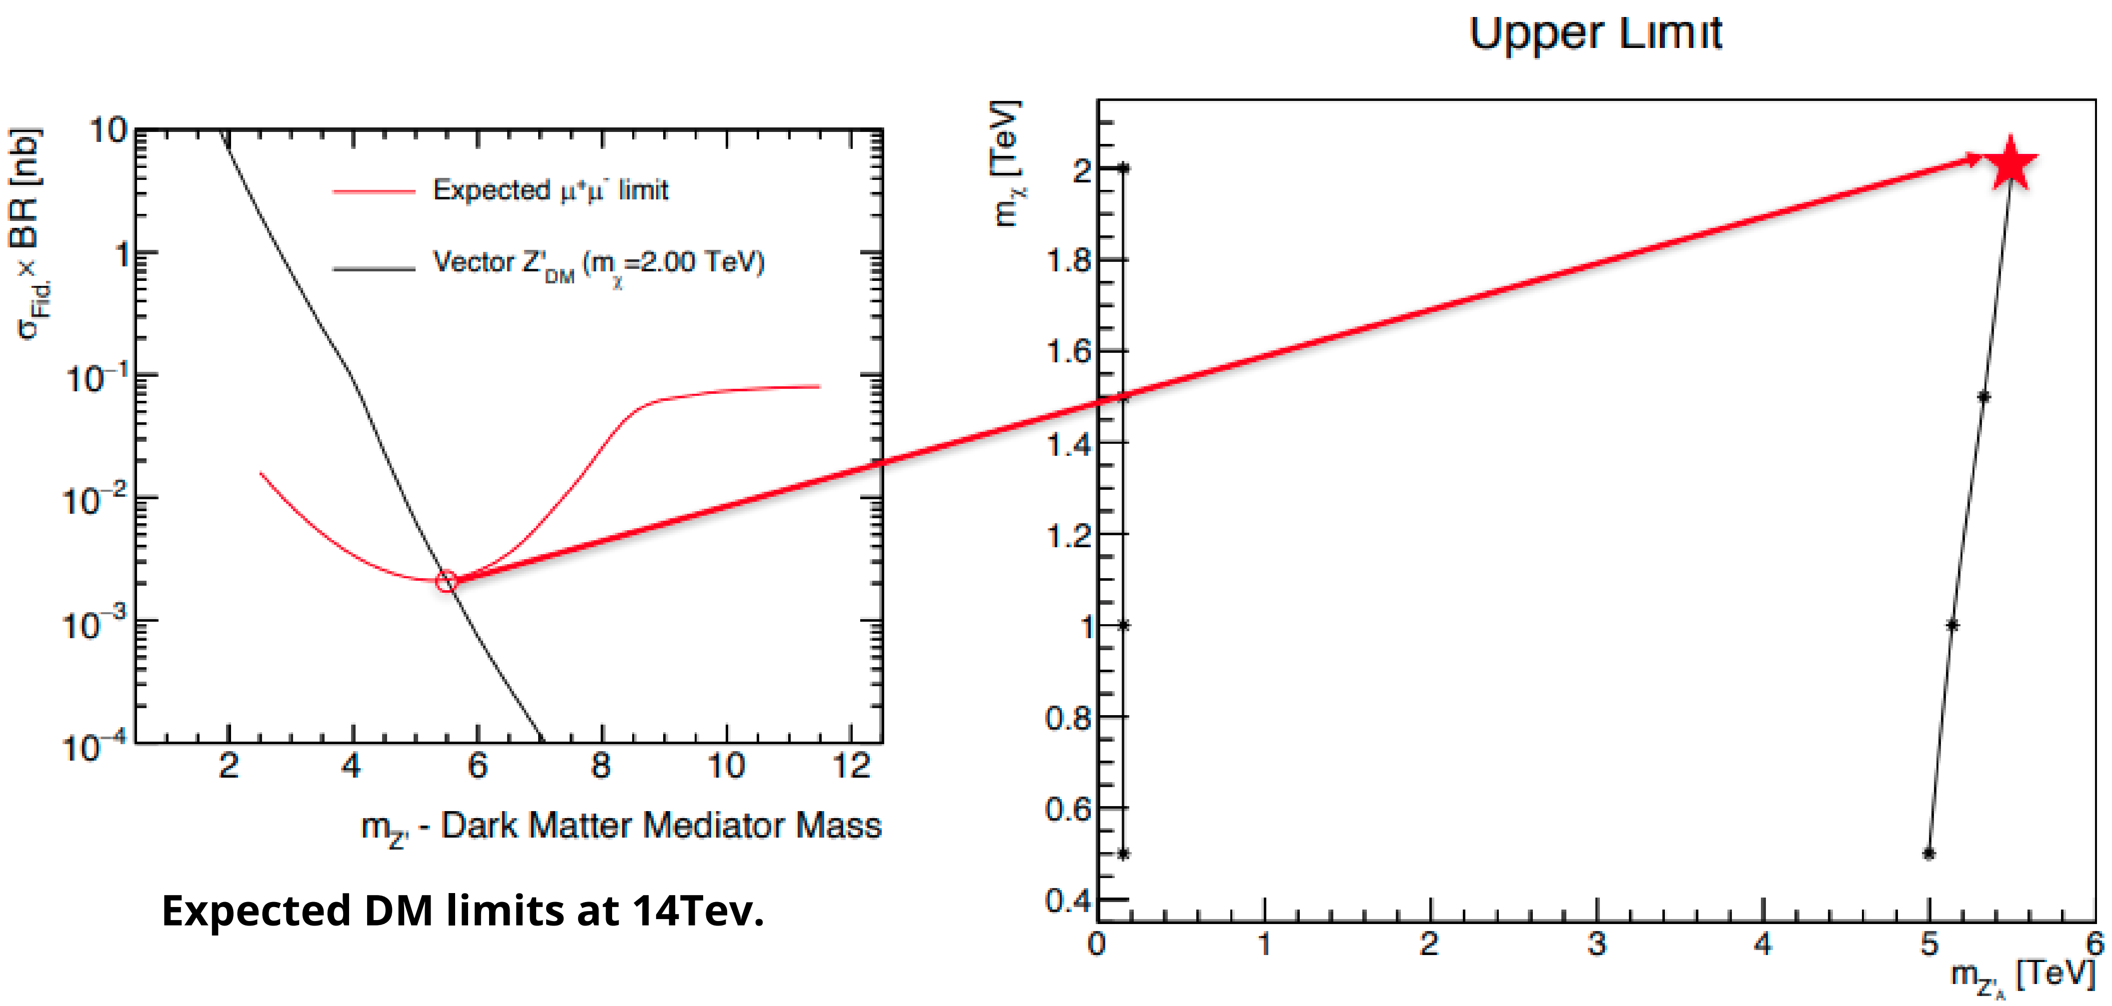

Expected DM limits at 14Tev.

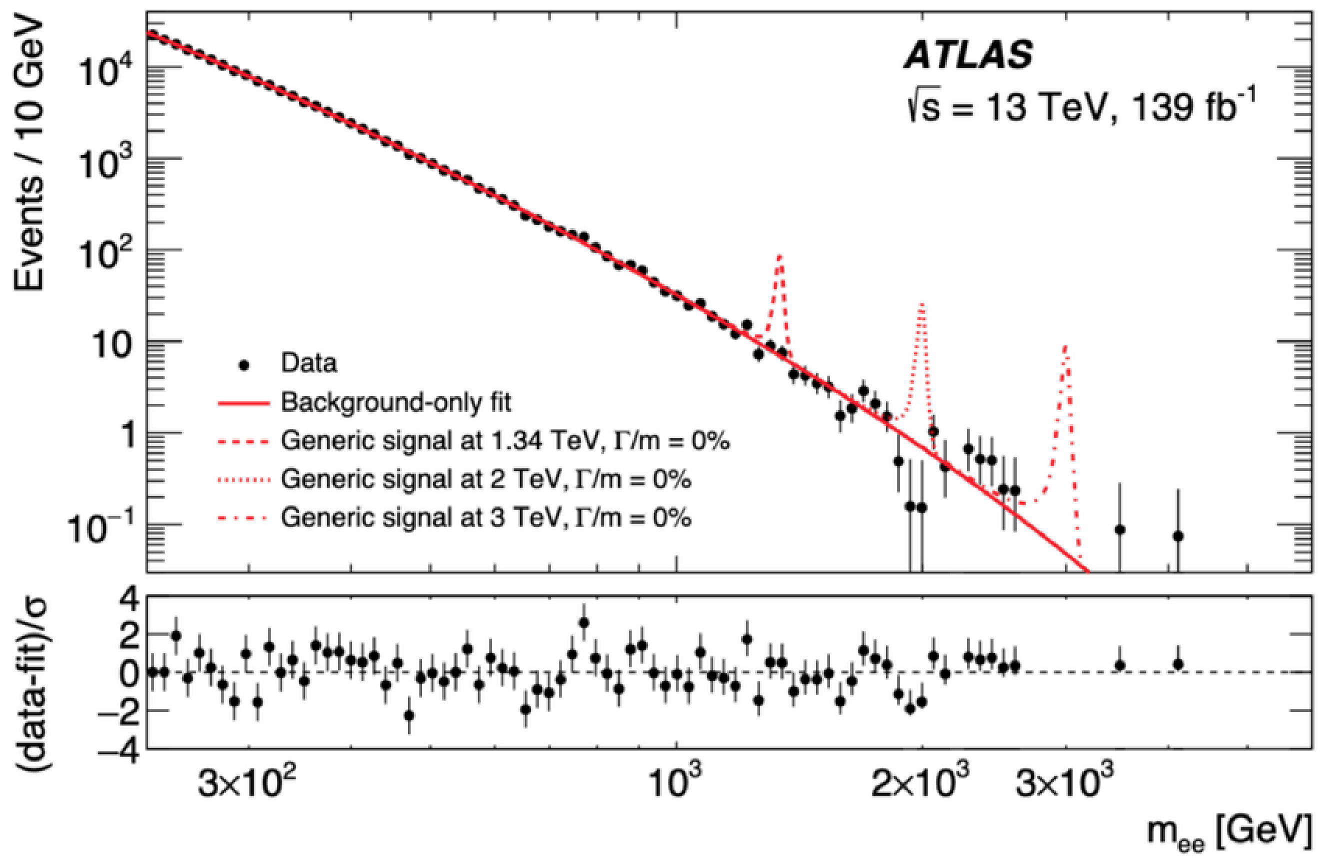

Dilepton Inclusive Search. Results of this analysis demonstrate good agreement with SM predictions.

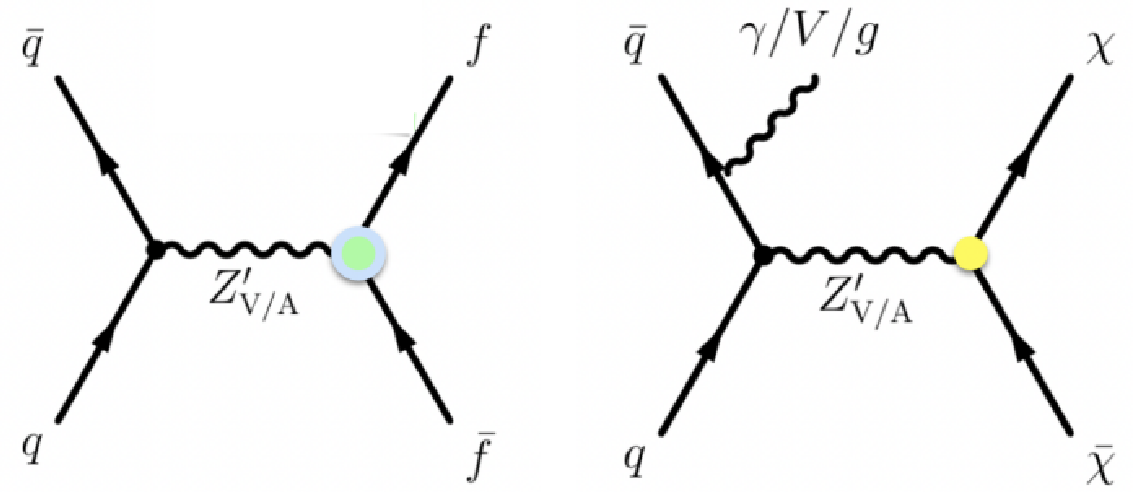

See yesterday's talk by J. Little

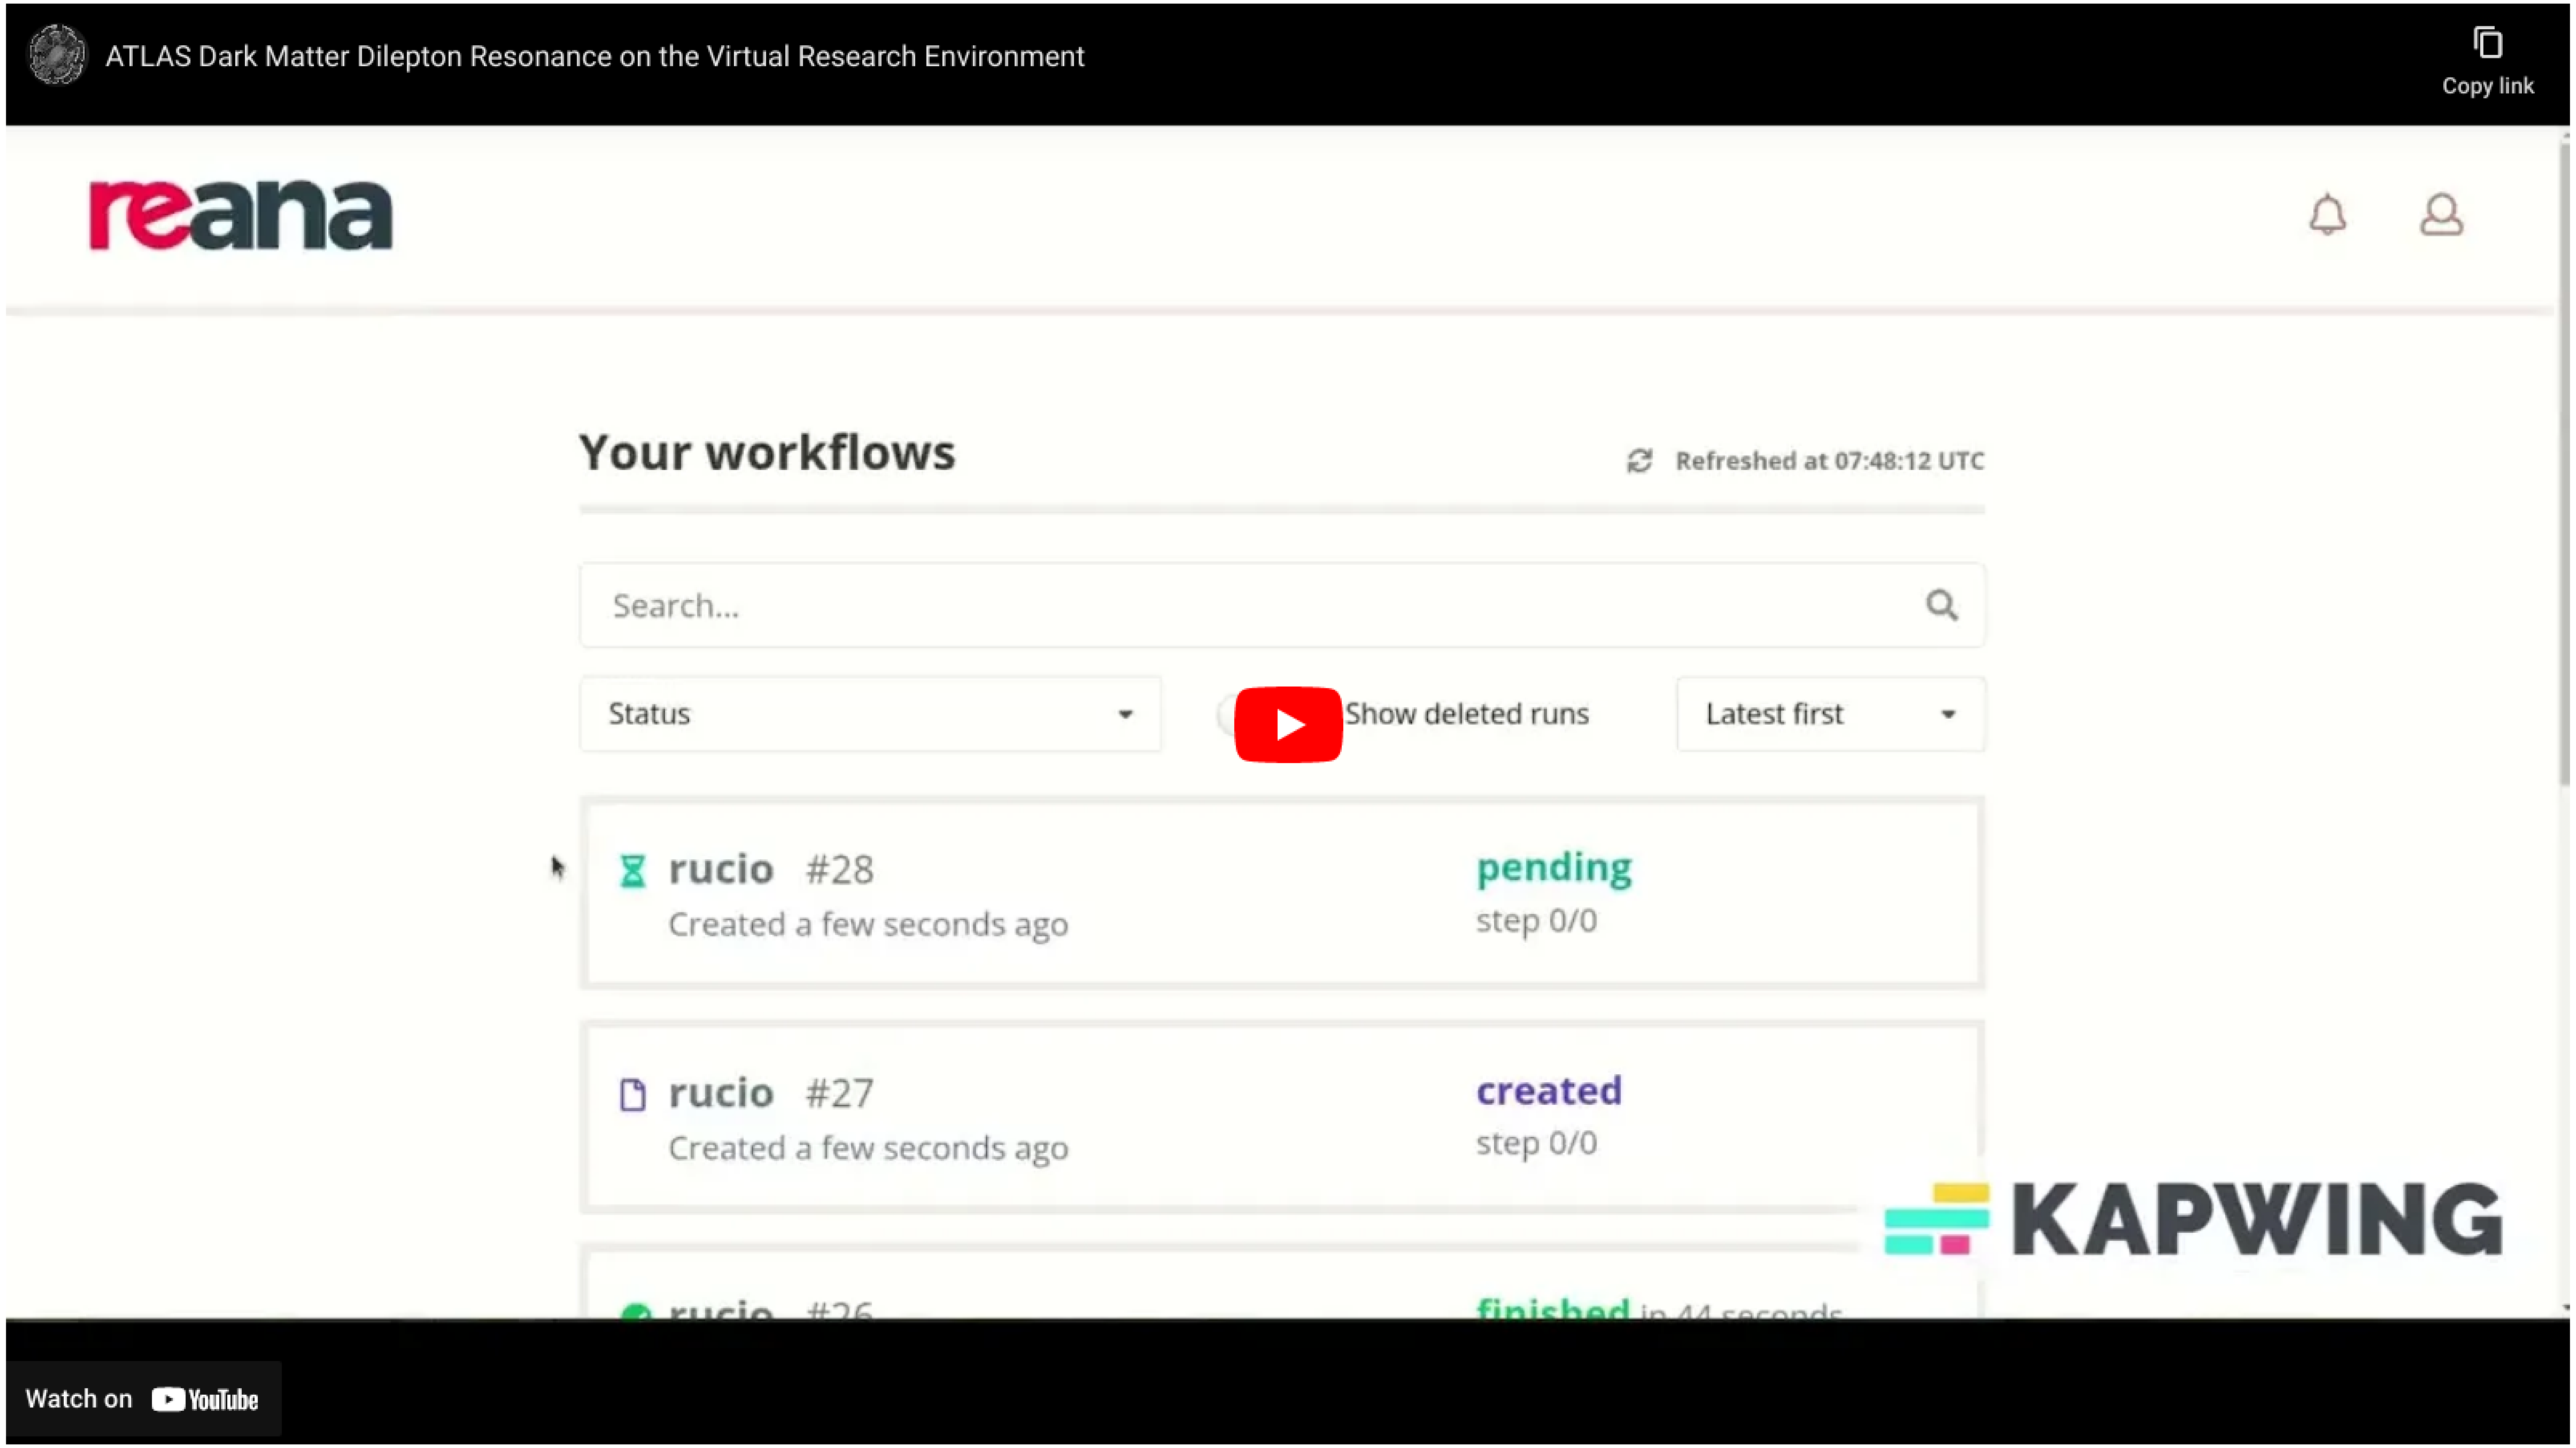

https://www.youtube.com/watch?v=hvlJLo_7xXc&ab_channel=ElenaGazzarrini

Behind the scenes…

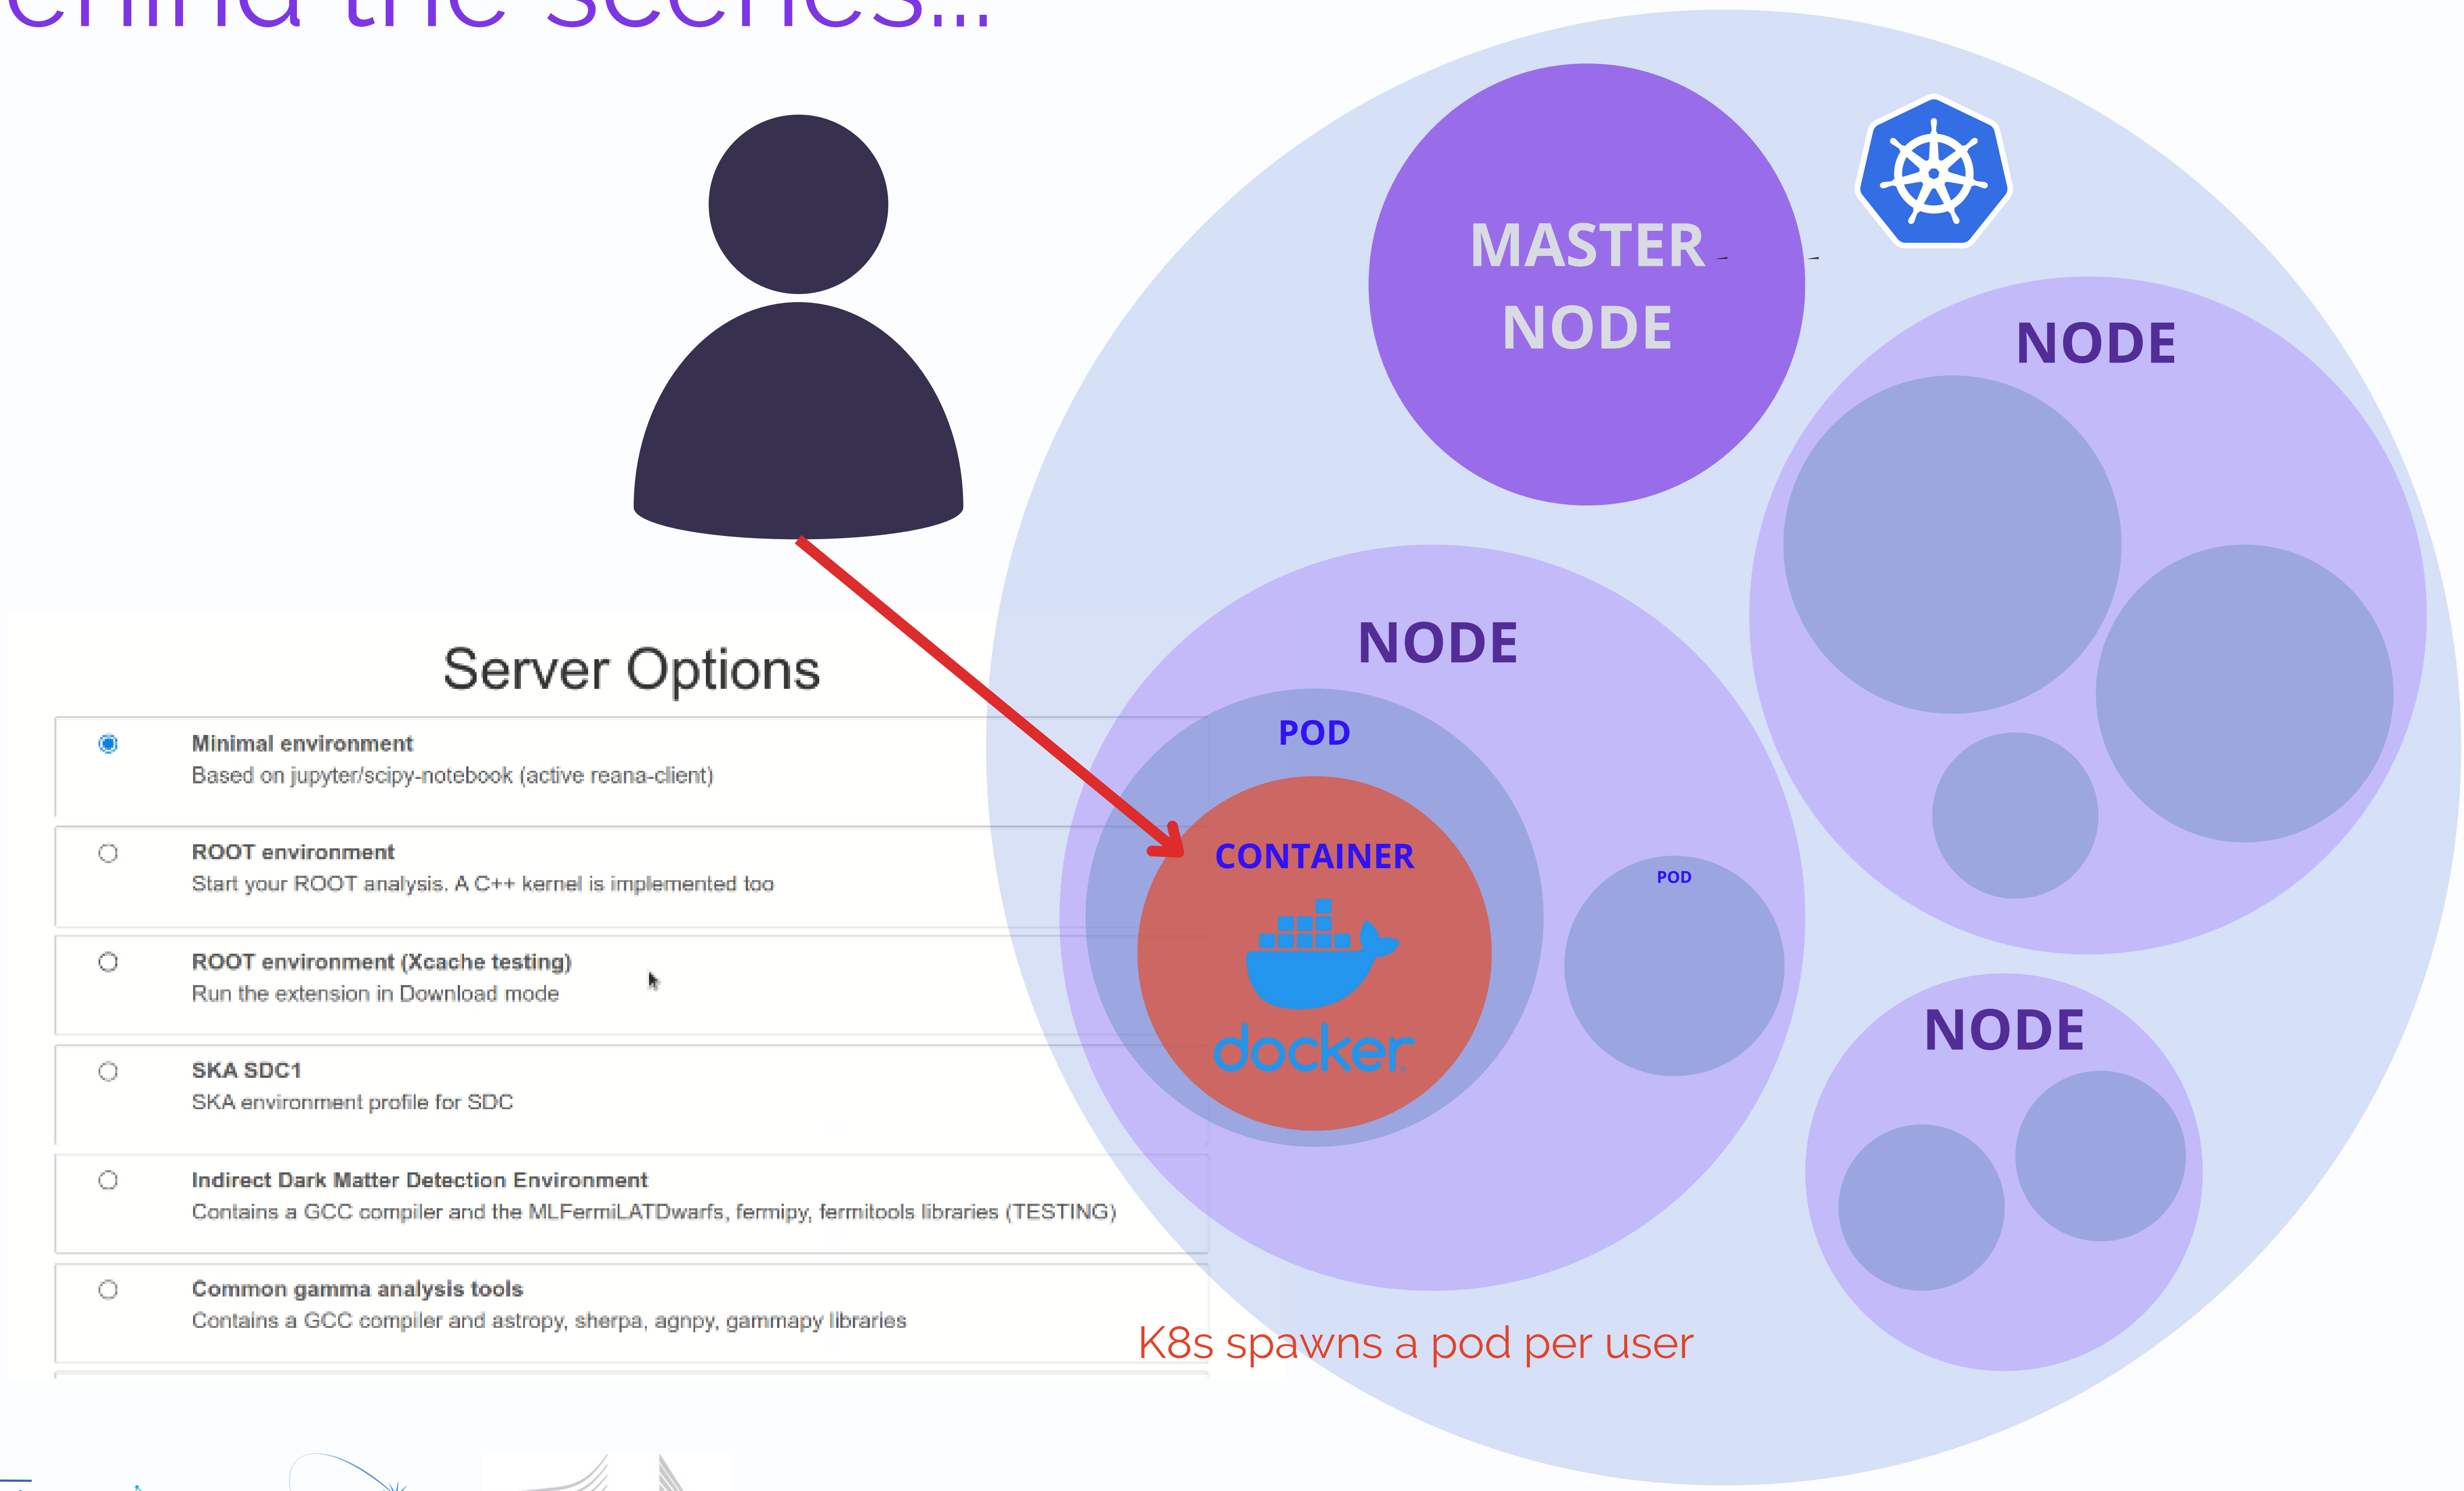

Behind the scenes…

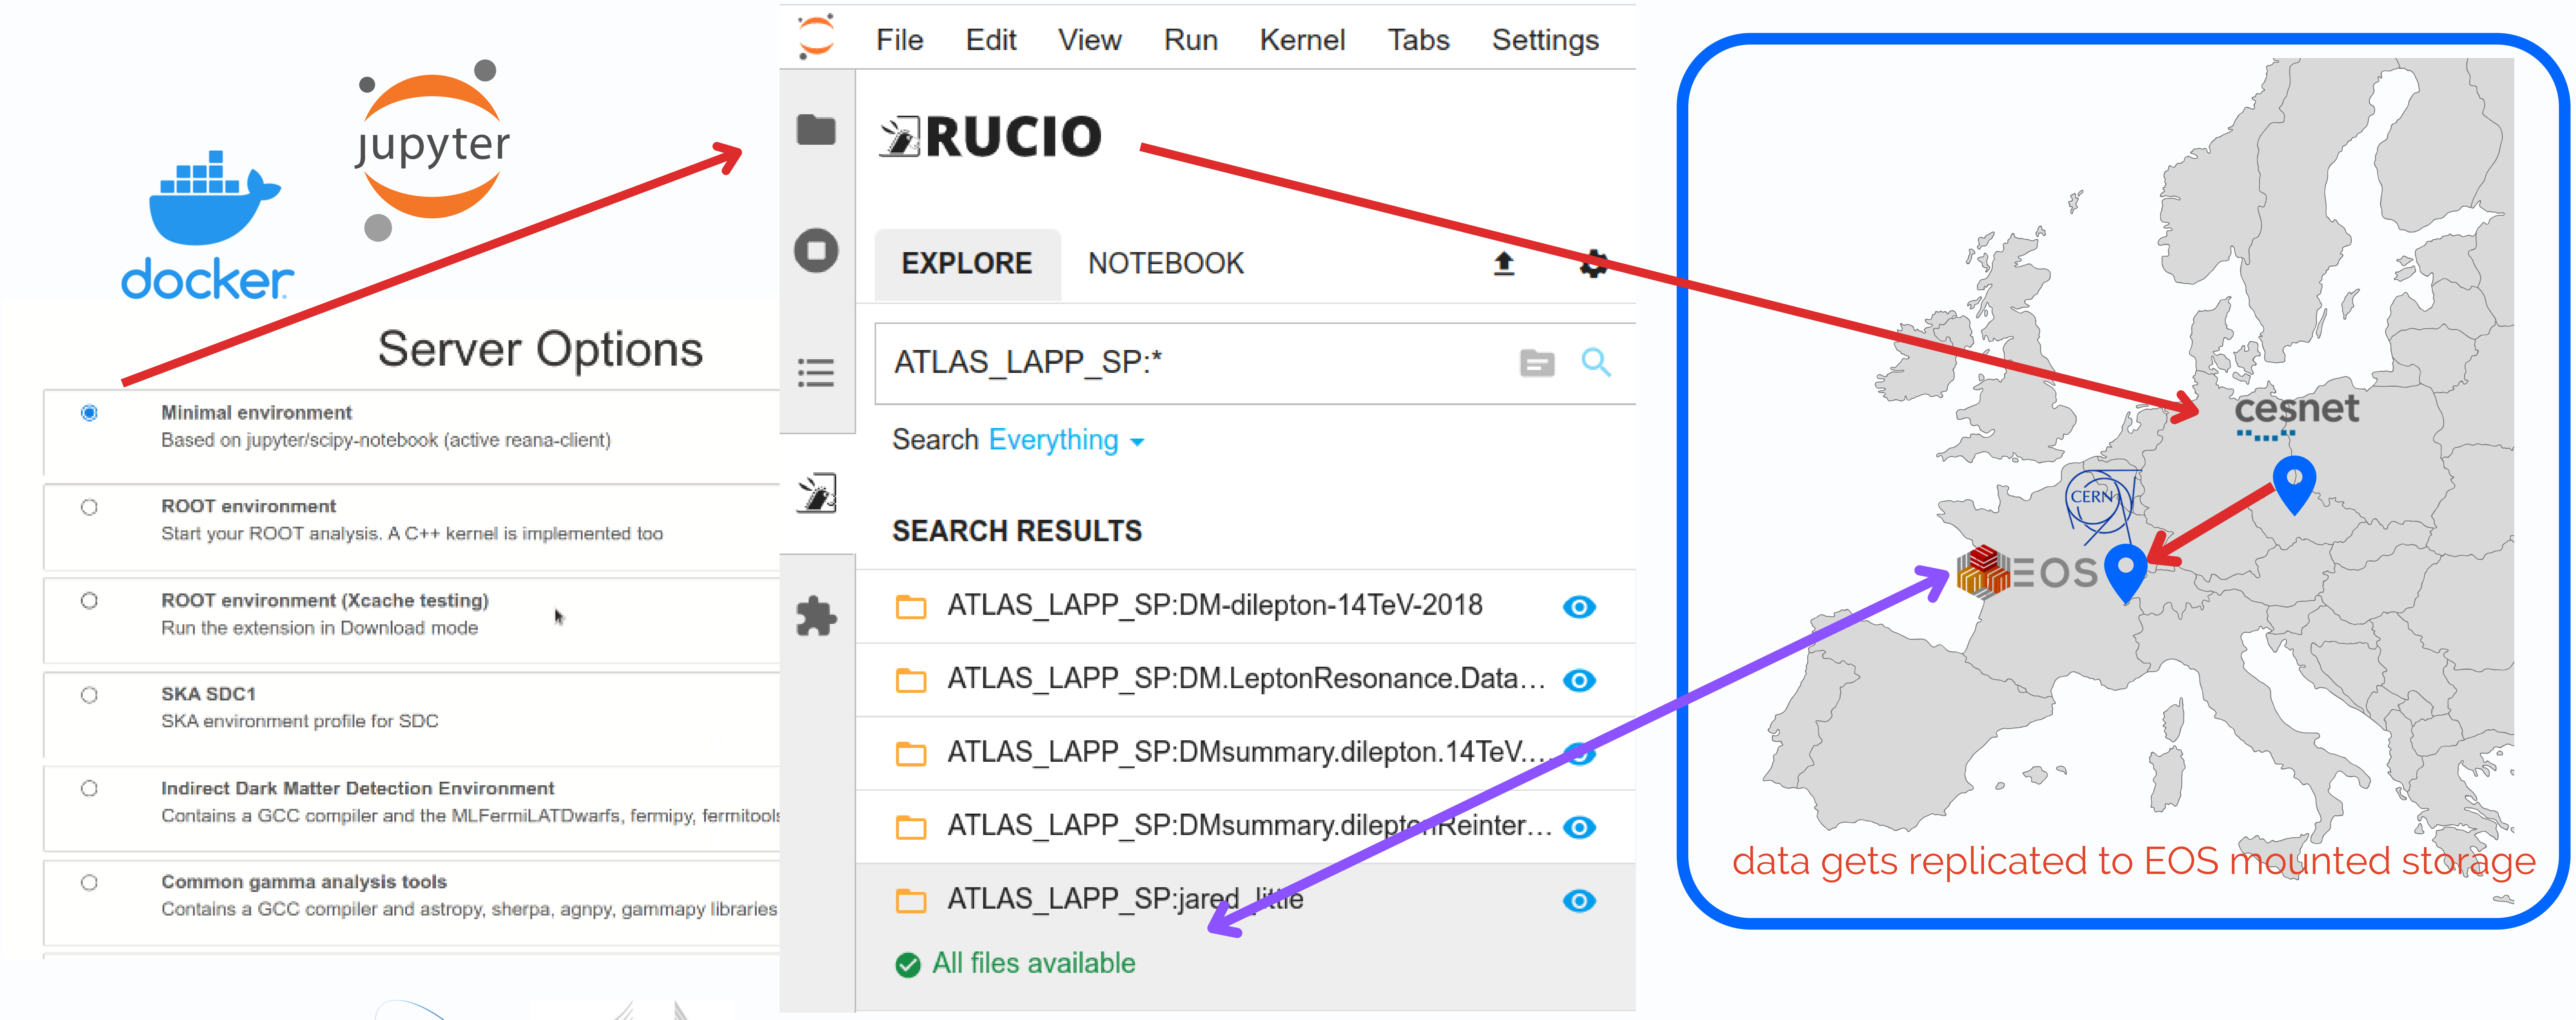

Behind the scenes...

reana.yaml

```yaml
version: 0.8.1
inputs:
  directories:
    - python/
workflow:
  type: serial
  specification:
    steps:
      - name: fetchdata-rucio
        voms_proxy: true
        rucio: true
        environment: 'projectescape/rucio-client'
        commands:
          - rucio whoami
          - rucio get ATLAS_LAPP_SP:DMsummary.dileptonReinterpretat
      - name: SetLimits
        environment: 'reanahub/reana-env-root6:6.18.04'
        compute_backend: kubernetes
        kubernetes_memory_limit: '9Gi'
        commands:
          - mkdir plots/
          - python python/MakeLimit.py
outputs:
  directories:
    - plots/
```

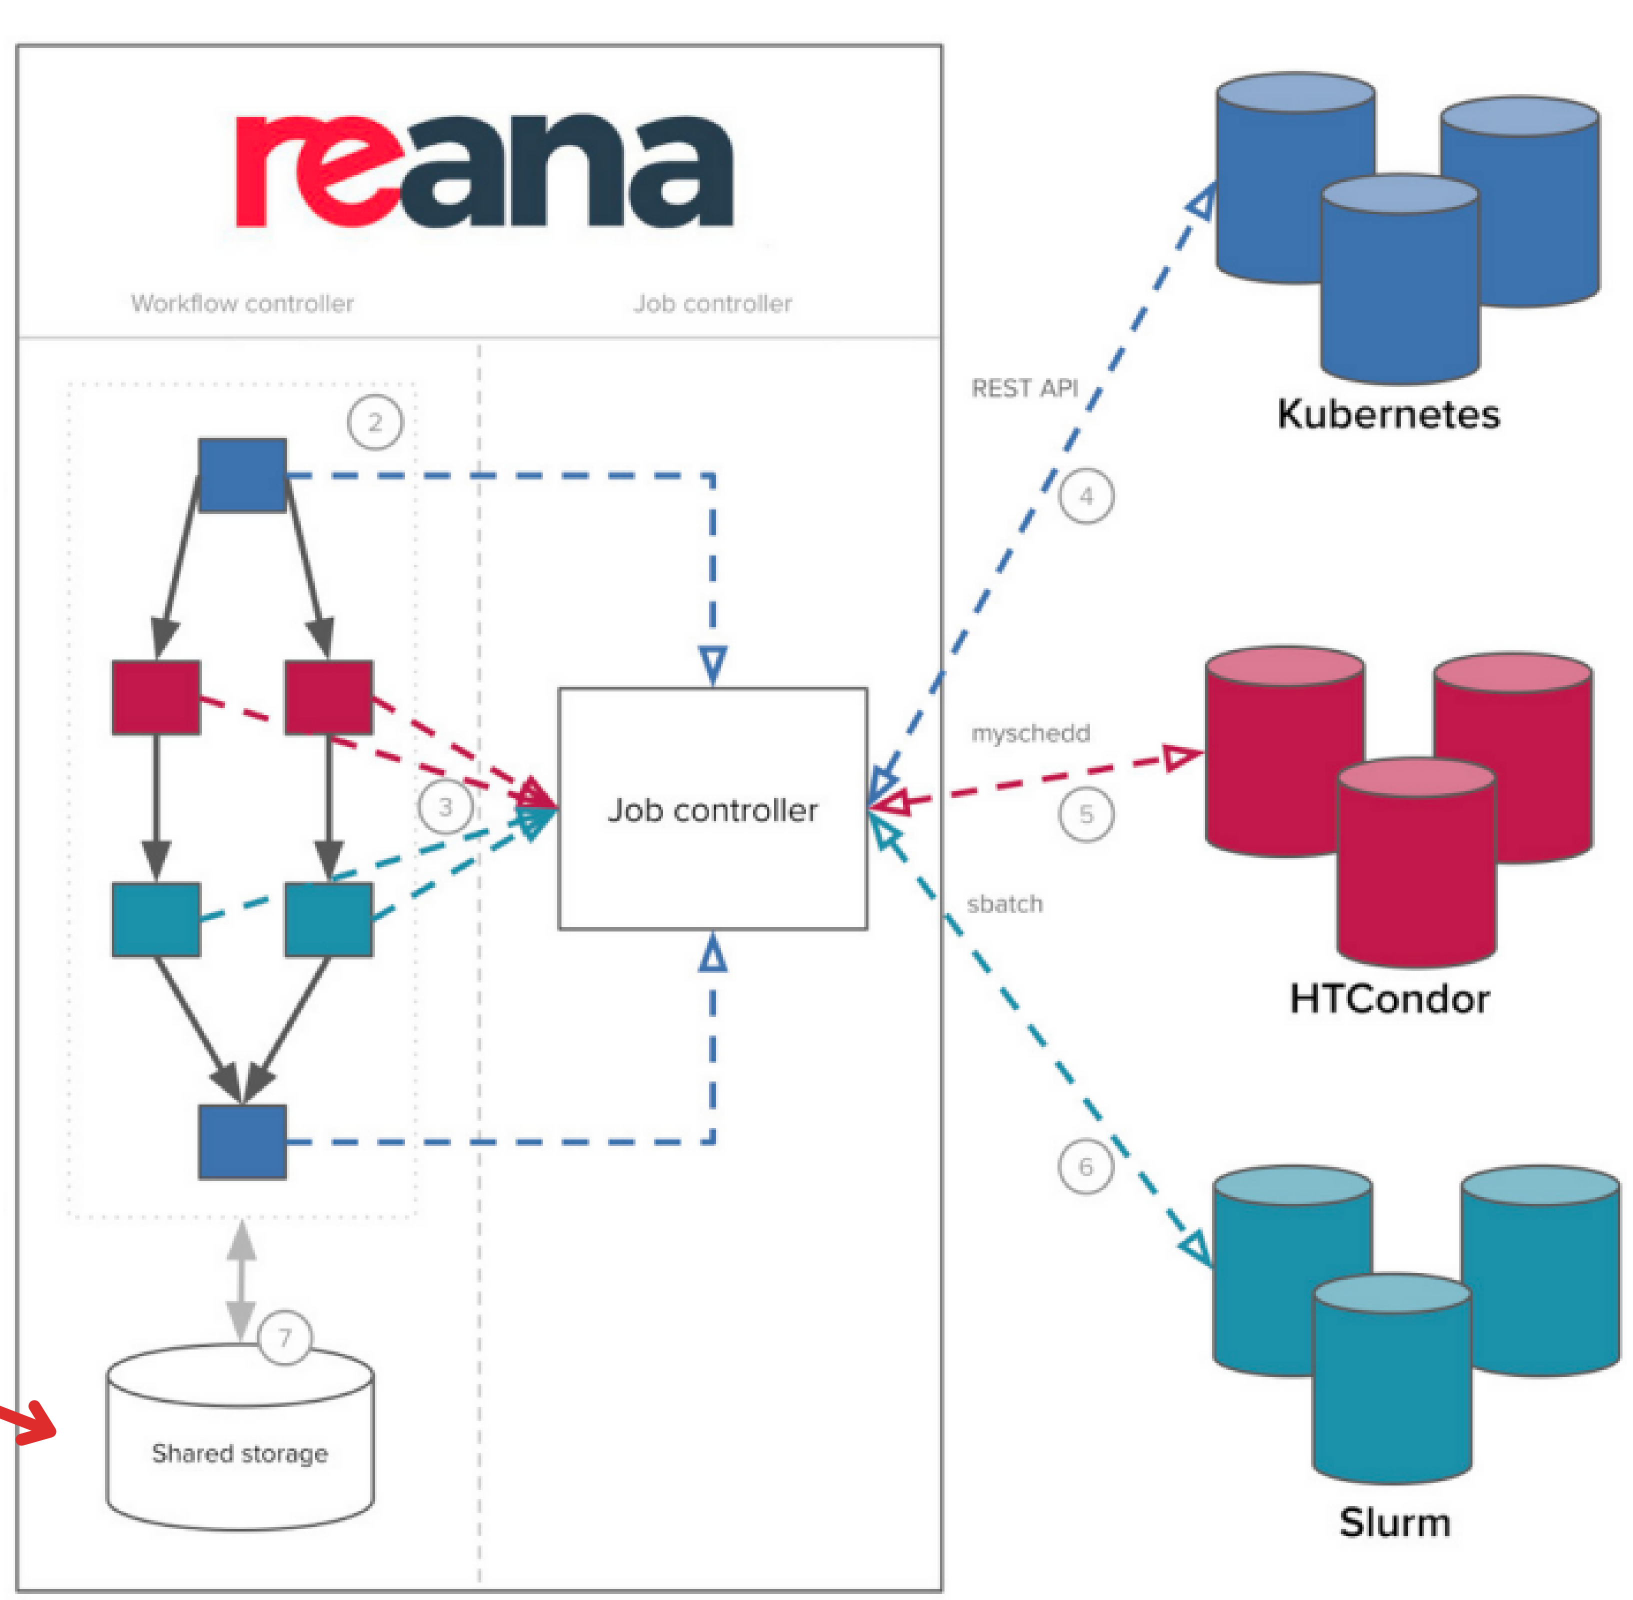

data gets downloaded on Reana storage

Future outlook

Success stories

- Escape Open Collaboration Agreement ensures the collaboration and joint common activities across scientific communities in the development of VREs
- VRE awoke interest from scientific domains who are in early-stage prototype phase
  - Einstein telescope (next generation gravitational waves detector)
  - NUCLEUS experiment (elastic neutrinos scattering)
  - VdR Würzburg - German centre for Data-Intensive Radio Astronomy
- Interest from new digital models (i.e. digital twins) developed within European projects

Future plans

- connection with HPCs, commercial clouds and other external computing resources
  - FENIX and the EuroHPC Joint Undertaking work (eg: FTS delivering files to Julich-HPC with S3 protocol)
- Use VRE as a performance evaluation between workload managers
- Caching data on distributed storage

The VRE is ...

- **modular**
  - integrates software, tools and packages
  - can be configured to connect to remote storage and computing resources
- **flexible**
  - ad-hoc workflows can be created via easily editable declarative files
  - can be installed on different machines
  - independent of CERN restrictions
- **reproducible/sustainable**
  - deployment is kept simple and documented to be used as a blueprint for other research infrastructures
  - allows analysis preservation
- useful for **sharing and collaboration**
  - common entry point with same authentication for all services

Thank you

special thanks to

Tibor Simko & Reana team, Martin Barisits & Rucio team, Xavier Espinal, Ian Bird, CNAF IAM team and all the Science Projects researchers (Jared Little, Caterina Doglioni, Christopher Eckner, Alexander Ekman, Axel Gallen, Mikhail Smirnov, Francesca Calore, Pooja Bhattacharjee, Valerio Ippolito, Estelle Pons, Elena Cuoco, Alberto Iess, Alessandro Parisi, Dany Vohl)

e-mail

elena.gazzarrini@cern.ch

Back up

Status

The VRE is an R&D project and it is not a production system. As such, the platform is maintained by a team of 3 people.

For the moment, ~ 230 users subscribed on the IAM platform and have therefore access to the resources.

VRE documentation and links to resources at: https://vre-hub.github.io/.

Links to useful related works are provided by clicking on the underlined text in the slides.

vCPUs	RAM (GB)	Masters	Nodes	Remote Storage (TB)	CephFS (TB)
184	335.8	3	23	646	1.8

25 Openstack machines

- 14.6GB RAM
- 8 VCPU
- 80GB Disk
- Fedora CoreOS 35
- LINUX

Two sides of the coin

A bipartite look at the ideal infrastructure ...

	SCIENTIST	IT ADMINISTRATOR
USEABILITY	Ergonomic (onboarding, documentation)	Maintenance, portability, modularity
DATA ACCESS	Various FAIR data/metadata types	Security, varied protocols and technologies
ANALYSIS	Performance	Cost, energy consumption
REPRODUCIBILITY / SUSTAINABILITY	Software and analysis steps preservation	Easy re - deployment