EOSC - Future Dark Matter Science Project

Elena Gazzarrini, Jared Little

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

Agenda

Introduction

Dark Matter

Virtual Research Environment

Science Projects

Demo

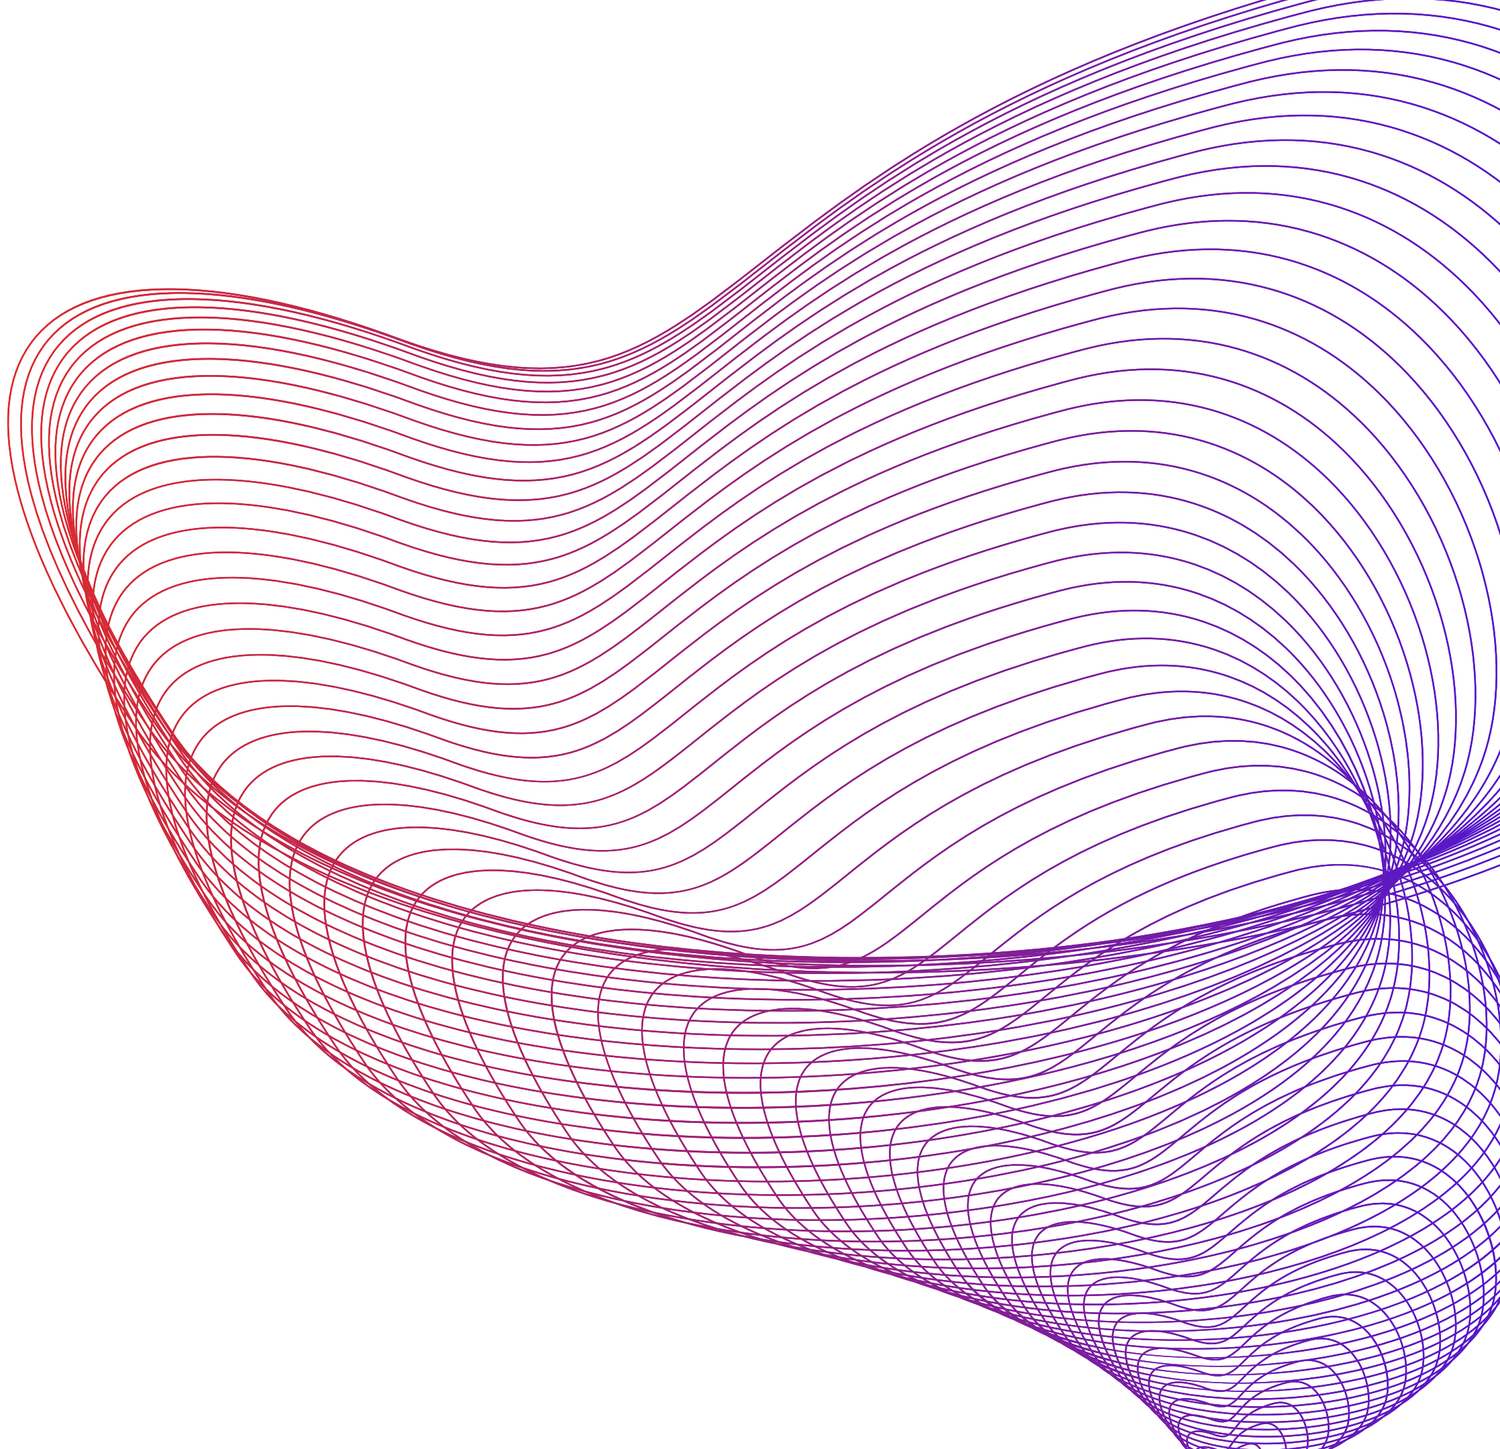

Introduction

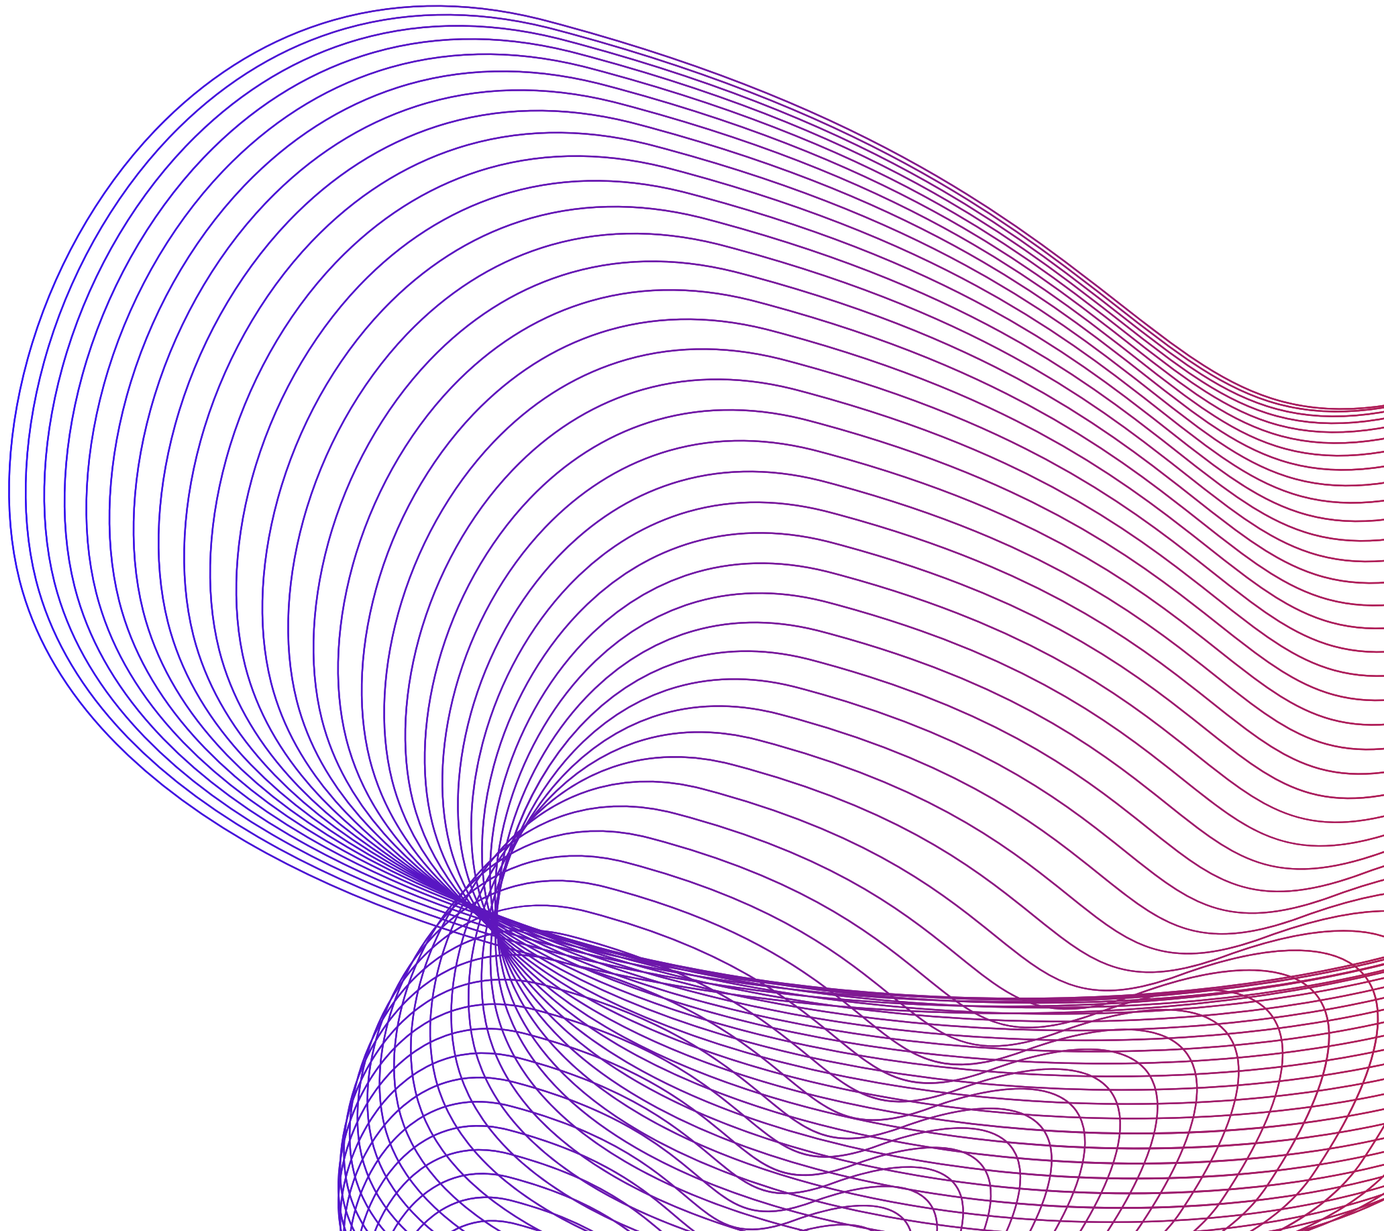

Data volumes are growing …

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.

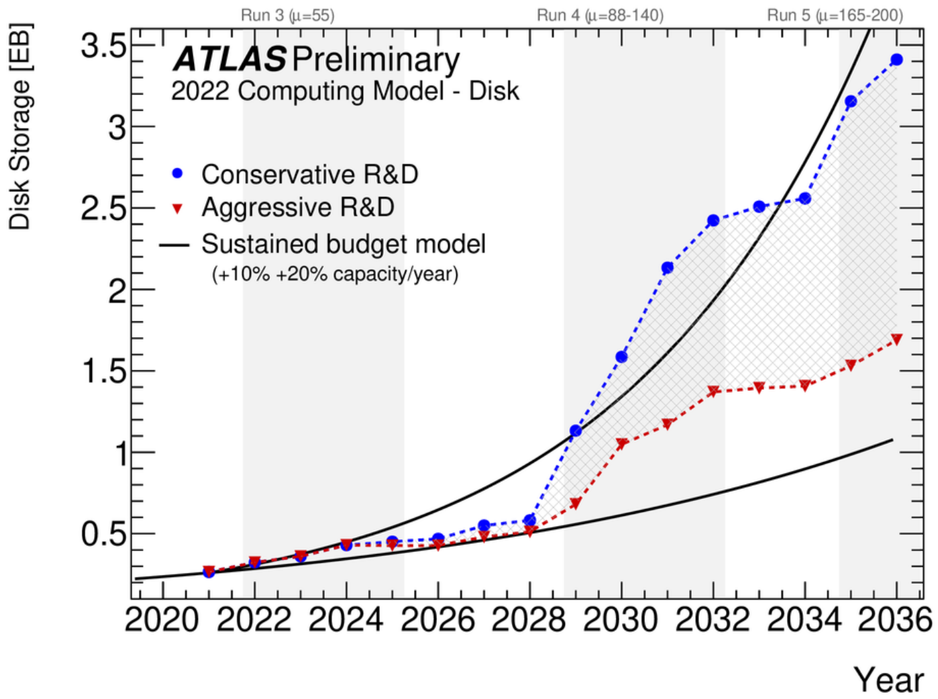

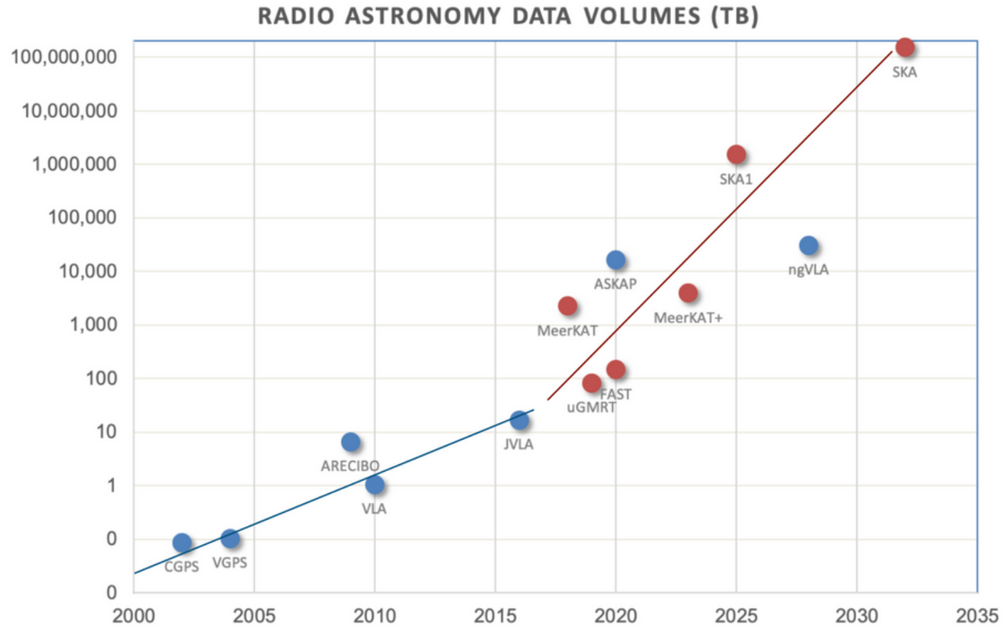

Taylor R. et al. Big Data Research Infrastructure Collaboration Toward the SKA (BRICSKA). doi: 10.1590/0001-3765202120201027. PMID: 34076205.

The challenge

A common infrastructure across Research Infrastructures would foster:

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

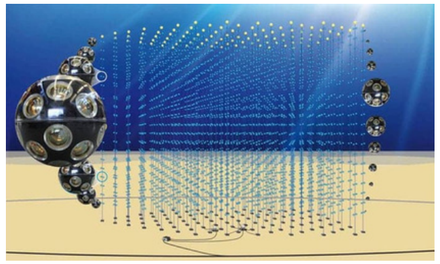
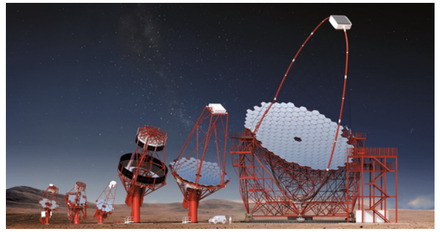
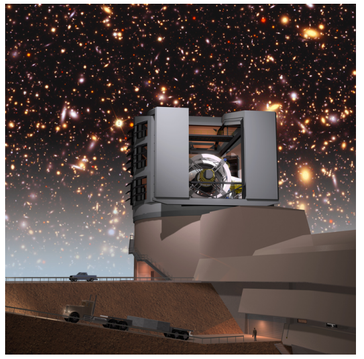


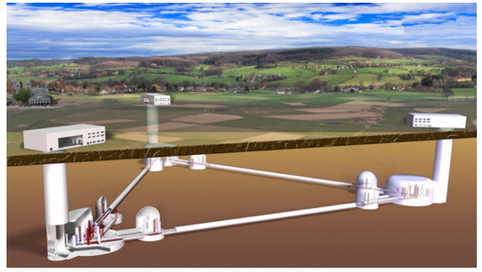
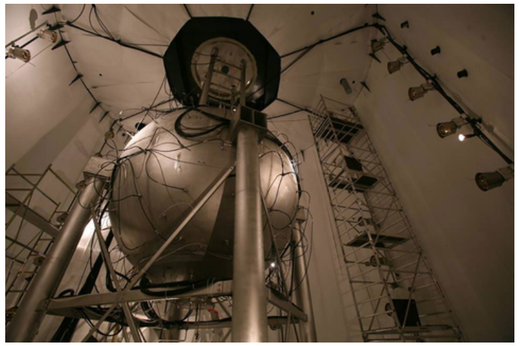

EOSC - Future Dark Matter Science Project Webinar - 6.6.2023 - E.Gazzarrini, J. little

Have you failed to reproduce a result?

70%

researchers tried and failed to reproduce others' results

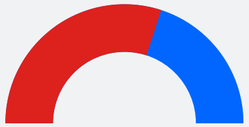

> 50%

researchers failed to reproduce own results

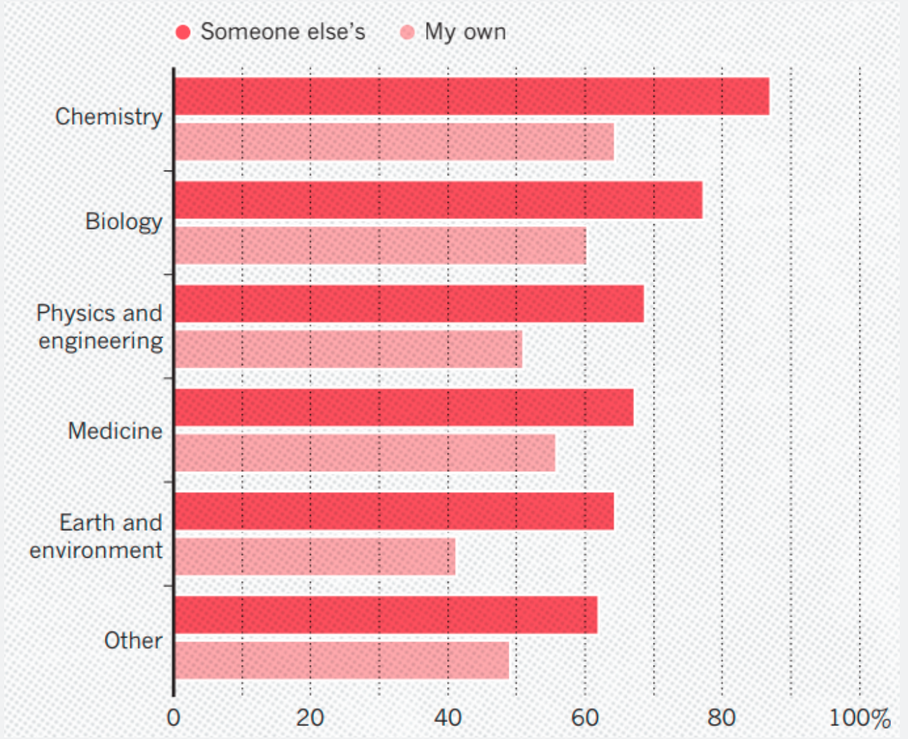

Baker, M. 1,500 scientists lift the lid on reproducibility. NatBure 533, 452–454 (2016).

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!

ESCAPE – Open Science Projects

- ESCAPE is an EU-funded project which aims to bring together different research infrastructures
  - **10 ESFRI** (CTA, EST, FAIR, HL-LHC, KM3NeT, SKA, LSST, VIRGO,ESO,JIVE)
  - **2 pan-European International Organisations** (CERN, ESO)
  - **4 supporting European consortia** (APPEC,ASTRONET,ECFA, NuPECC)
- ESCAPE services will contribute to the European Open Science Cloud (EOSC) through the **EOSC-Future** project
- 2 Science Projects produce cutting edge results: **Dark Matter** and Extreme Universe
  - Demonstrate open science capabilities
  - Improve productivity of researchers
  - Gain new insights and innovation across disciplines
  - Bring together the services implemented within ESCAPE

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

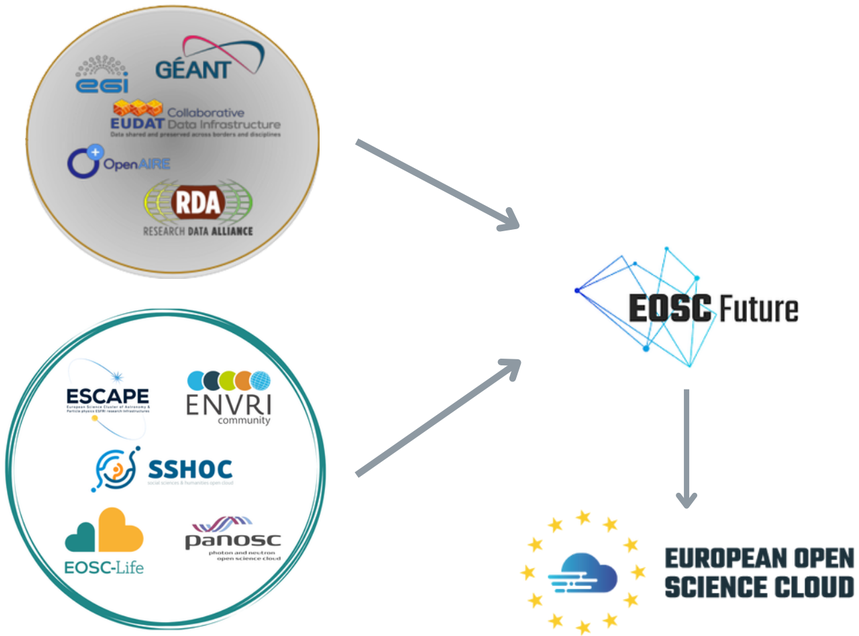

Timeline

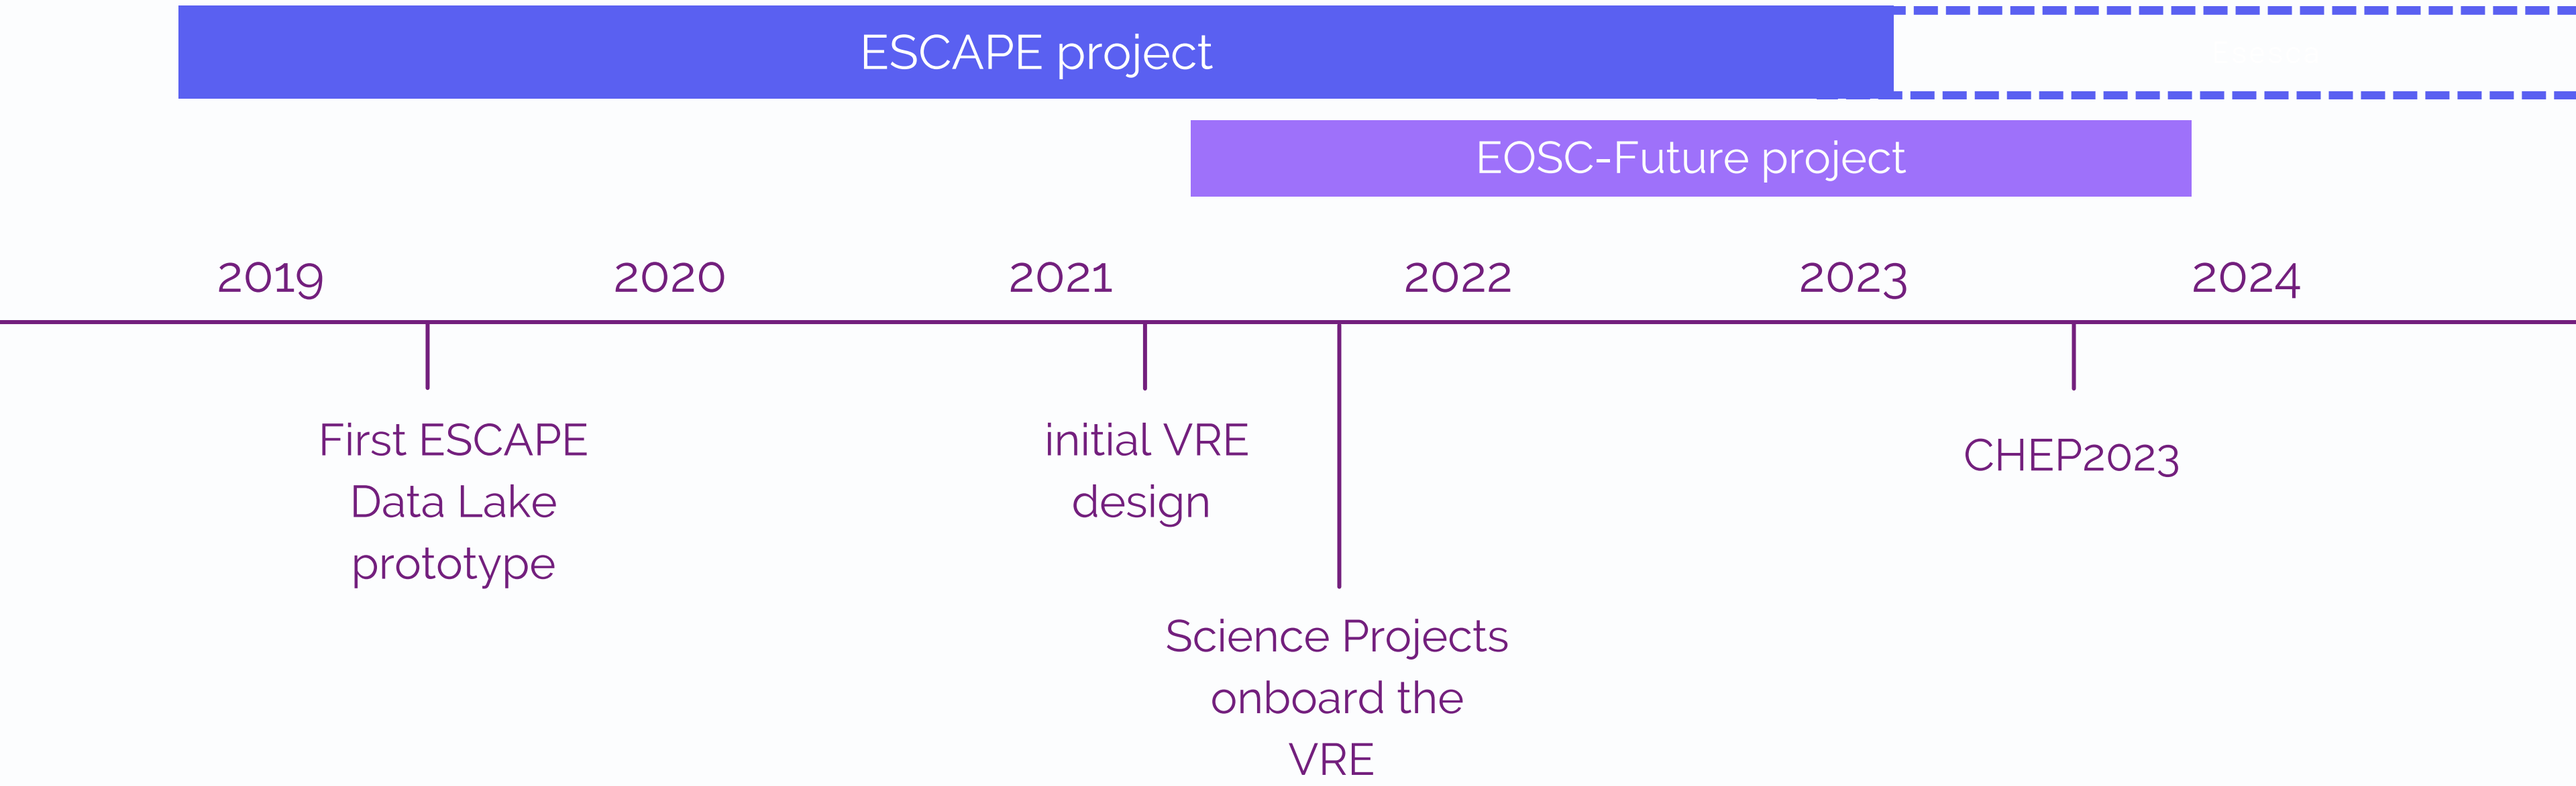
ESCAPE project
EOSC-Future project
2019
2020
2021
2022
2023
2024
First ESCAPE Data Lake prototype
initial VRE design
Science Projects onboard the VRE
CHEP2023

Dark Matter

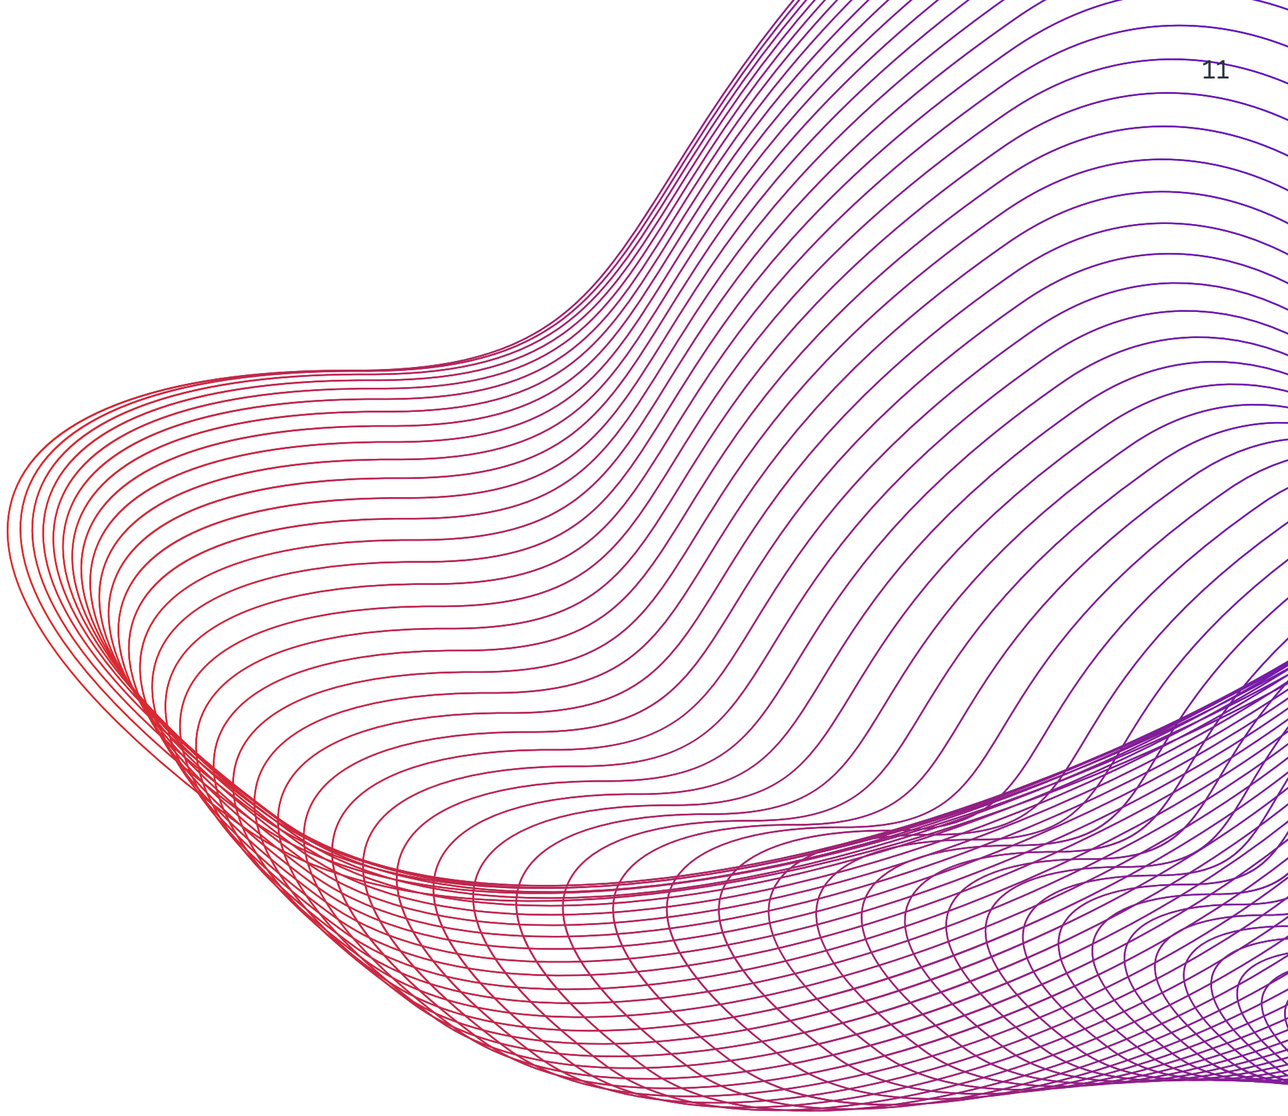

Big science question: Dark Matter

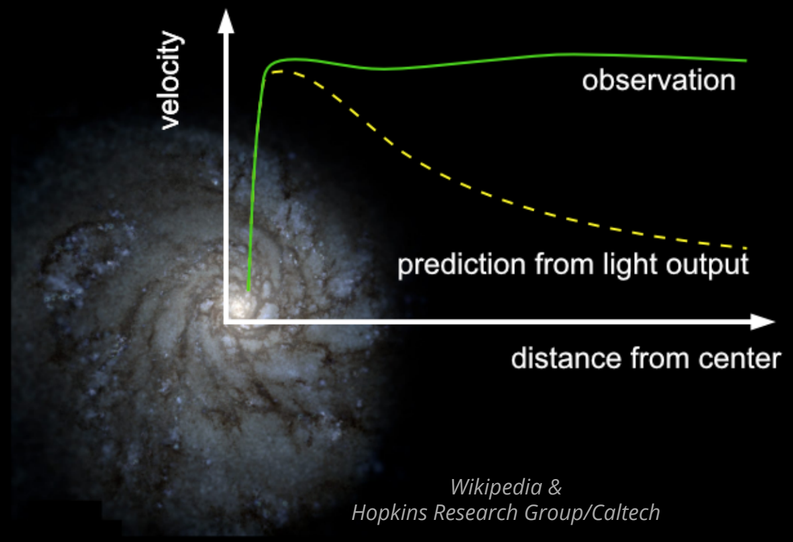

Wikipedia &
Hopkins Research Group/Caltech

Vera Rubin,
© Washington Times & Zuma

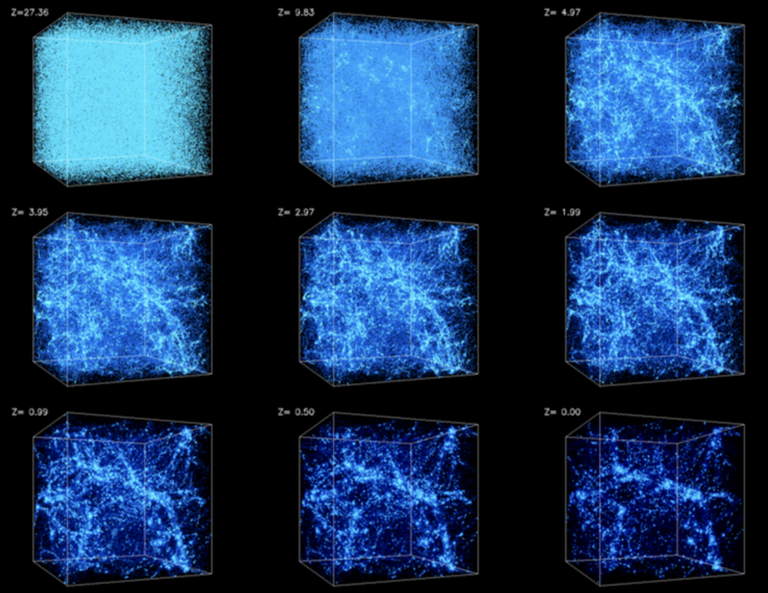

Simulations were performed at the National Center for Supercomputer Applications by A. Kravtsov and A. Klypin.

visible matter

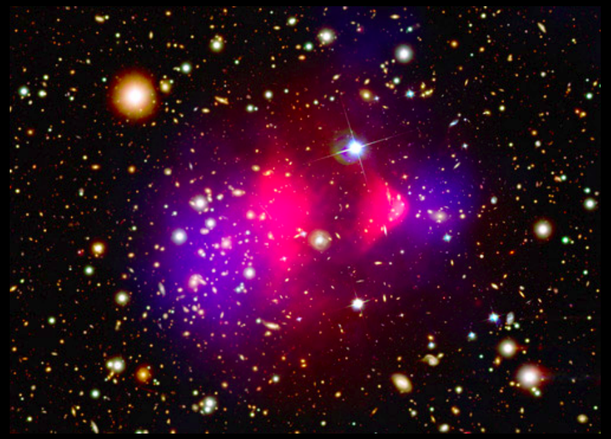

“invisible” matter

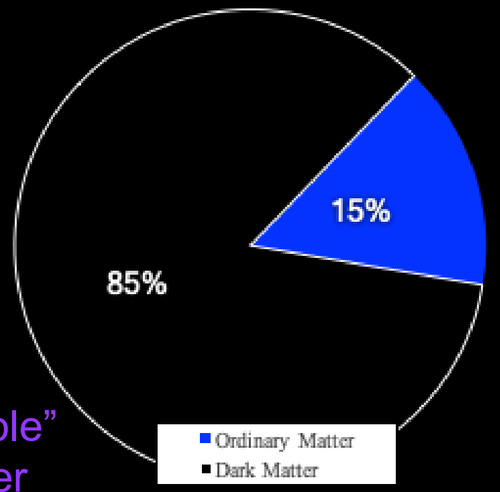

Dark Matter Experiments

- large, complex, costly experiments
- only one or a few experiments of each type worldwide

→

Maximising each experiment's science outputs is imperative

- **create** and store new analyses, datasets and results
- **combine** multiple results studying the same question
- **reinterpret** existing studies for new questions

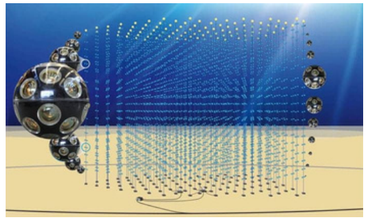

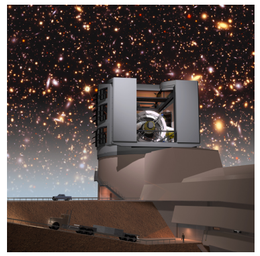

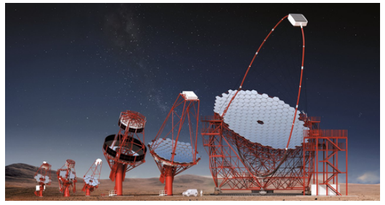

cta — cherenkov telescope array — the observatory for ground-based gamma-ray astronomy

HiLumi HL-LHC PROJECT

DARKSIDE

KM3NeT

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

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

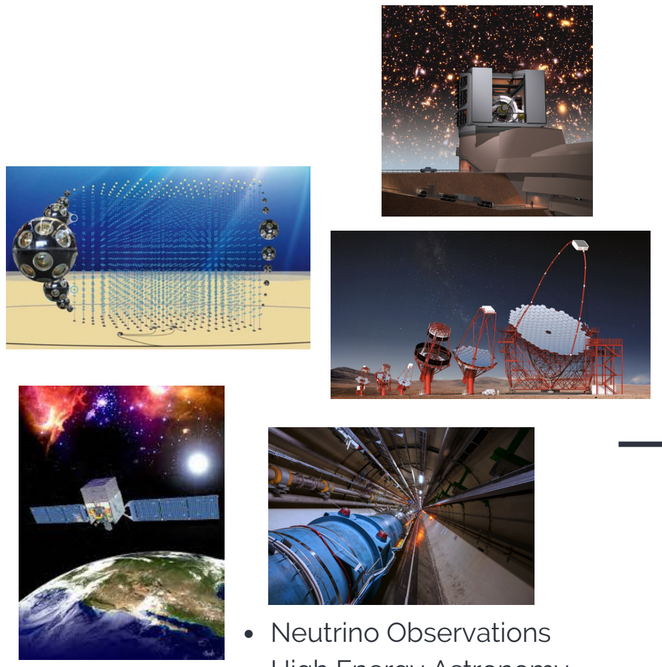

- Neutrino Observations
- High Energy Astronomy
- Low Frequency Telescopes
- Particle Physics

DARK MATTER Science Project

Virtual Research Environment

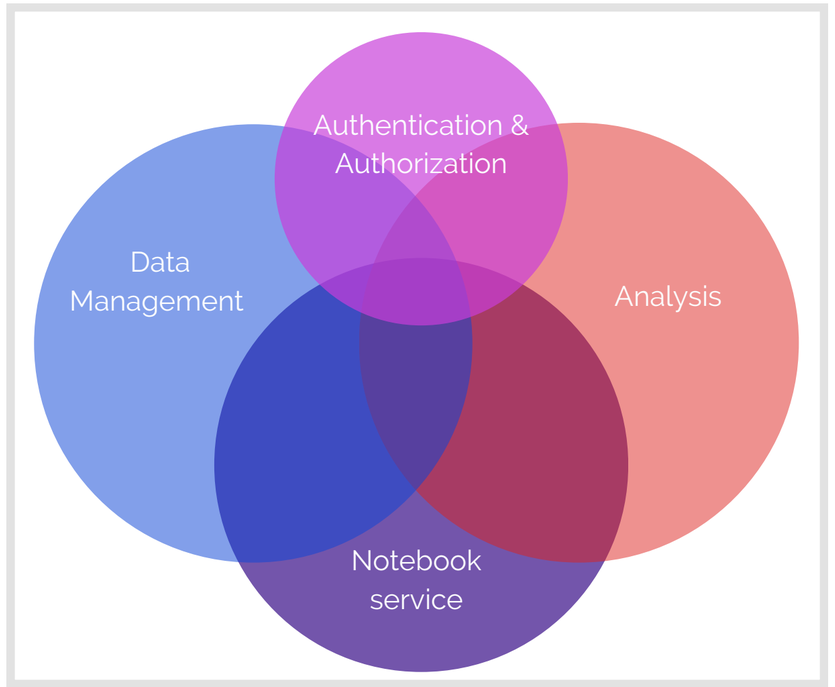

Dark Matter: Complementary Approach

Focus: Looking for **W**eakly **I**nteracting **M**assive **P**articles (**WIMP**s)

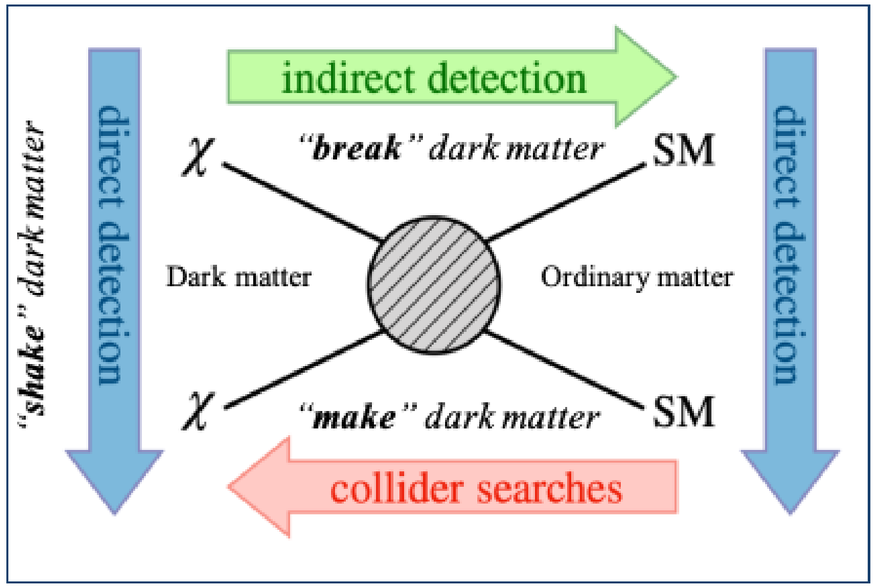

A joint discovery of the nature of dark matter requires different experiments and inputs

Experiments have **different** data sizes, workflows, data, and result sharing policies

Example of a **discovery scenario**

Late 2020s

Direct detection experiment sees a hint of a signal, with characteristics compatible with WIMP DM

Mid 2030s

2040s

Inspired by:
Dark Matter Complementarity (Snowmass report), arXiv:2210.01770
T. Slatyer's "Paths to discovery" talk at Snowmass 2022

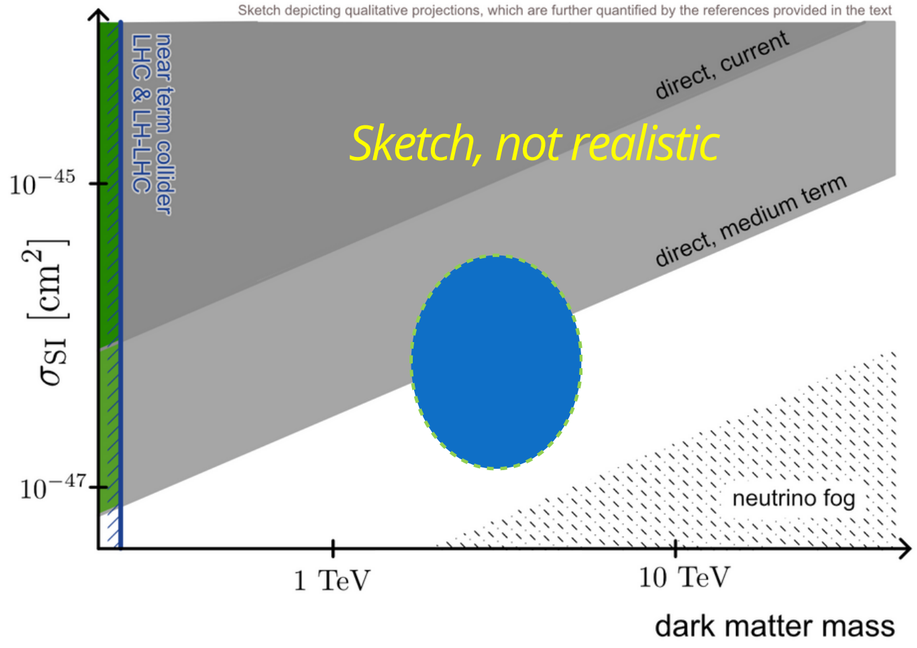

Example of a **discovery scenario**

Late 2020s

Direct detection experiment sees a hint of a signal, with characteristics compatible with WIMP DM

Mid 2030s

Direct detection experiment (using another technique) confirms these hints

2040s

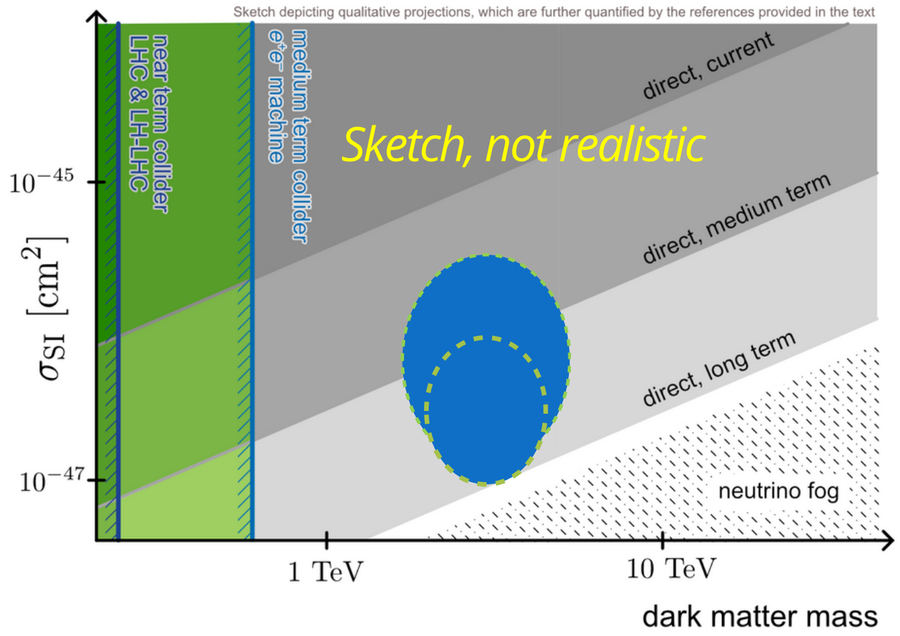

Inspired by:
Dark Matter Complementarity (Snowmass report), arXiv:2210.01770
T. Slatyer's "Paths to discovery" talk at Snowmass 2022

Example of a **discovery scenario**

Late 2020s

Direct detection experiment sees a hint of a signal, with characteristics compatible with WIMP DM

Mid 2030s

Direct detection experiment (using another technique) confirms these hints

Indirect detection experiment observes signals of DM annihilation

2040s

Inspired by:
Dark Matter Complementarity (Snowmass report), arXiv:2210.01770
T. Slatyer's "Paths to discovery" talk at Snowmass 2022

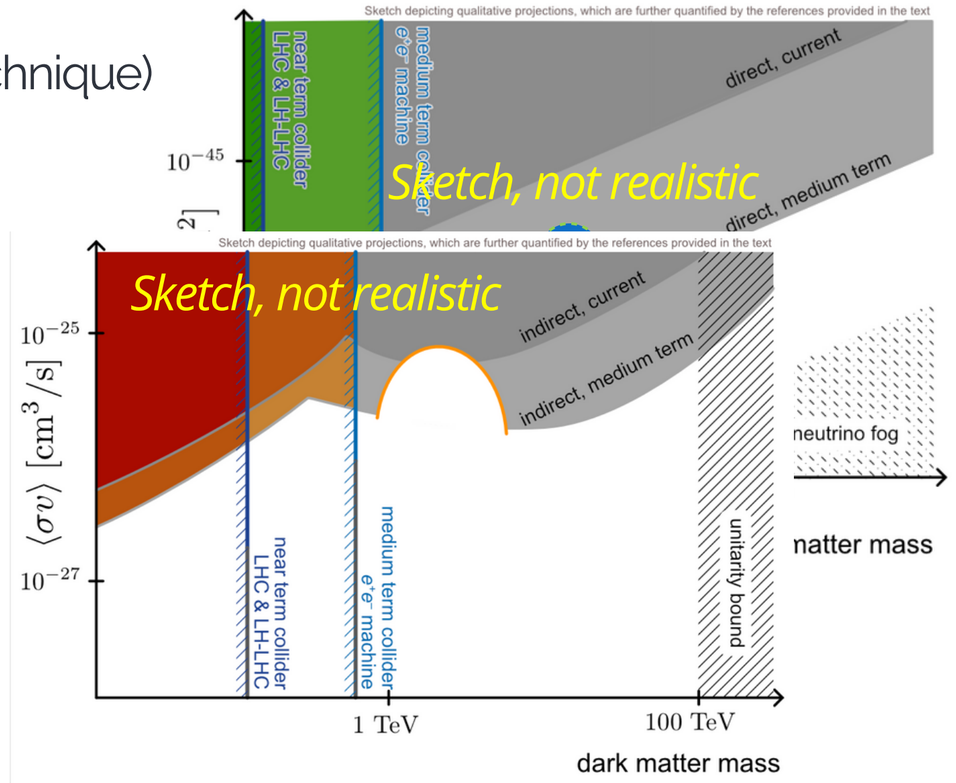

Example of a **discovery scenario**

Late 2020s

Direct detection experiment sees a hint of a signal, with characteristics compatible with WIMP DM

Direct detection experiment (using another technique) confirms these hints

Mid 2030s

Indirect detection experiment observes signals of DM annihilation

2040s

Future collider, built to target particles with the mass of the putative DM candidate, sheds light on interactions between DM and ordinary matter

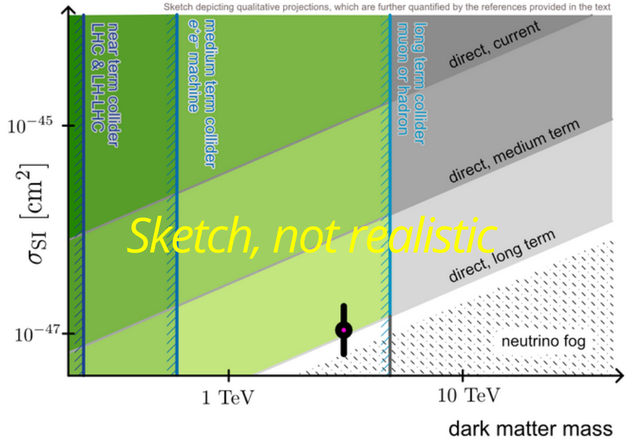

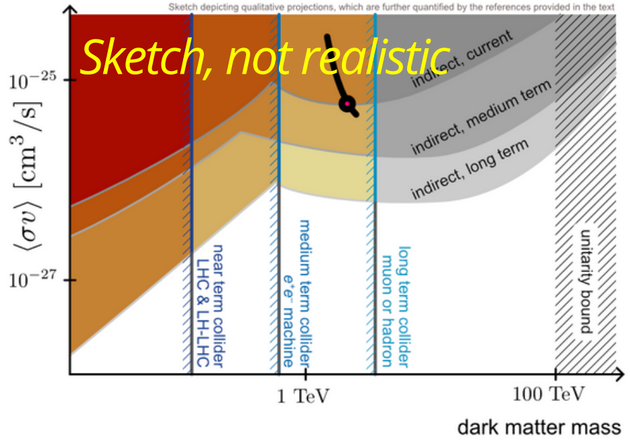

Inspired by:
Dark Matter Complementarity (Snowmass report), arXiv:2210.01770
T. Slatyer's "Paths to discovery" talk at Snowmass 2022

Example of a **discovery scenario**

Late 2020s

Direct detection experiment sees a hint of a signal, with characteristics compatible with WIMP DM

Mid 2030s

Direct detection experiment (using another technique) confirms these hints

Indirect detection experiment observes signals of DM annihilation

2040s

Future collider, built to target particles with the mass of the putative DM candidate, sheds light on interactions between DM and ordinary matter

Inspired by:
Dark Matter Complementarity (Snowmass report), arXiv:2210.01770
T. Slatyer's "Paths to discovery" talk at Snowmass 2022

Such a scenario requires **interoperable** and **reproducible** analyses

- **comparison** and **combination** of results from different experiments
- **end-to-end workflows** available for cross-checks

with the **Dark Matter Science Project**, we build a **prototype** that fulfils these requirements

Experiments involved in the Dark Matter Science Project

Direct detection: DarkSide

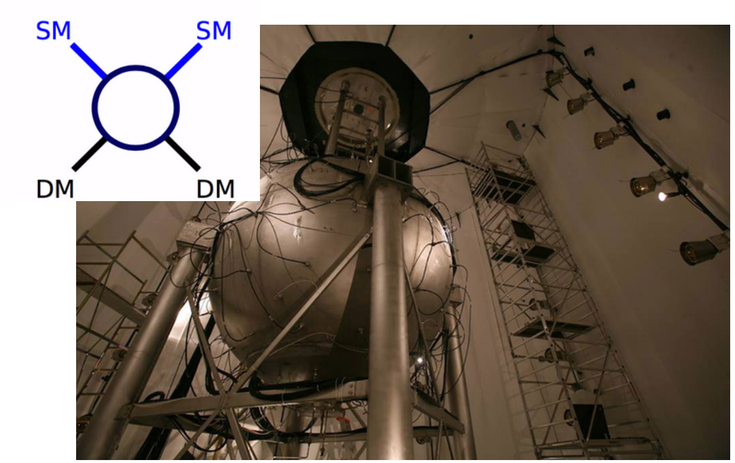

Colliders: ATLAS @ LHC

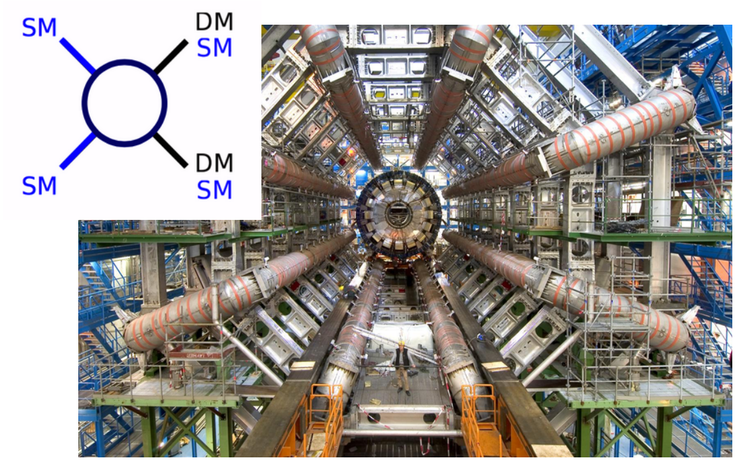

Indirect detection: FermiLAT, KM3NeT

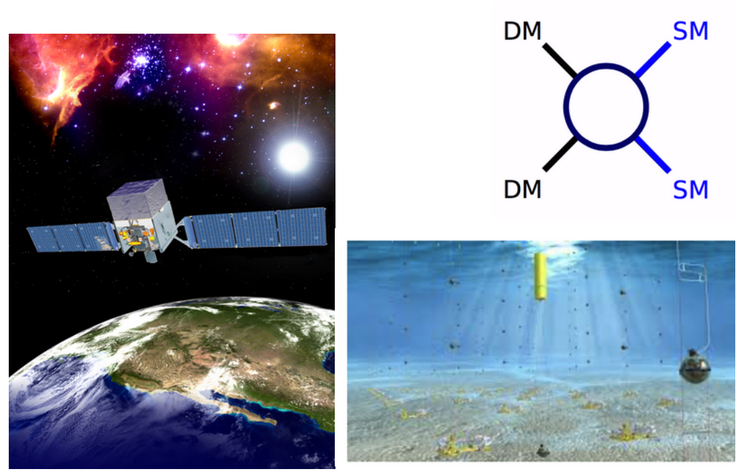

...and their evolutions: **DarkSide-20k / Argo, ATLAS @ HL-LHC, CTA**
Some of the **analysis & ML tools** necessary for these evolutions are also part of this Science Project

With the Dark Matter Science Project,
we understand the computing and analysis challenges of some of the future DM experiments

Science outputs of the Dark Matter SP

- Individual results and publications
- Plots highlighting complementarity of different experimental efforts
- Data and software objects + pipelines
- Data on the Data Lake, and software on the ESCAPE Software Catalogue
- Pipelines accessible via VRE
- **Combination of experimental results**

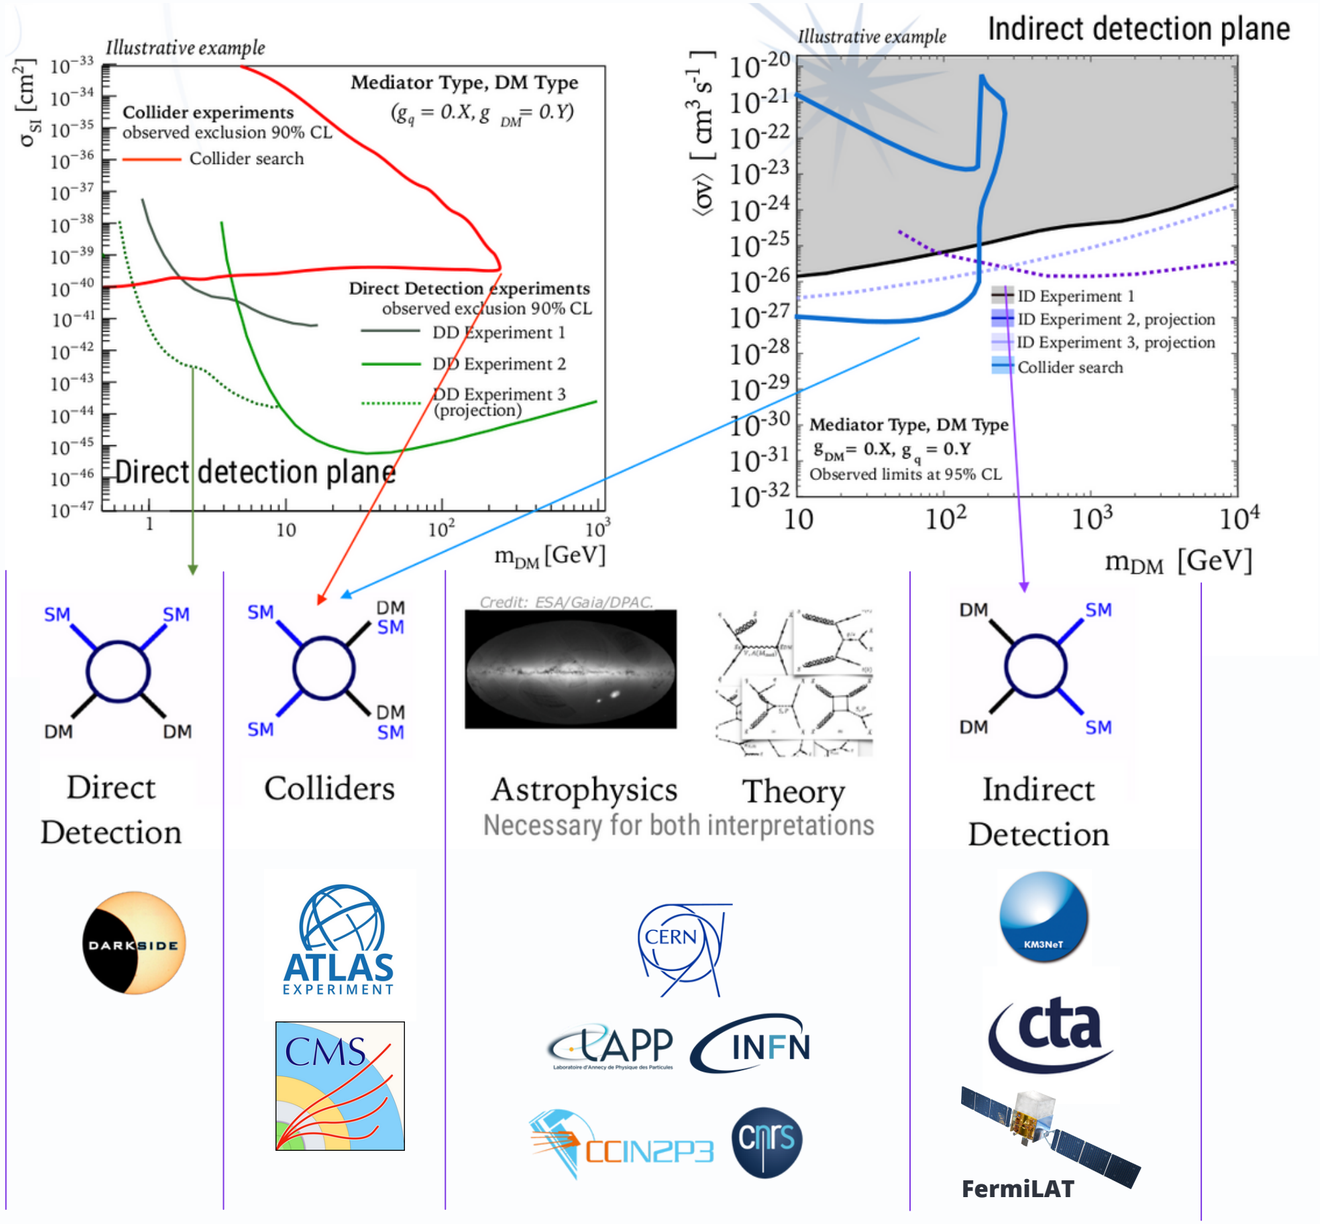

Analysis Workflows for the DM Science Project

Generation & simulation

Expt. data

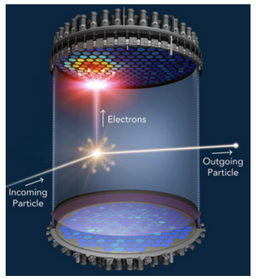

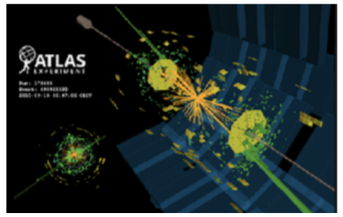

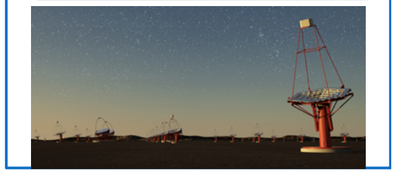

Credit: SLAC/LZ/ATLAS/CTA

Data **processing**

JupyterHub
(smaller workflows)

REANA
(larger workflows)

Analysis of events/distributions

Interpretation of results

Combination of results with other searches/experiments

Comparison of results with other searches / experiments

Virtual Research Environment

ESCAPE Data Lake

ESCAPE Software Catalogue

Data sharing and data processing

Data analysis, preservation and interpretation

Virtual Research Environment

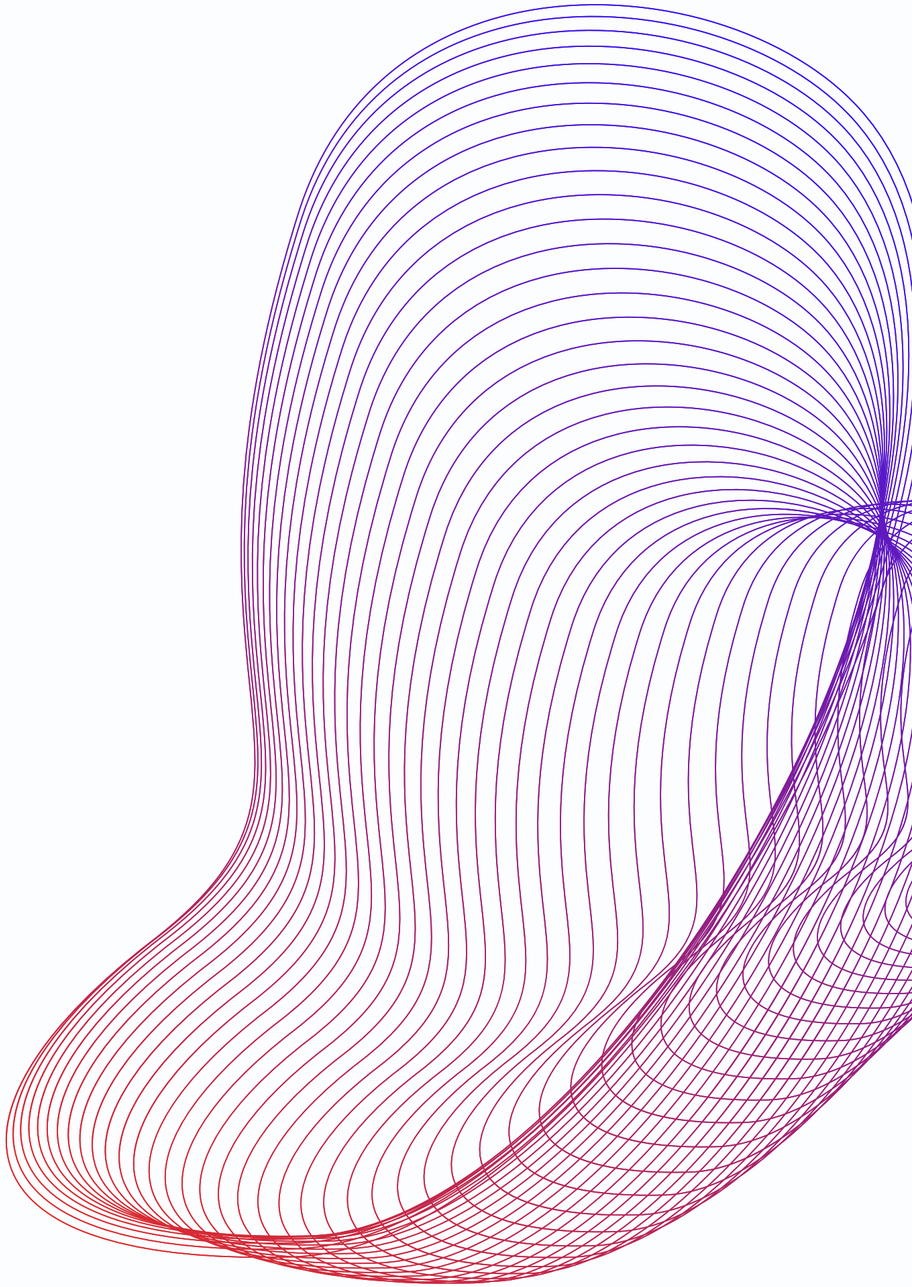

The building blocks

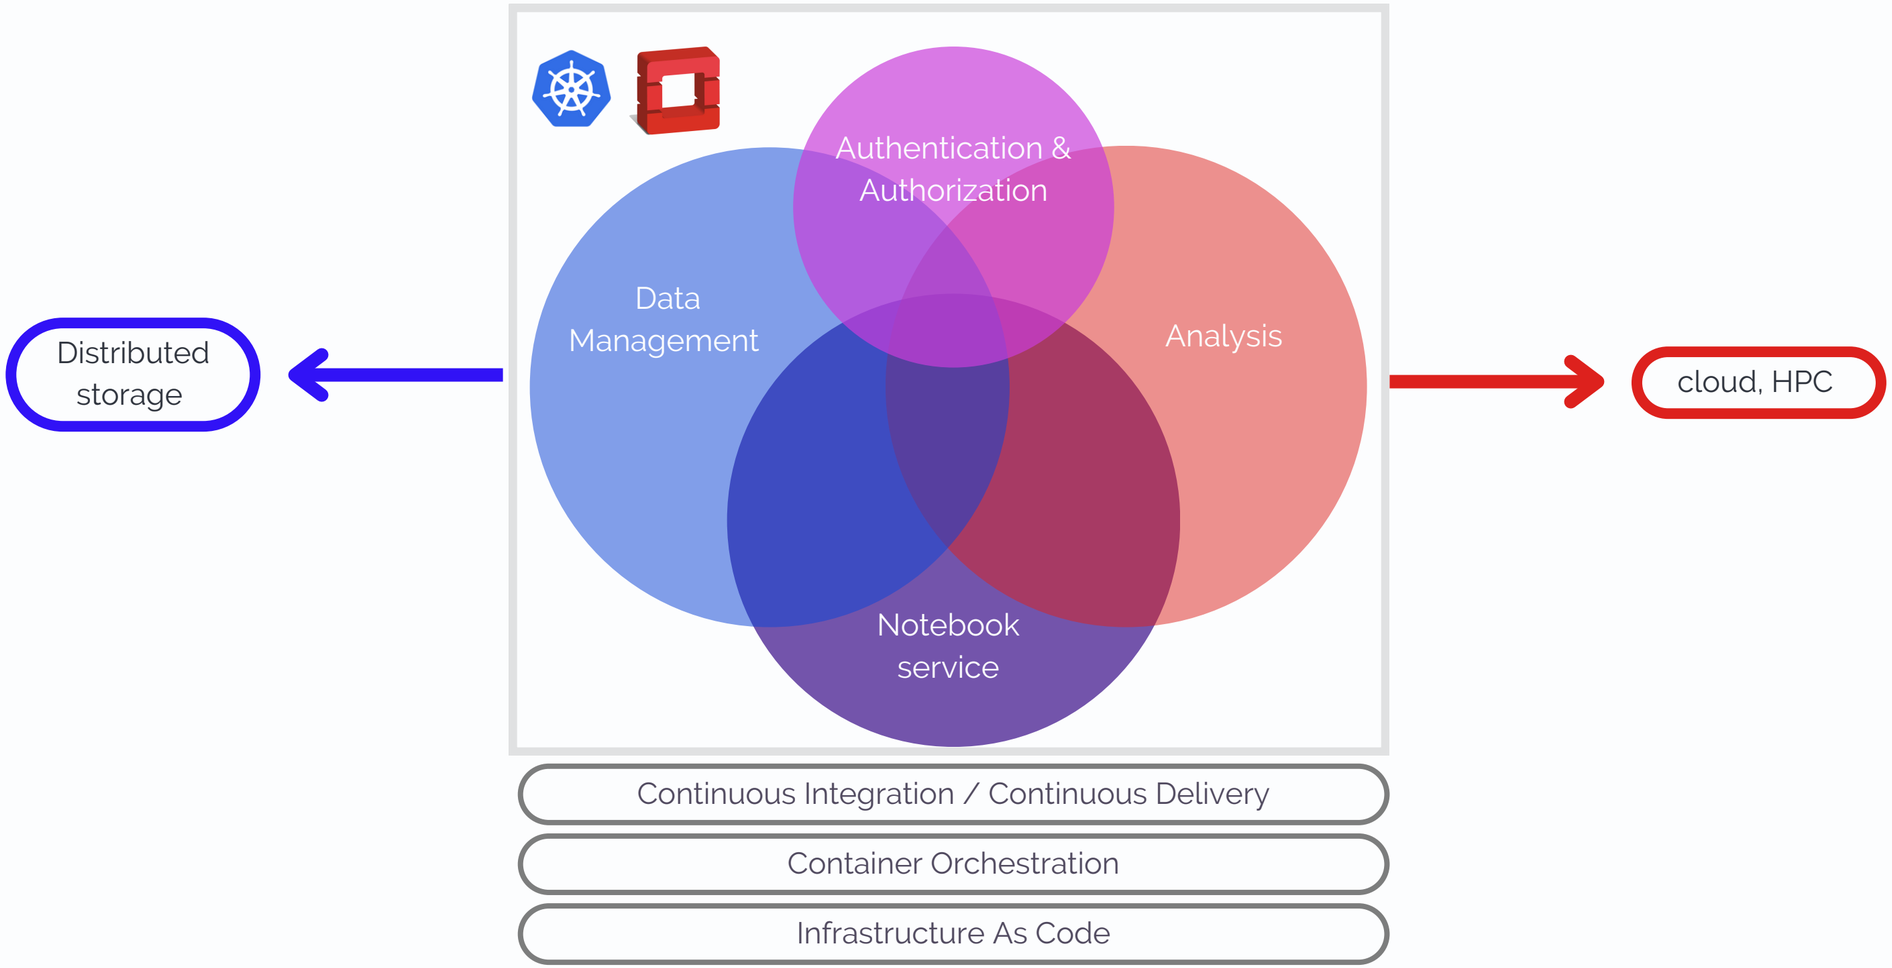

Authentication & Authorisation

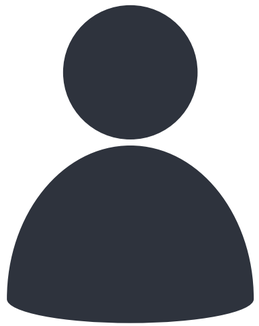

INDIGO Identity and Access Management (IAM) - adopted by WLCG for token

- OIDC tokens
- X.509 certificates / one VO for all the experiments

@ INFN CNAF

Authentication & Authorization

subject mapping cronjob

Data Management

Notebook service

Analysis

Data Management

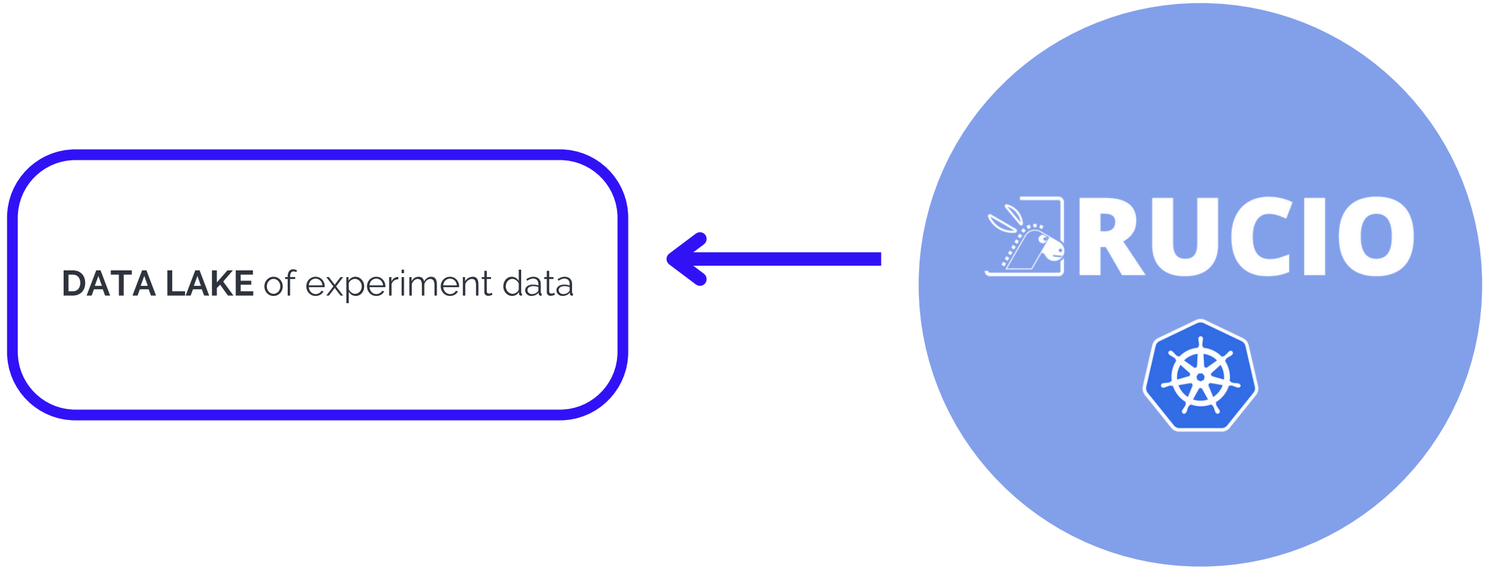

Rucio is an **open-source data management and orchestration** project initially developed by the ATLAS experiment to manage large volumes of data. It is now used by various CERN and non-CERN communities.

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

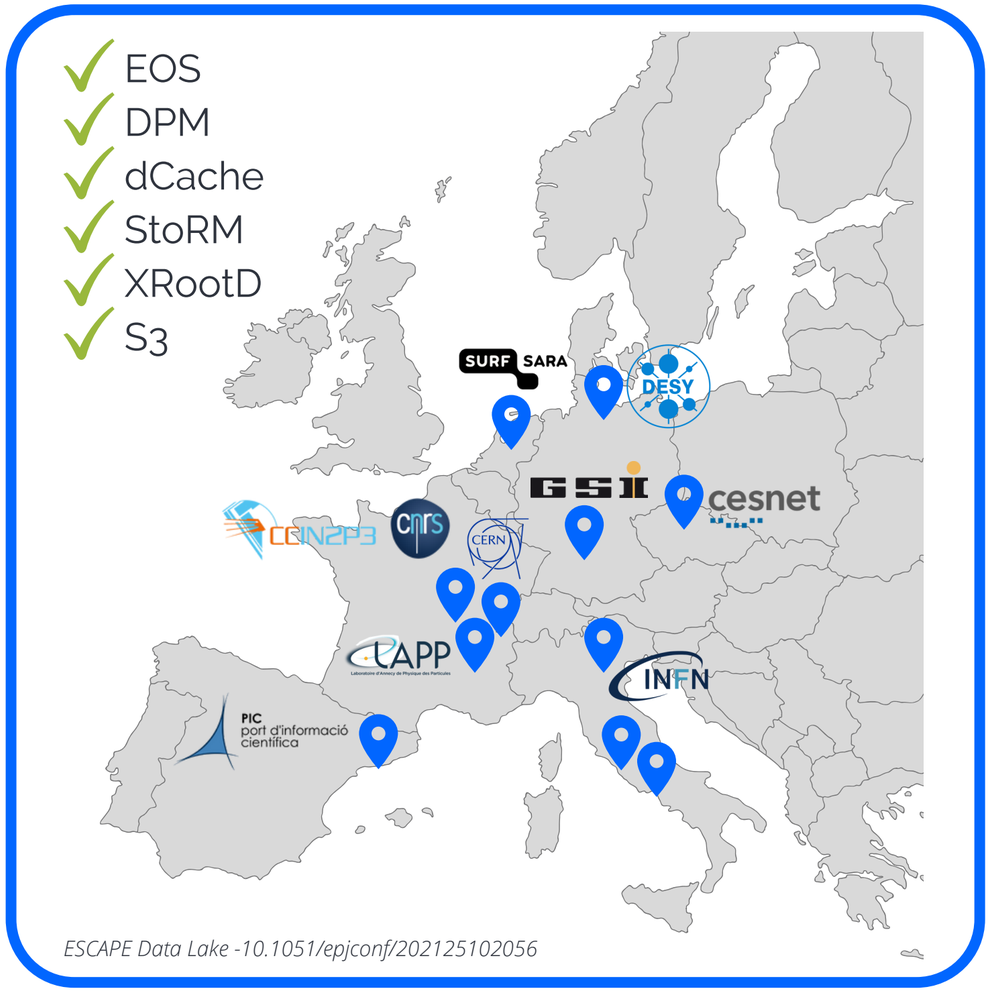

ESCAPE Data Lake -10.1051/epjconf/202125102056

Central Relational Database

FTS3and GFAL2

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

Main & authorisation servers

Notebook Service

To facilitate interactive analysis.

interface to run preliminary analysis

containerised environments on public repositories

CERN Virtual Machine FS (CVMFS) installed

client libraries and software installed to interact with underlying services

CephFS volumes provided as shared, temporary storage solution

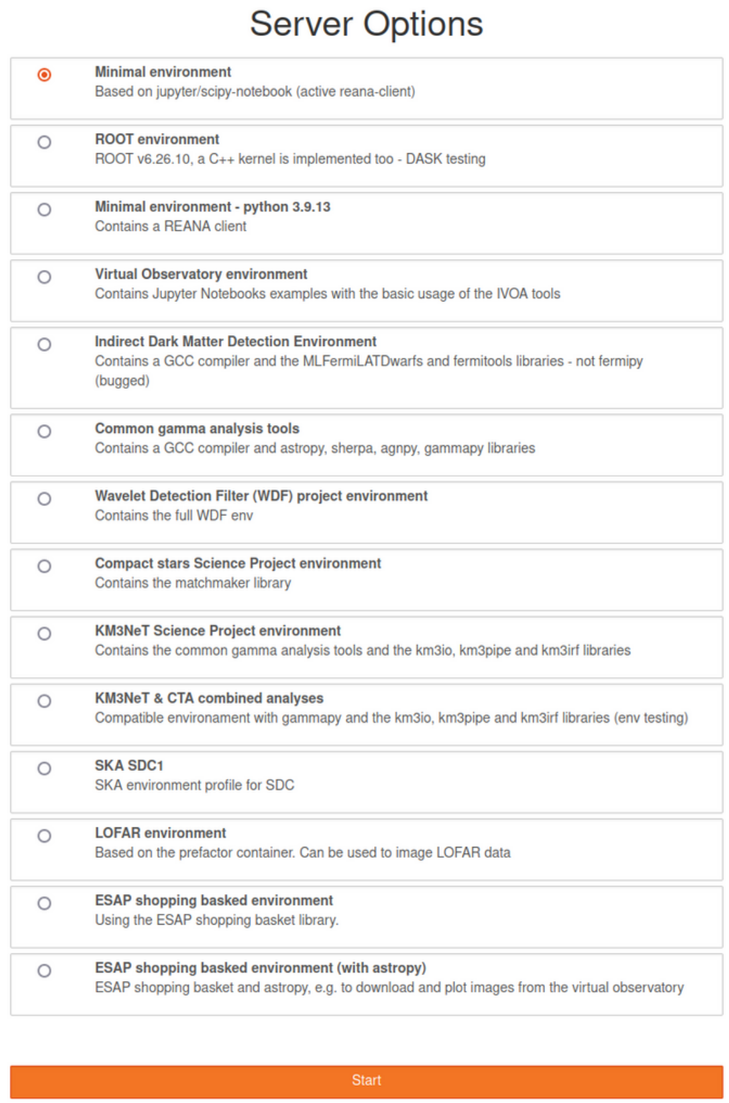

Data into the notebook

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

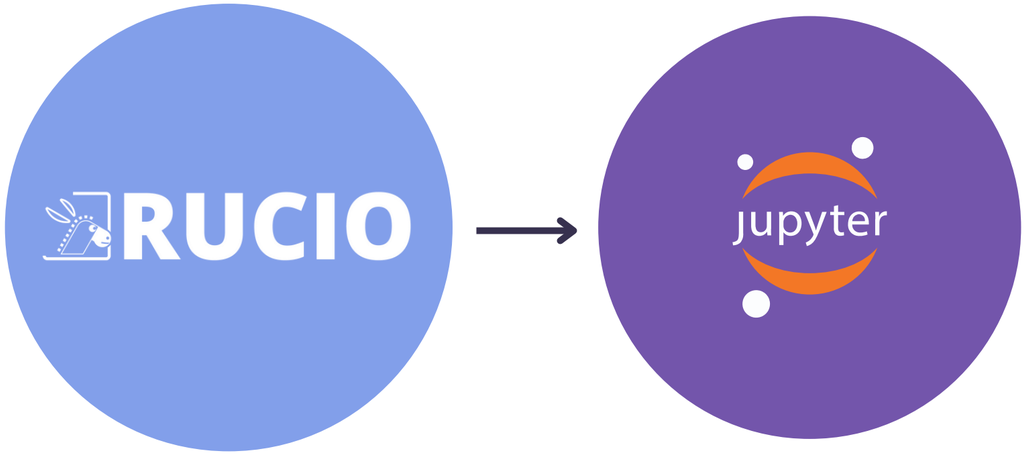

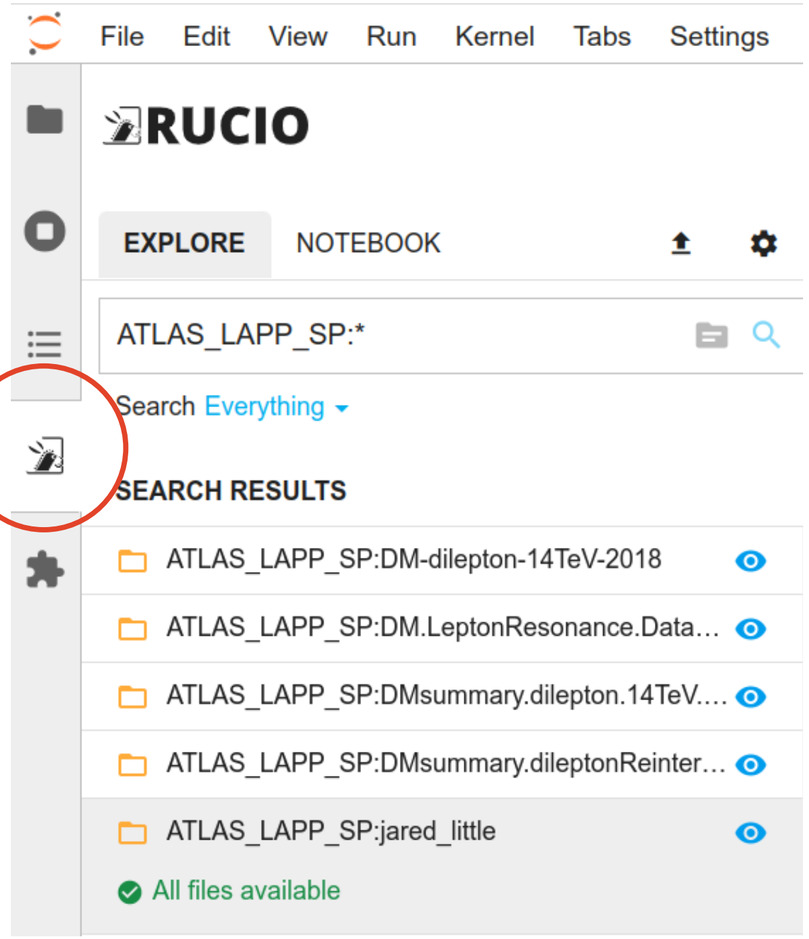

Data into the notebook

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?

Openstack
Jupyterhub node

FUSE mount

Data Lake

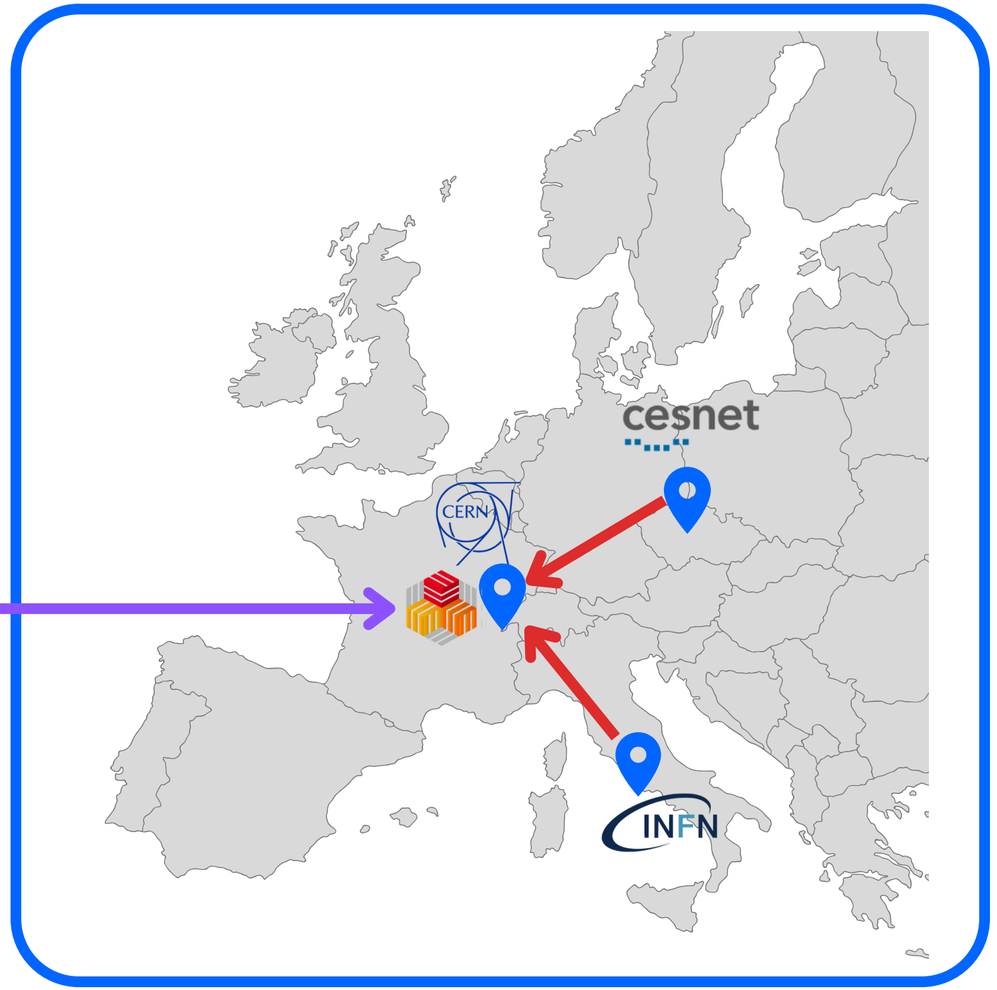

Computing

- **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)

Authentication & Authorization

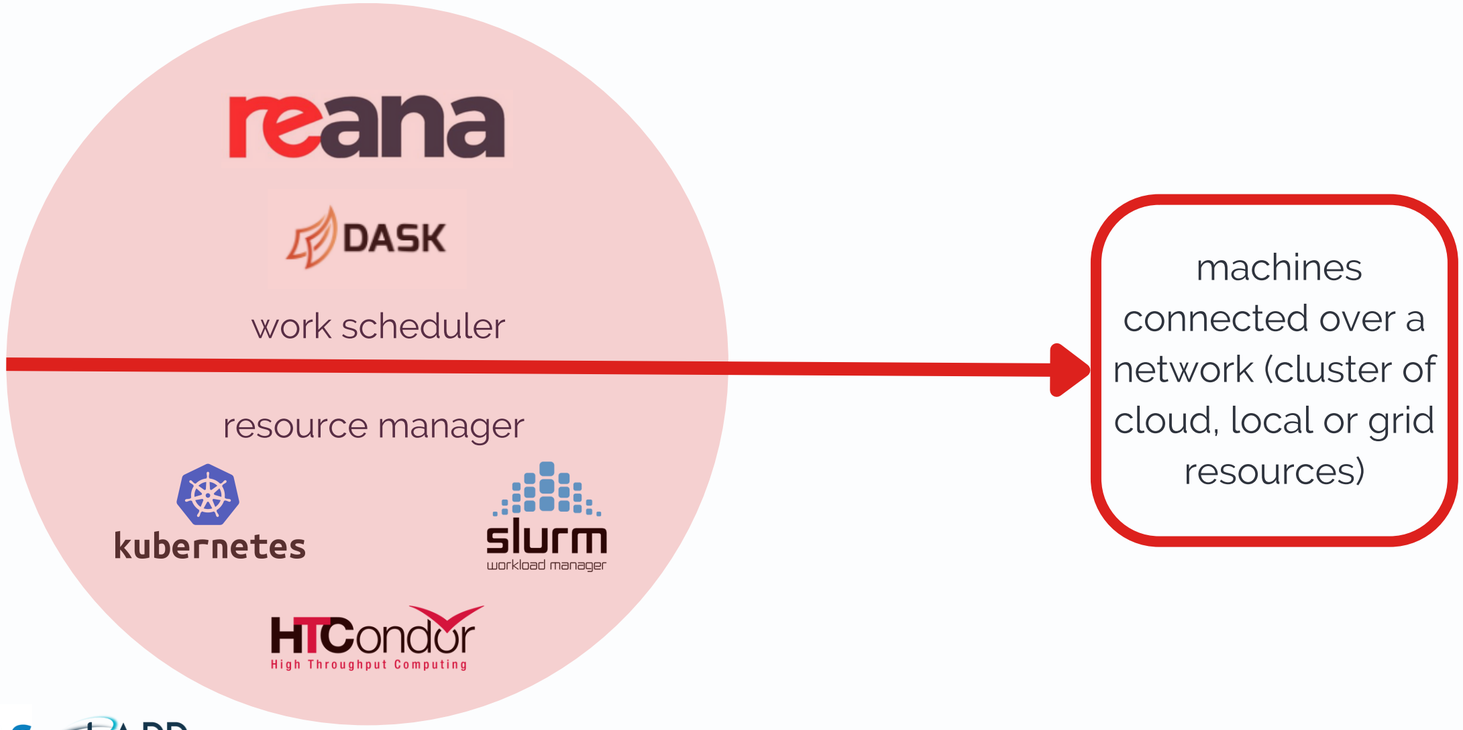

Analysis preservation and distribution

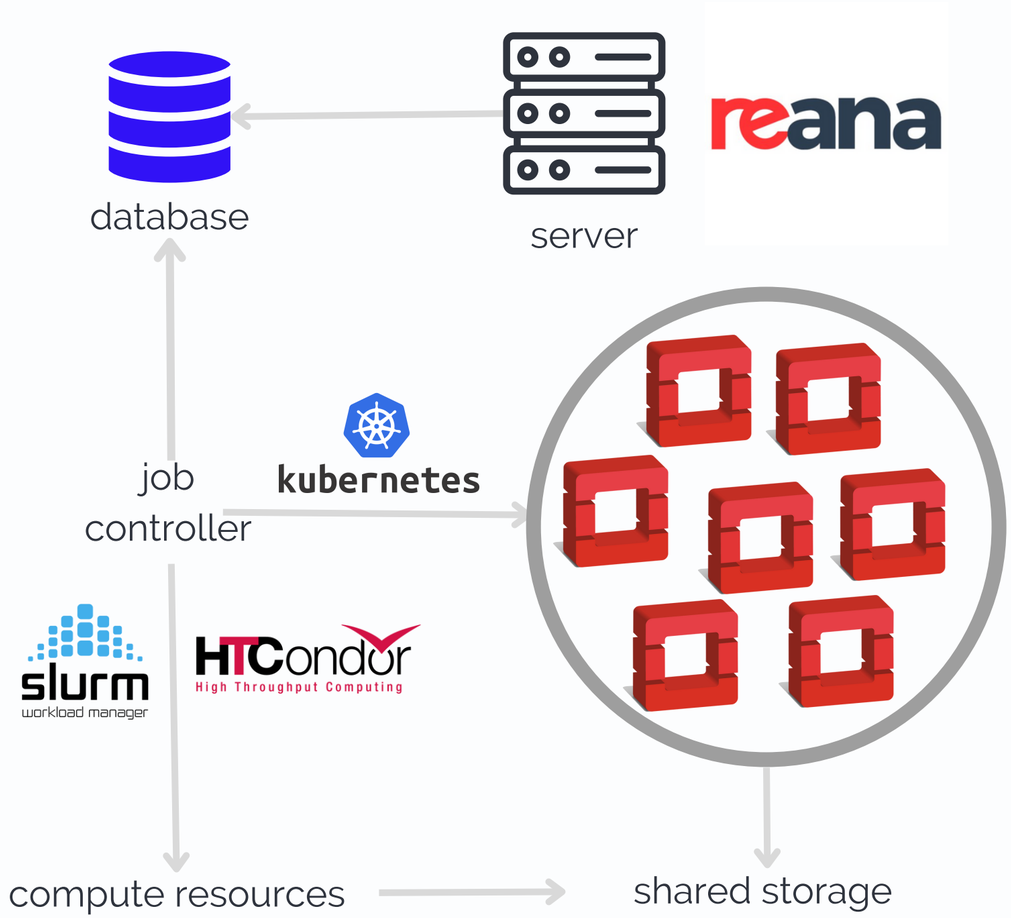

Reana is a reproducible analysis project developed at CERN, to make the **preservation** of heavier analyses seamless.

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

Non-local analysis preservation

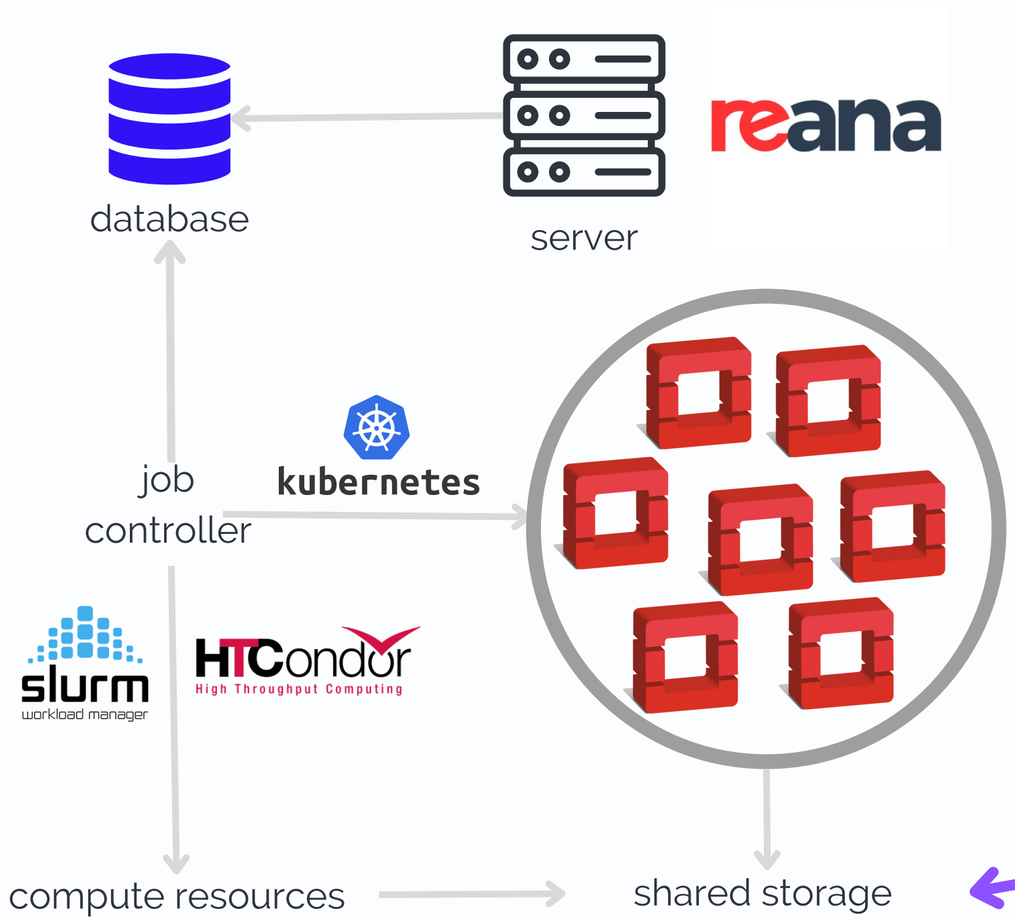

From the Reana client:

- 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**

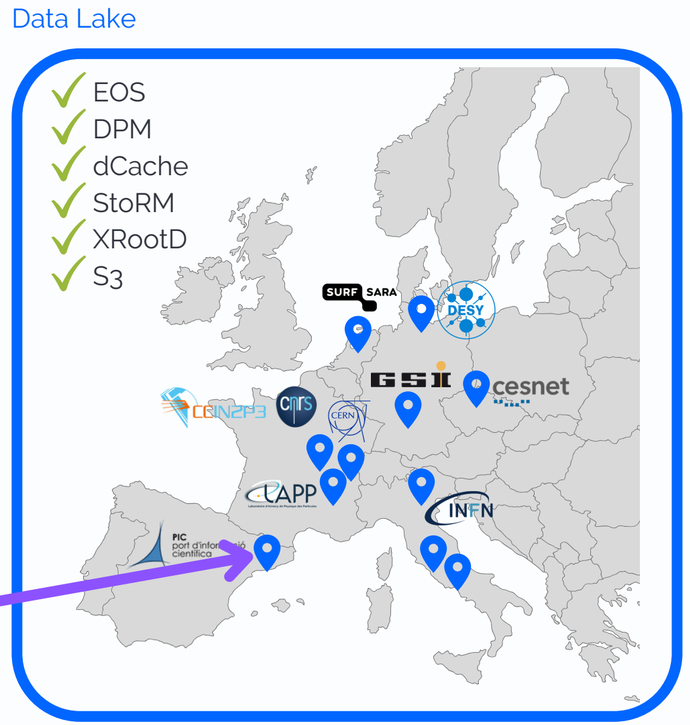

Workflow distribution with Dask

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

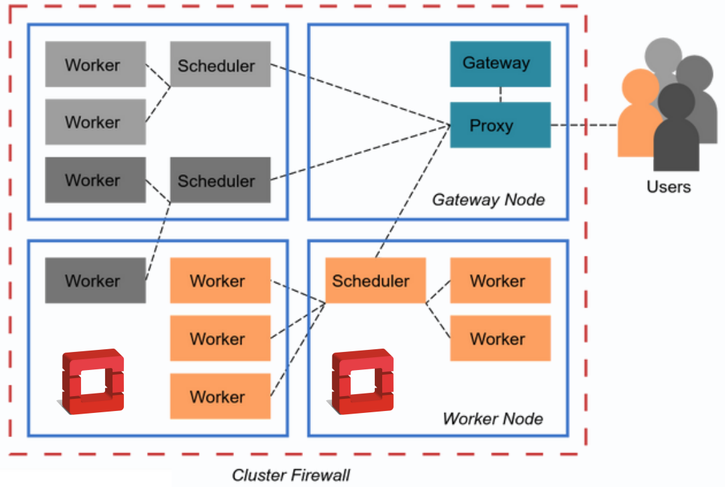

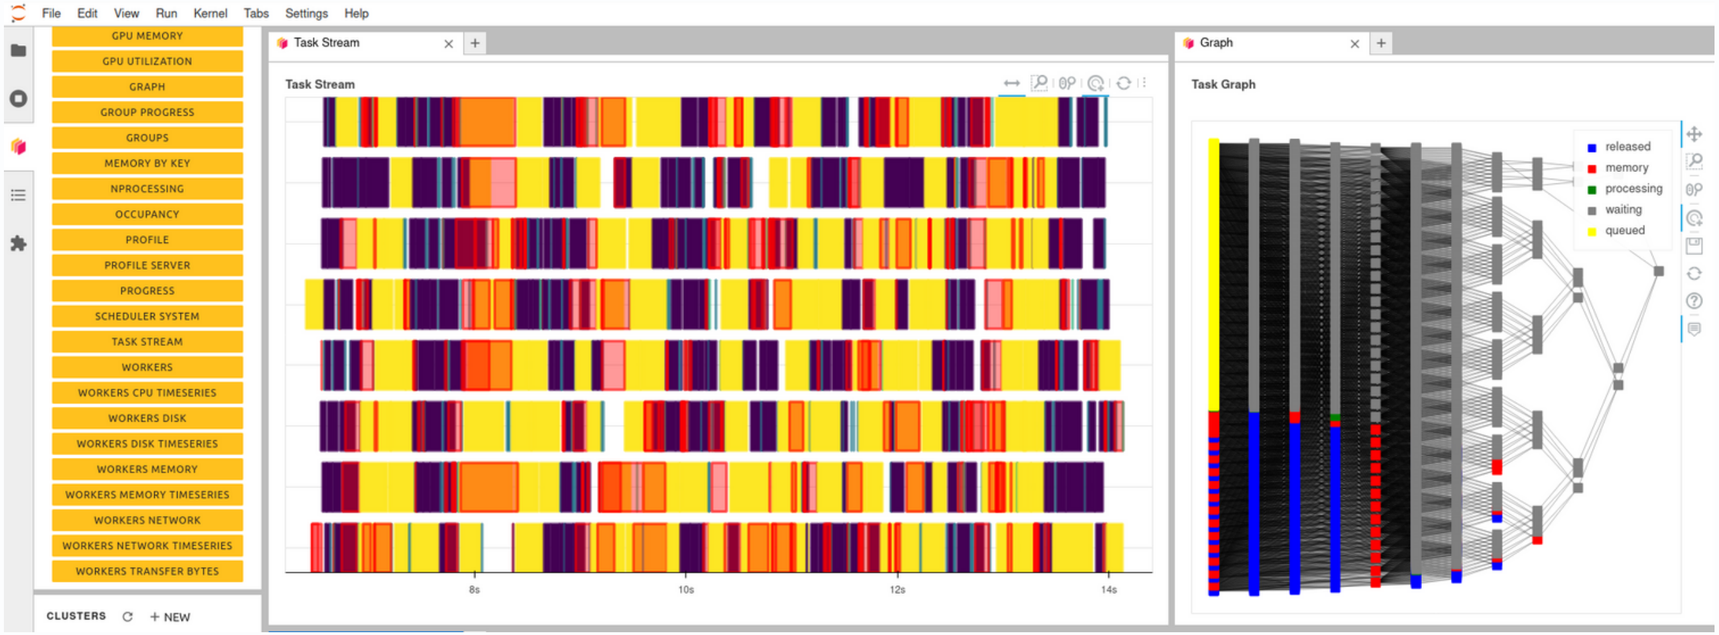

Monitoring, testing, dashboards, on-boarding

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

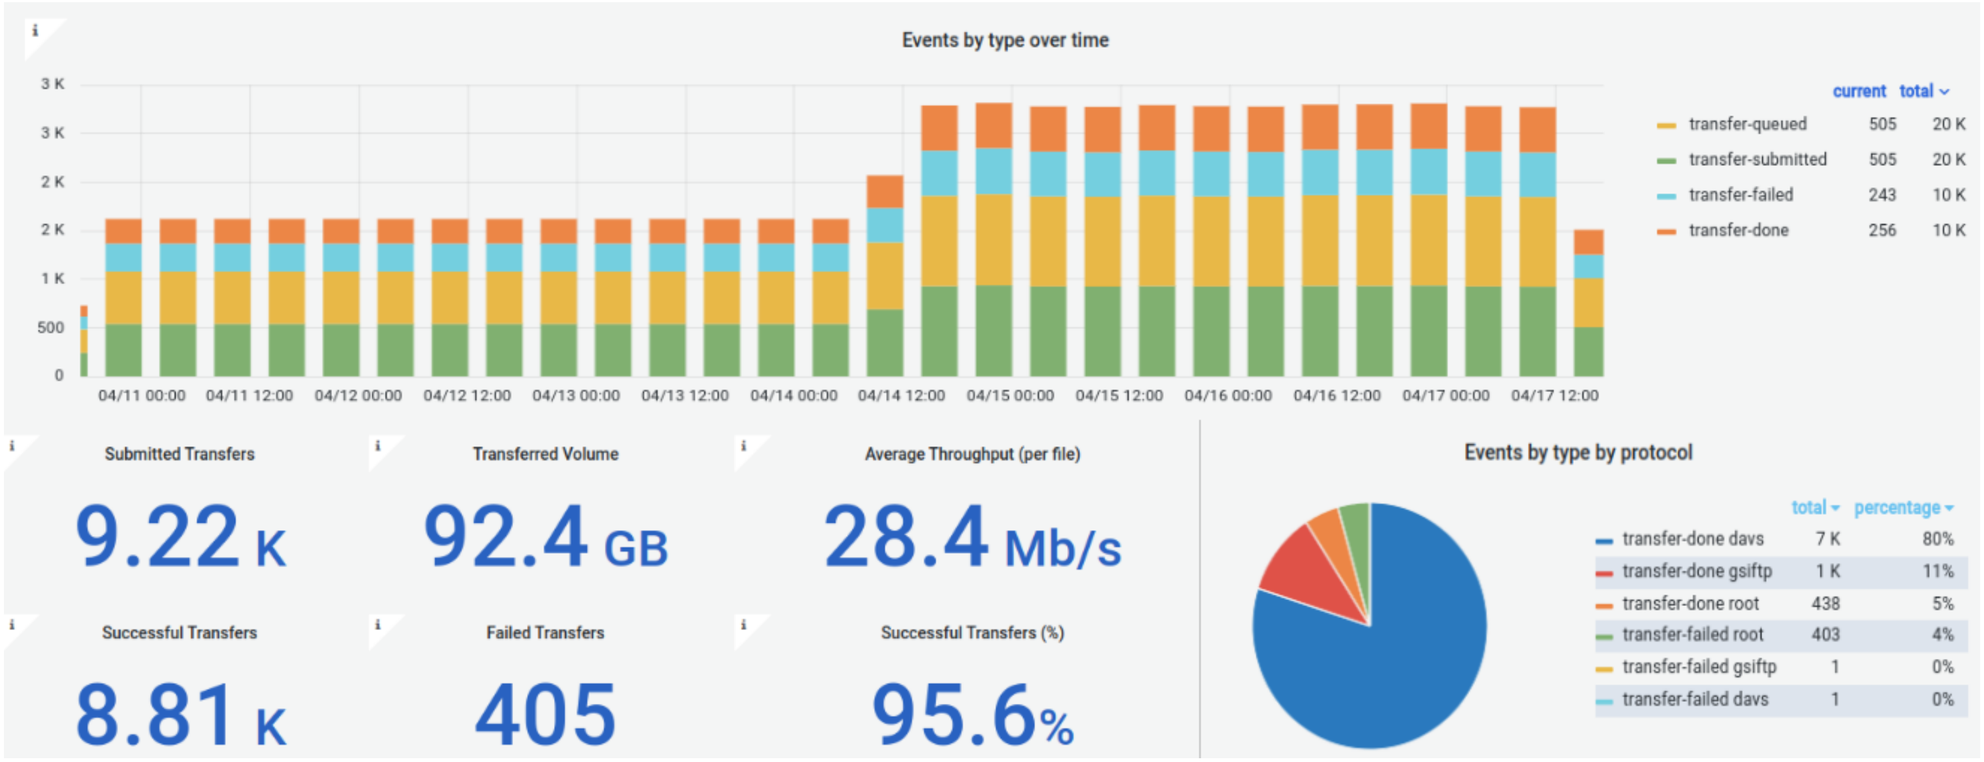

Monitoring, testing, dashboards, on-boarding

- **Rucio and Reana UI** interfaces deployed with K8s allow to explore and debug failed transfers and workflows.

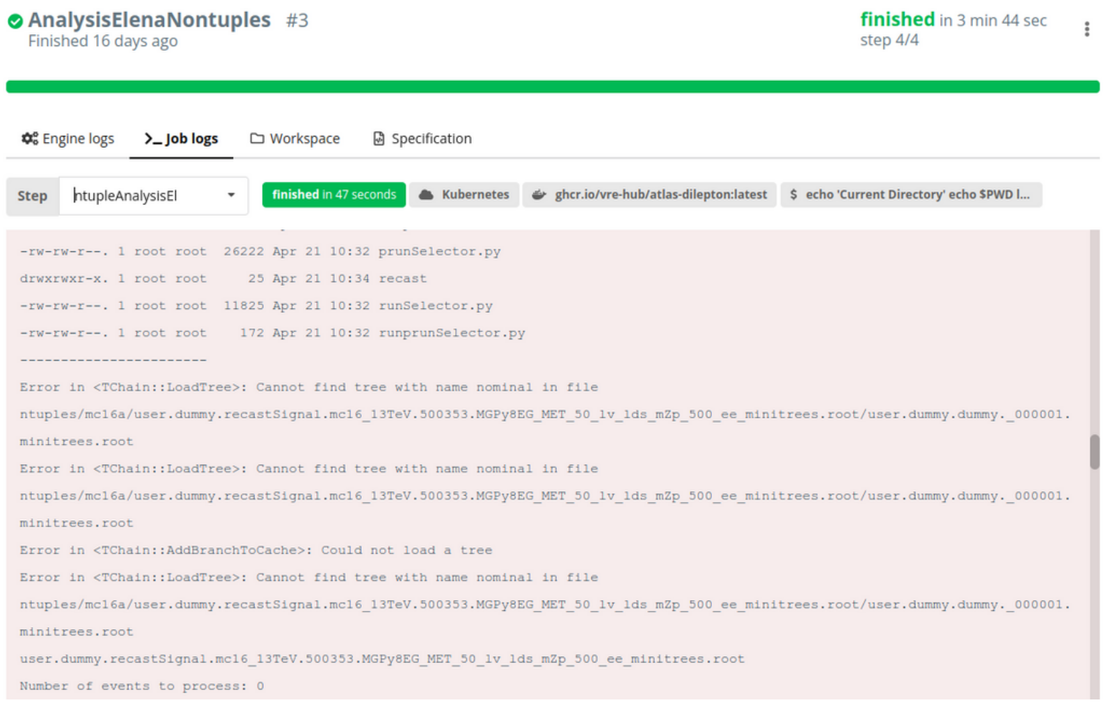

Name	Account	RSE Expression	Creation Date	Remaining Lifetime	State
elena_test:2023.03.16-11.19.03.txt	egazzarr	EULAKE-1	2023-05-07T13:22:23.000Z	7d	STUCK
user.ron:test_from_CERN-030523_1643.txt	garcia	SURF-IOP-EXP	2023-05-04T10:35:14.000Z	-	STUCK
user.ron:test_from_CERN-030523_1643.txt	garcia	EULAKE-1	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

- **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
- **Public Github repository** hosts
  - cloud deployment of the infrastructure components with Helm, Flux, Terraform and K8s
  - Science Projects software to reproduce the analyses
  - forums and discussions

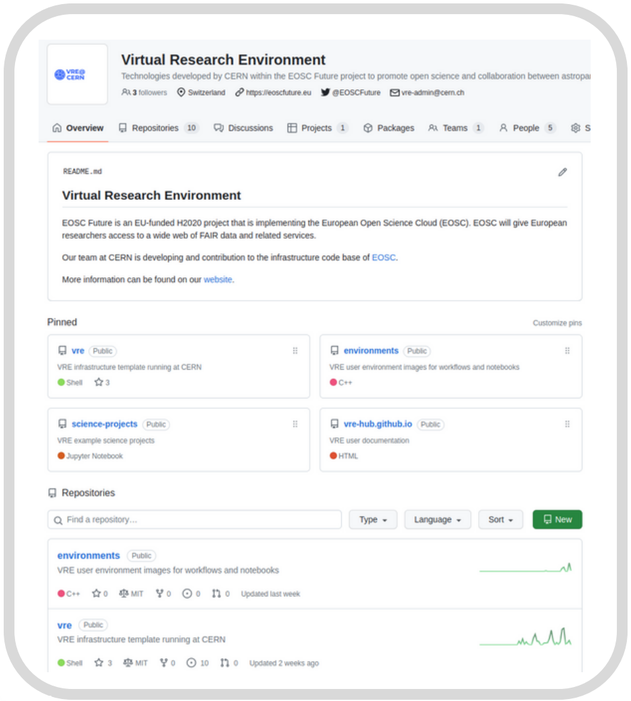

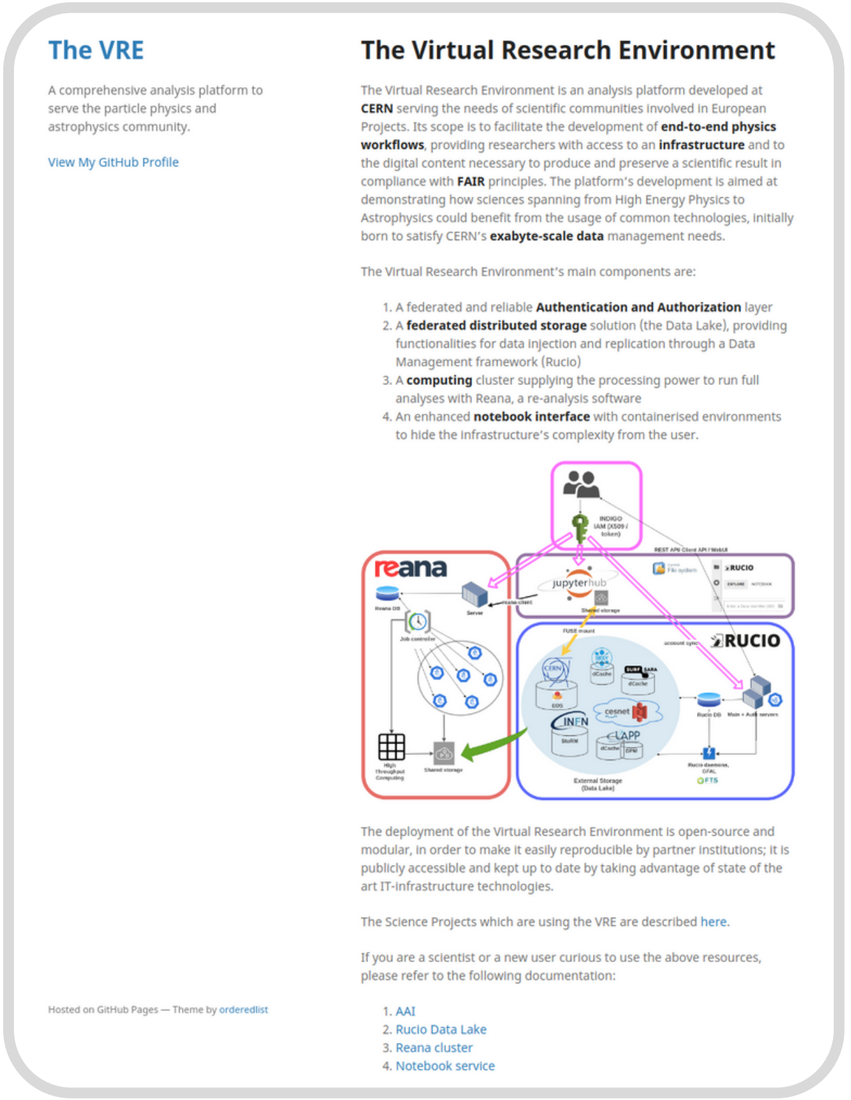

vre-hub.github.io

Science Projects

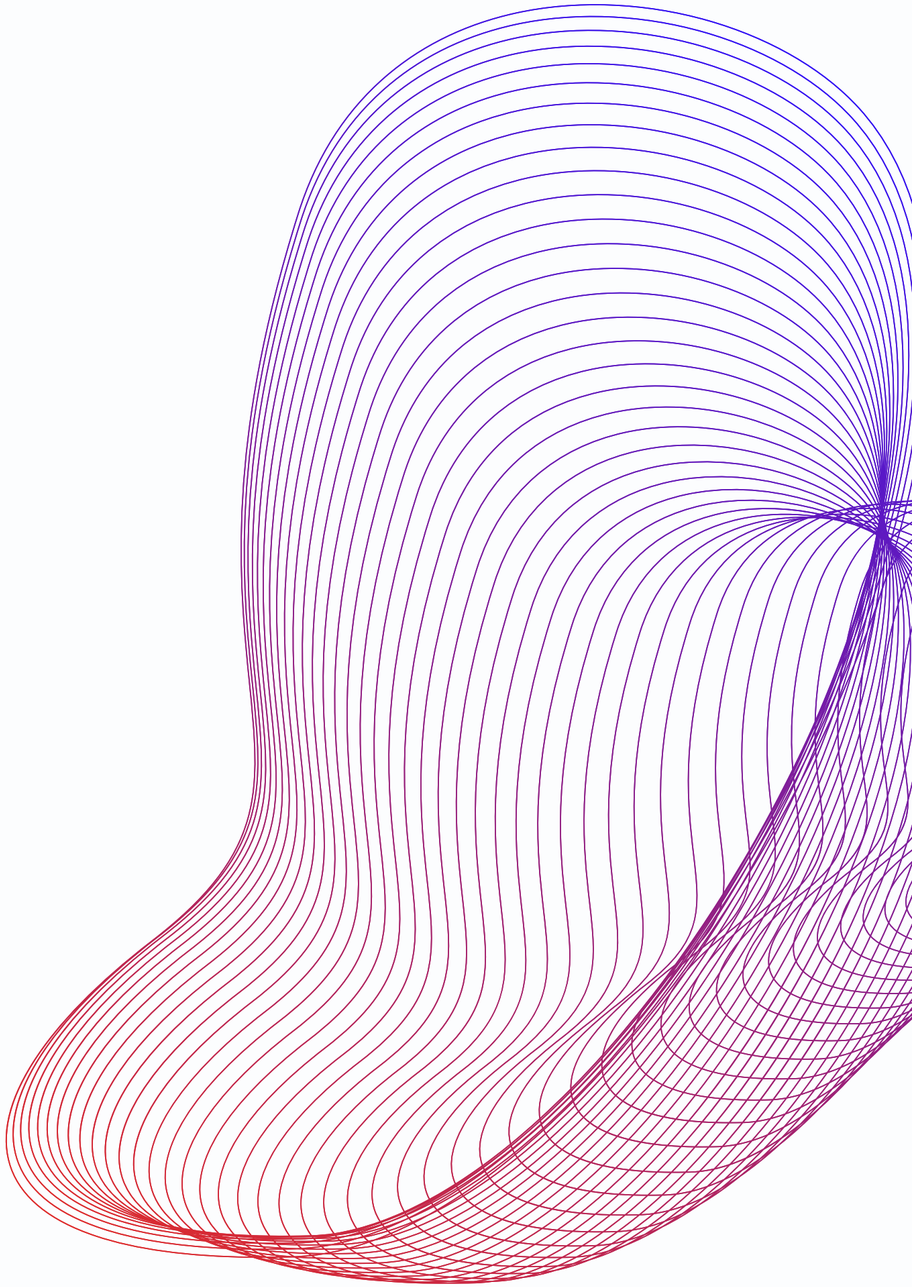

Dark matter at particle colliders: searches in the ATLAS experiment

Jared Little
Laboratoire d'Annecy De Physique Des Particules (L.A.P.P.)

Supervised by:
Tanya Hrn'ova and Stephane Jezequel (LAPP),
Caterina Doglioni
(University of Manchester and Lund University)

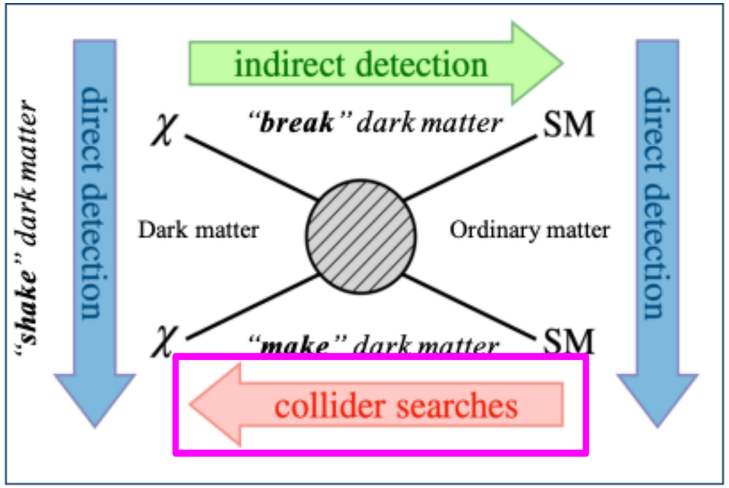

DM Science Project - ATLAS

The ATLAS Experiment, along with CMS, are two general purpose detectors located on the Large Hadron Collider.

Wide range of physics investigated:

- Higgs discovered in 2012.
- Precision measurements on Standard Model properties.
- Searches for new physics, including particles that make up dark matter.

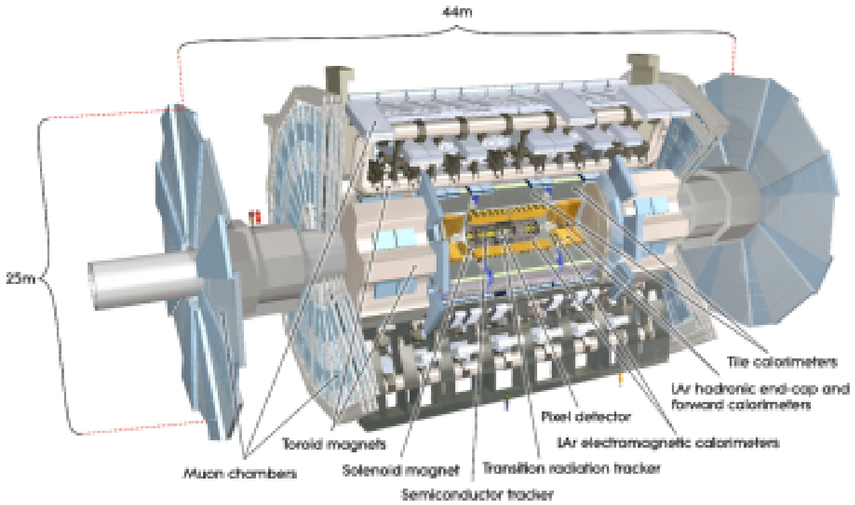

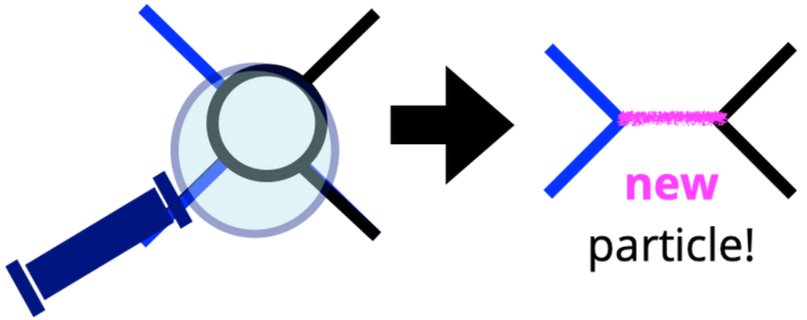

At the LHC, we are trying to “make” dark matter.
By probing the interactions with ordinary matter, we can better understand the nature of DM.

Inclusive Dilepton Resonance Search

Looking for a bump (= new particle) over the background of known particles

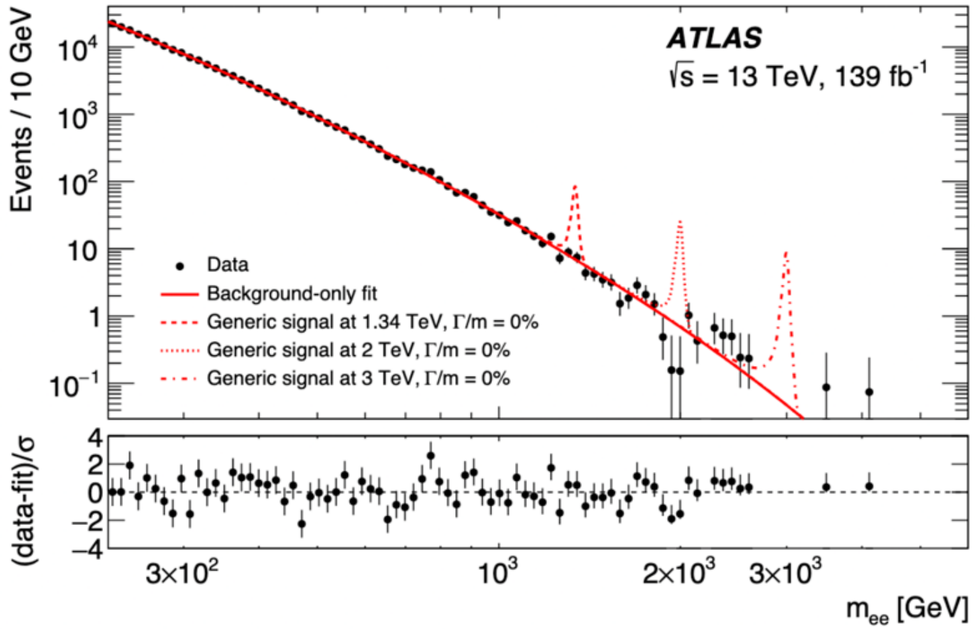

DM mediator decays in two electrons
→ search in di-electron final state

- No signal → constraints on the fiducial cross-section of a new Z' particle.

Two projects within this TSP:

1. Reinterpretation of inclusive resonance search in terms of dark matter mediators ✓
2. Exclusive Z'+MET analysis

Reinterpretation of the Resonance Search

Use the dilepton resonance search to constrain dark matter mediators.

- Assuming a non-zero coupling to leptons, a neutral mediator associated with a dark sector would produce an excess in the dilepton invariant mass distribution.

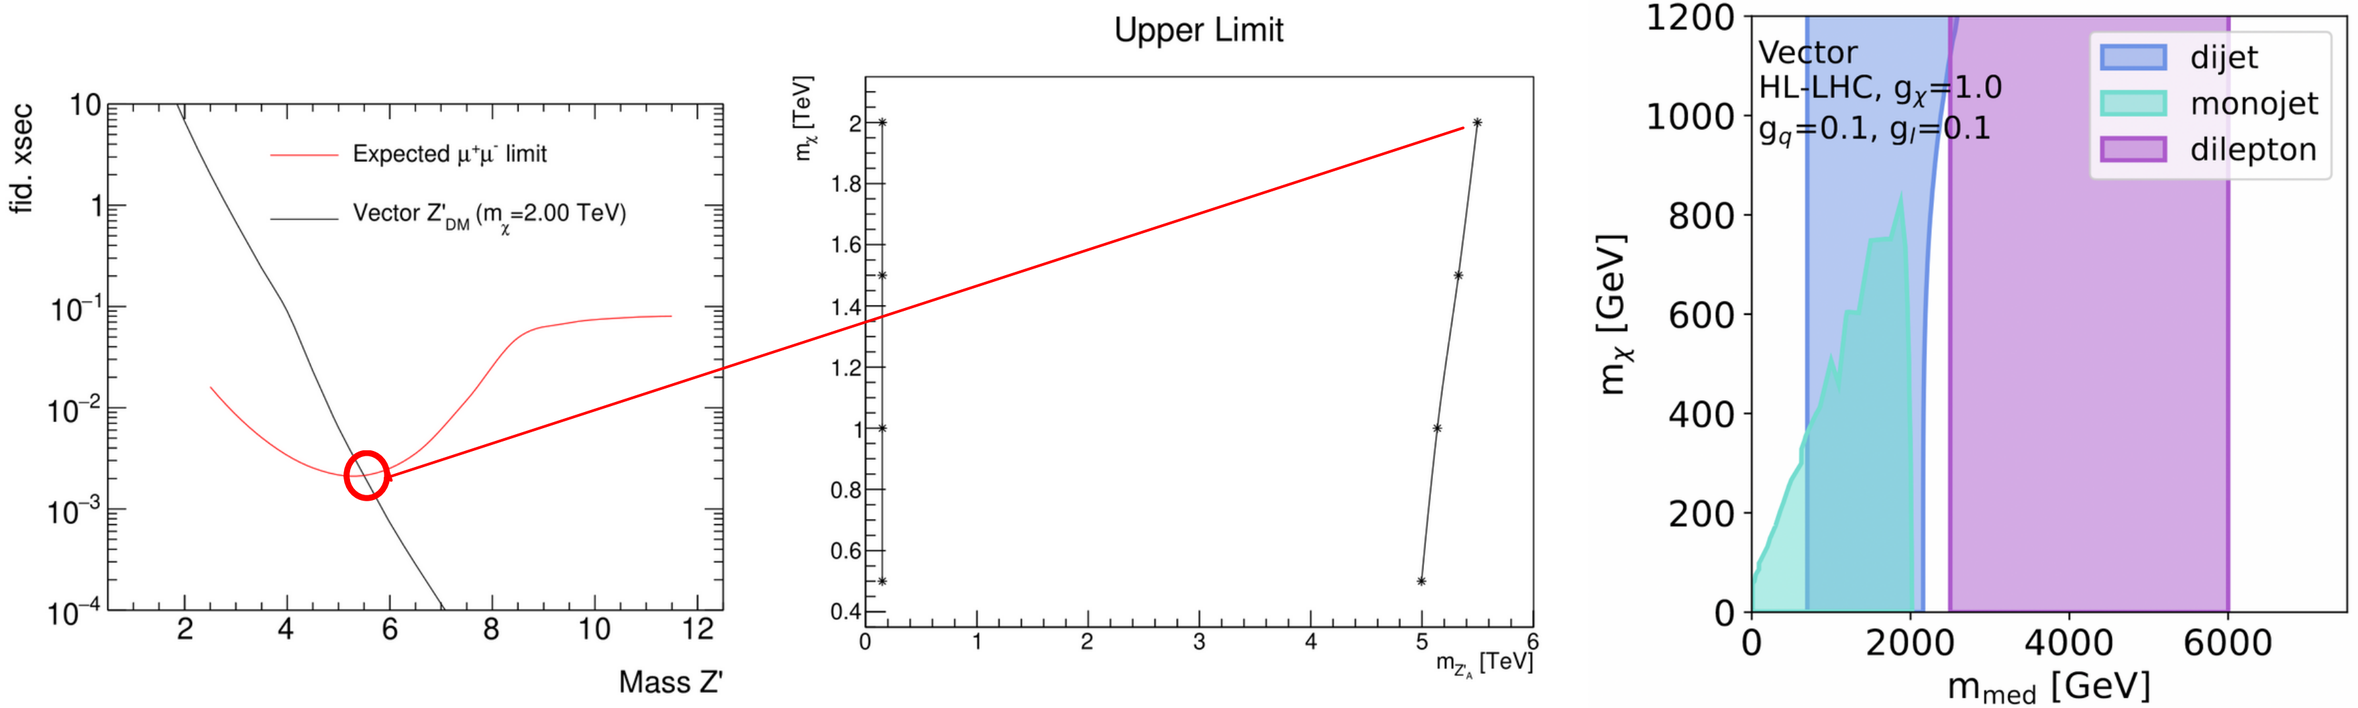

Results included in this paper: https://arxiv.org/abs/2206.03456
(prepared within the US prioritization effort "Snowmass")

Reinterpretation of the Resonance Search

This reinterpretation was set up with **REANA**, sending the jobs to a remote computer from the **VRE**.

- Multiple stage workflows can be sent, passing the output to the following stage.

reana

Reproducible research data analysis platform

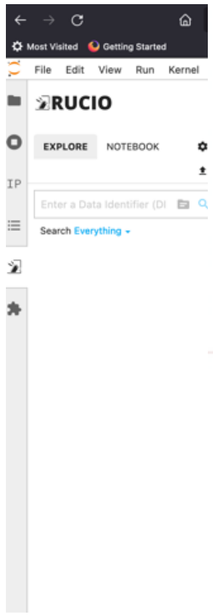

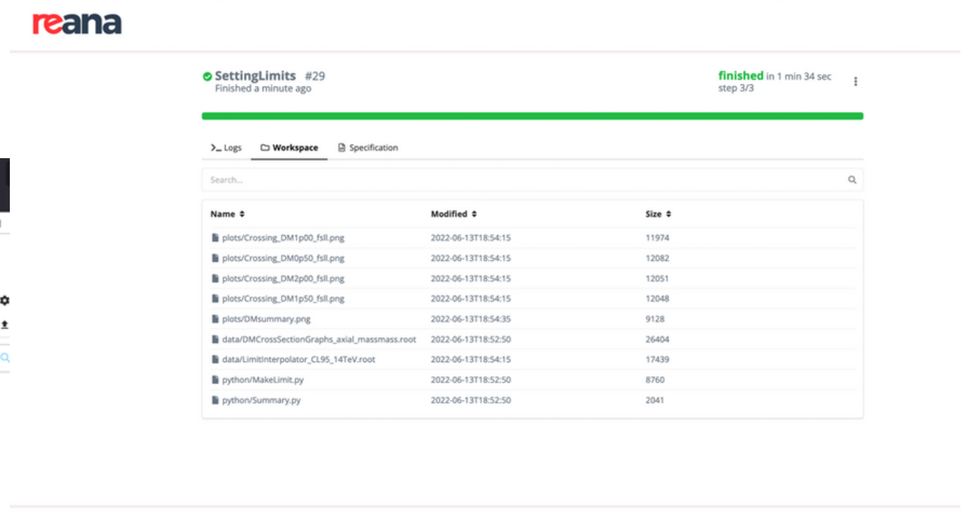

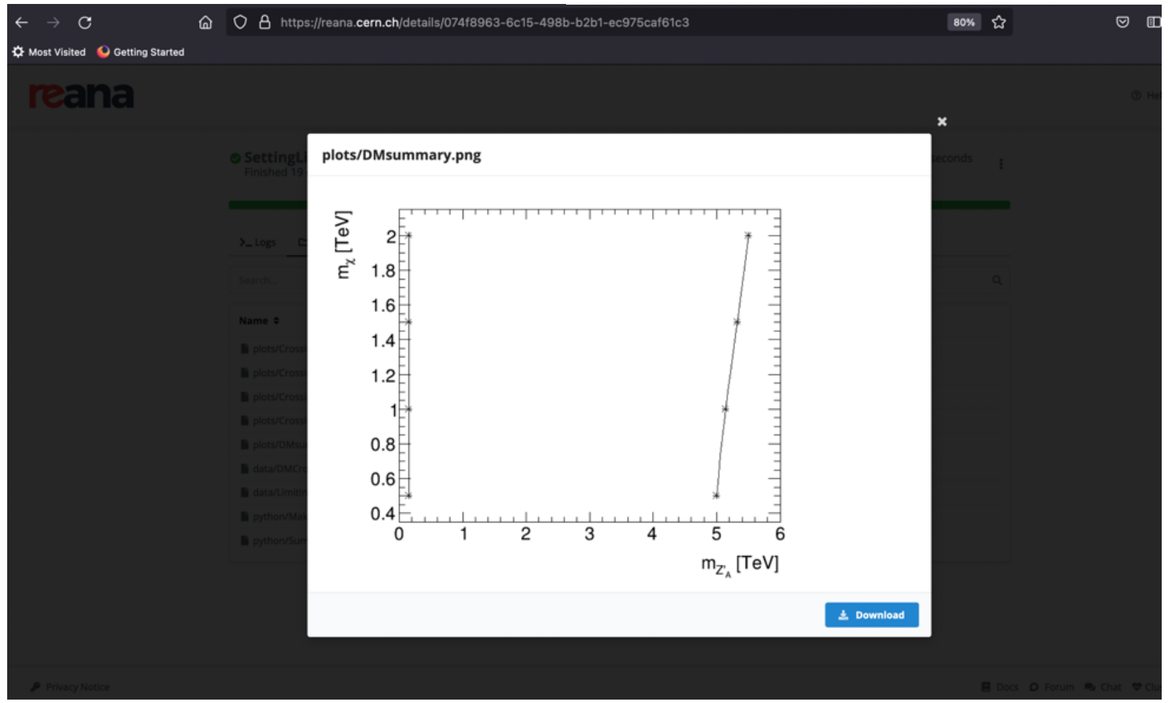

Exclusive Z'+MET Analysis

Search targeting dilepton resonances in the ll+MET final state.

- Searching for well-motivated models that could have escaped detection up to this point.
- Benchmark models help guide our analysis techniques, but we aim to stay as general as possible.
  - **Reproducible and reinterpretable results** are necessary for **collaboration.**
- By targeting dilepton events with MET in the final state, we will be more sensitive in the low-mass regions where the dilepton analysis was dominated by Standard Model events.
  - **Results expected soon.**

Dark Higgs

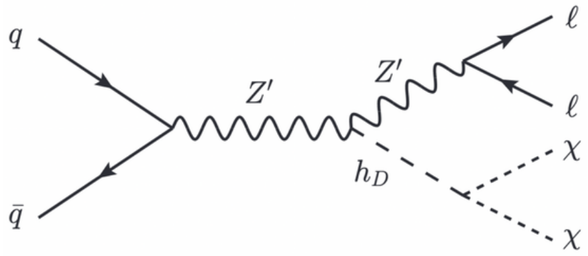

Light Vector

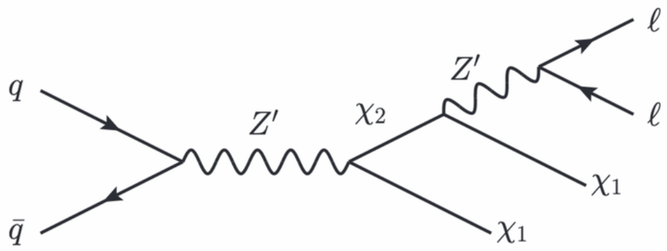

https://arxiv.org/pdf/1504.01386.pdf

Exclusive Z'+MET Analysis

DxAOD

Calculate Kinematic Variables

ntuples

Selections / Validation

ntuples/histograms

Statistical Analysis

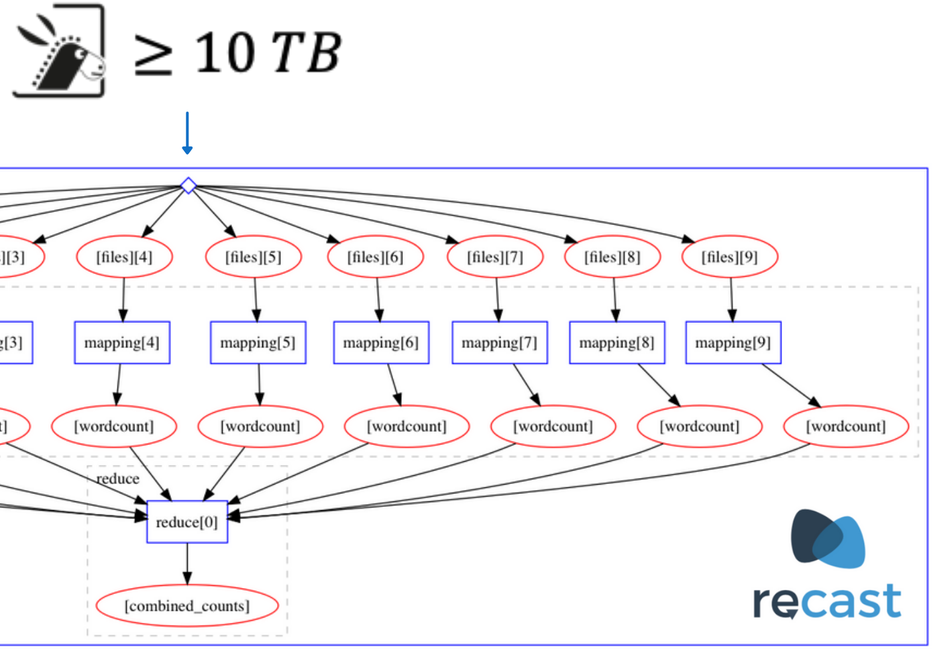

Dark matter Direct Detection: Darkside plans and results

Maria Adriana Sabia (INFN/La Sapienza)
Paolo Salomone (INFN/La Sapienza)
Marco Rescigno (INFN)
Valerio Ippolito (INFN)

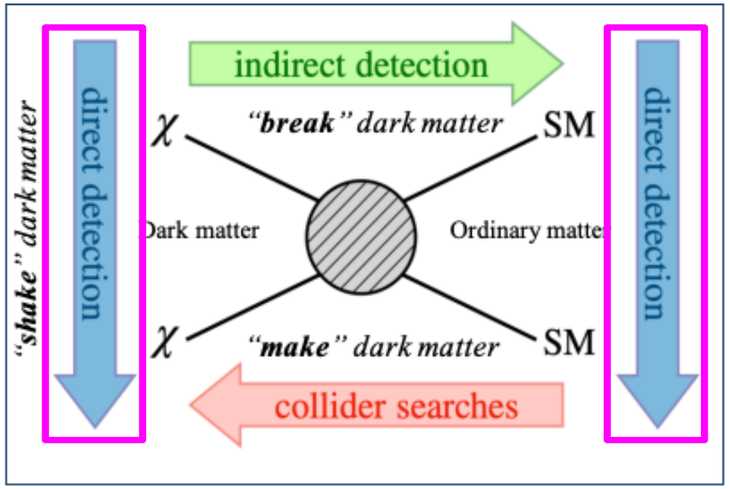

Direct Detection with a LAr TPC

- DM as WIMP-like particle produces a **nuclear** or an **electron recoil**.
- Elastic scattering with Argon Nuclei results in **Scintillation & Ionization**.

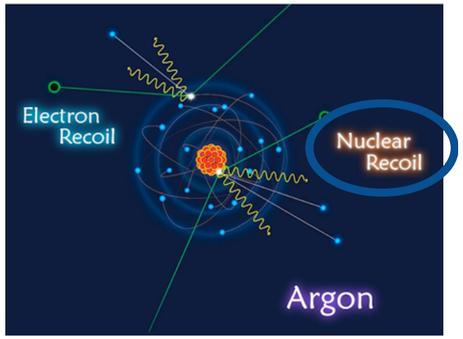

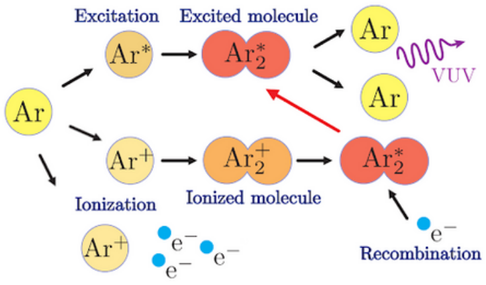

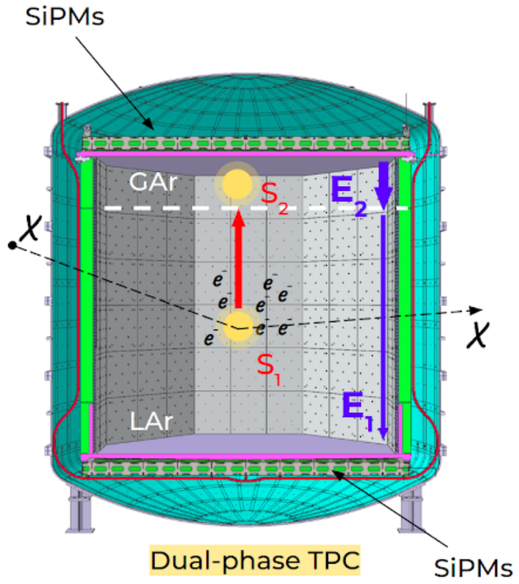

Evolution of the DarkSide Experiment

2012	2013 - 2018	2026 - 2036	2030+
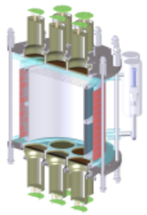	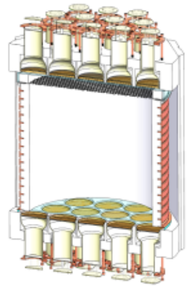	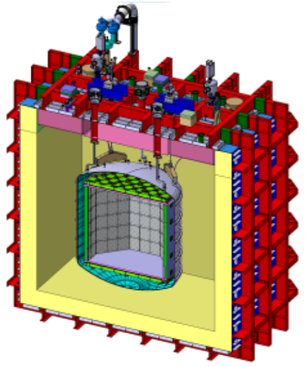	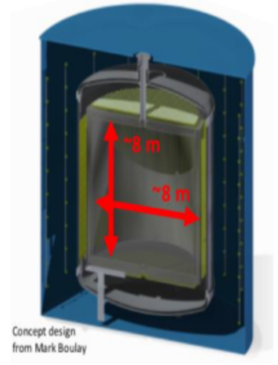
DarkSide-10	**DarkSide-50 16 t d**	**DarkSide-20k 200 t yr**	**Argo 3000 t yr**

DarkSide Plans

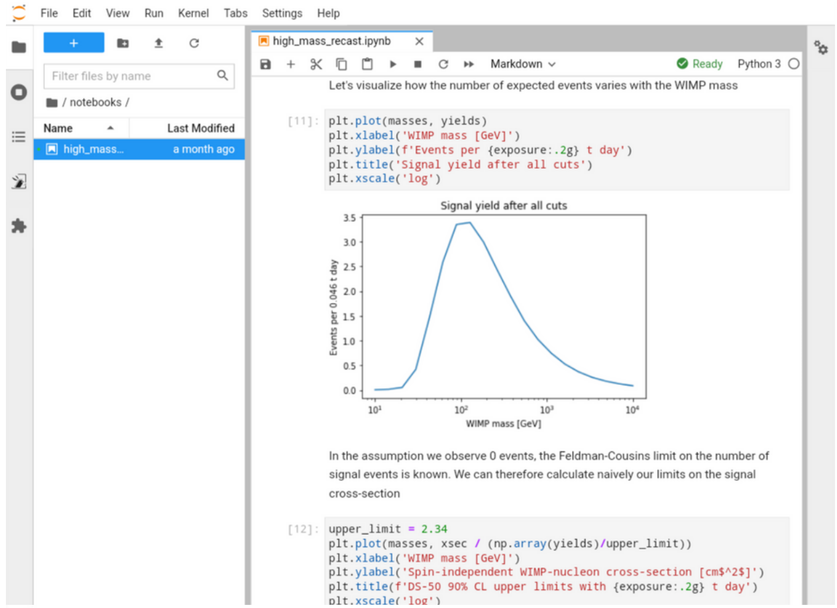

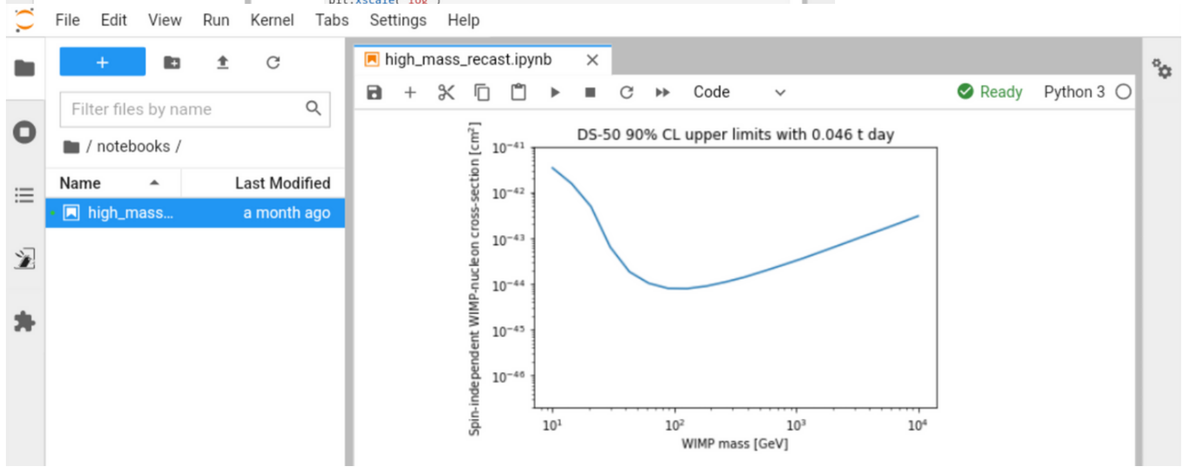

Implemented a reanalysis tool for a high-mass search on the VRE platform.

- Output: **DarkSide50** exclusion curve for WIMP-nucleon cross section.

Further work is ongoing.

- Low mass analysis to be implemented.
- Different **theoretical models** (WIMP halo, argon response…) can be inserted by the user to produce different limit results.
- Working towards first open implementation.

Indirect dark matter search with gamma rays

... and its association with the VRE platform via open-science tools

Pooja Bhattacharjee, Christopher Eckner
Laboratoire d'Annecy De Physique Des Particules (L.A.P.P)

Supervised by:
Francesca Calore
Laboratoire d'Annecy-le-Vieux de Physique Théorique (L.A.P.Th)

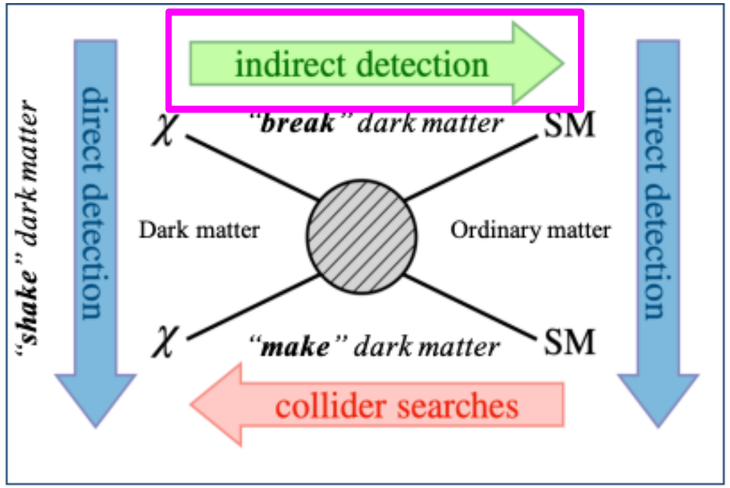

Indirect Searches for Dark Matter

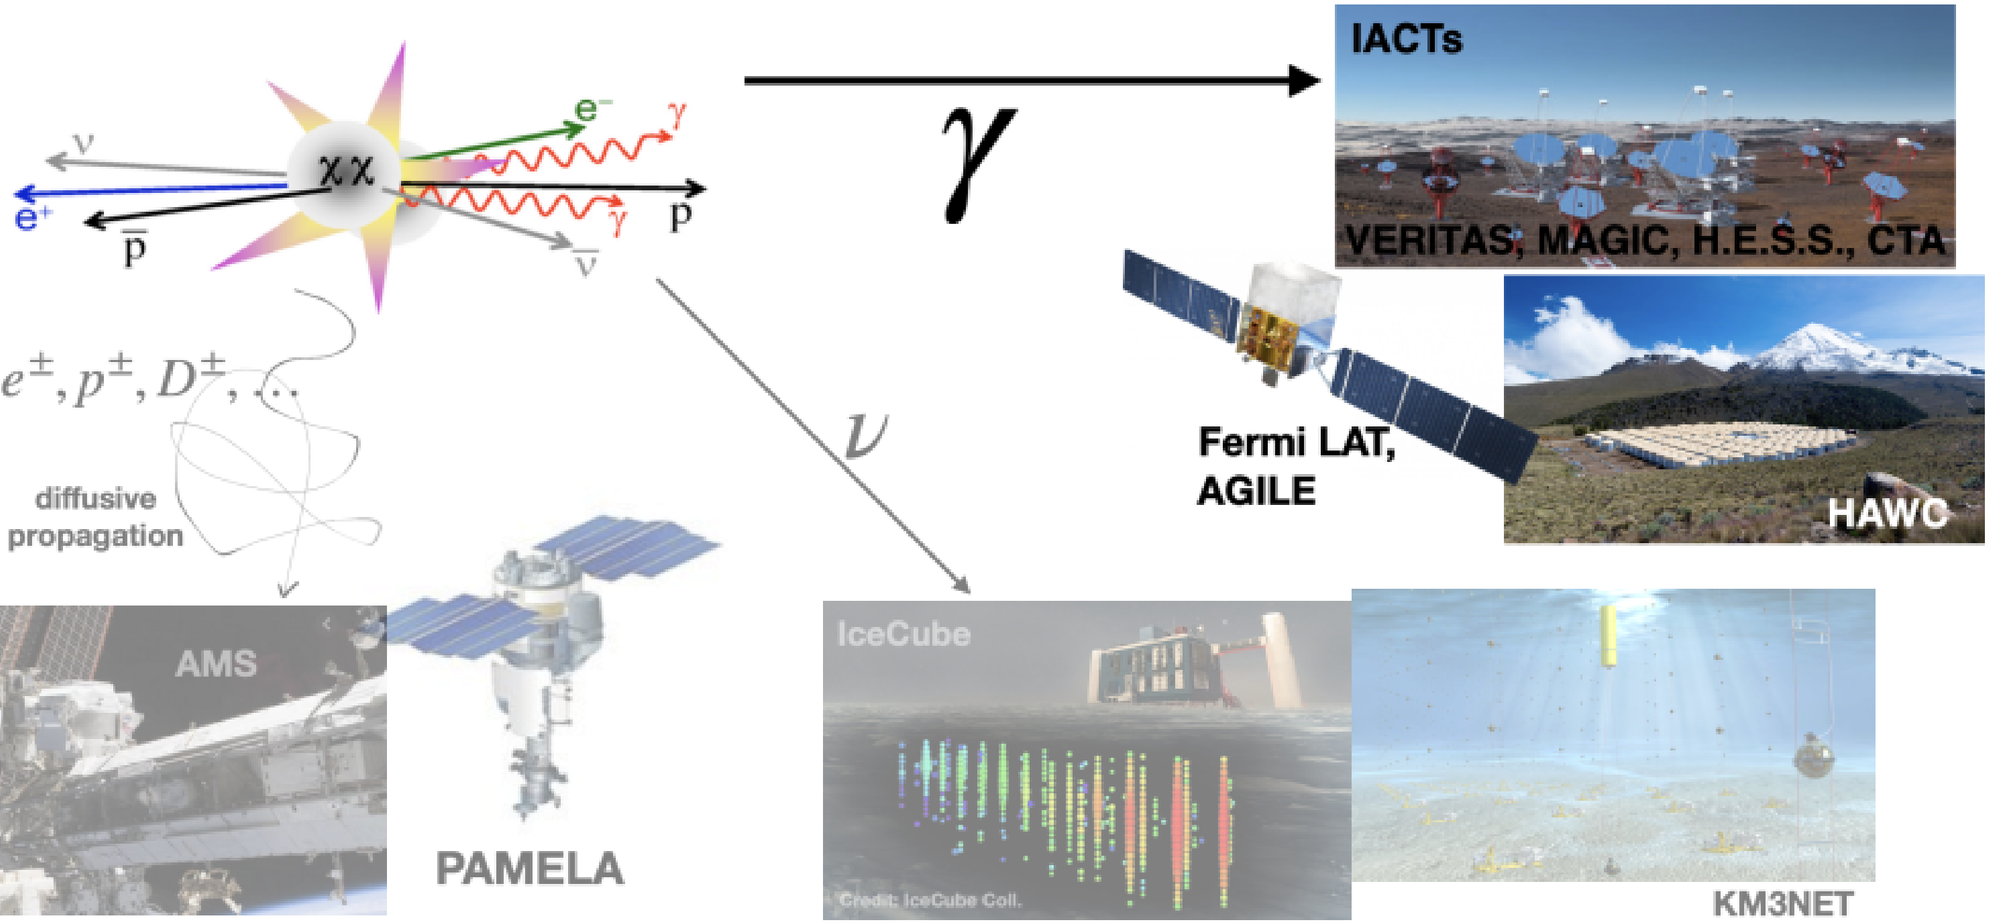

C. Eckner

Fermi Large Area Telescope

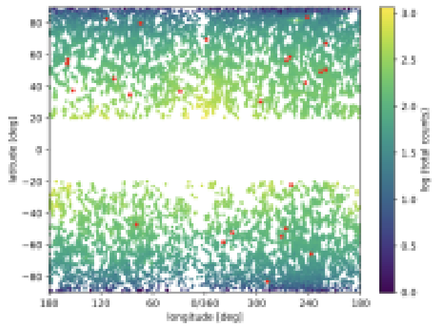

MLFermiDwarfs

Learn to predict the gamma-ray background over the entire sky via training data based on real gamma-ray measurements from Fermi-LAT in a **machine learning based approach**

(performance demonstrated in [F. Calore et al. JCAP10 (2018) 029],[A. Alvarez et al. JCAP09 (2020) 004])

Indirect Dark Matter Detection on the VRE

- The data and main processing software (Fermi Science Tools) are publicly accessible, and now fully available in the VRE.
- Code is entirely written in python 3 using well-known packages like scikit-learn.
- Package can be optimized from the command line enabling a quick check of the viability of a user-defined Dark Matter model.

Fermi Large Area Telescope

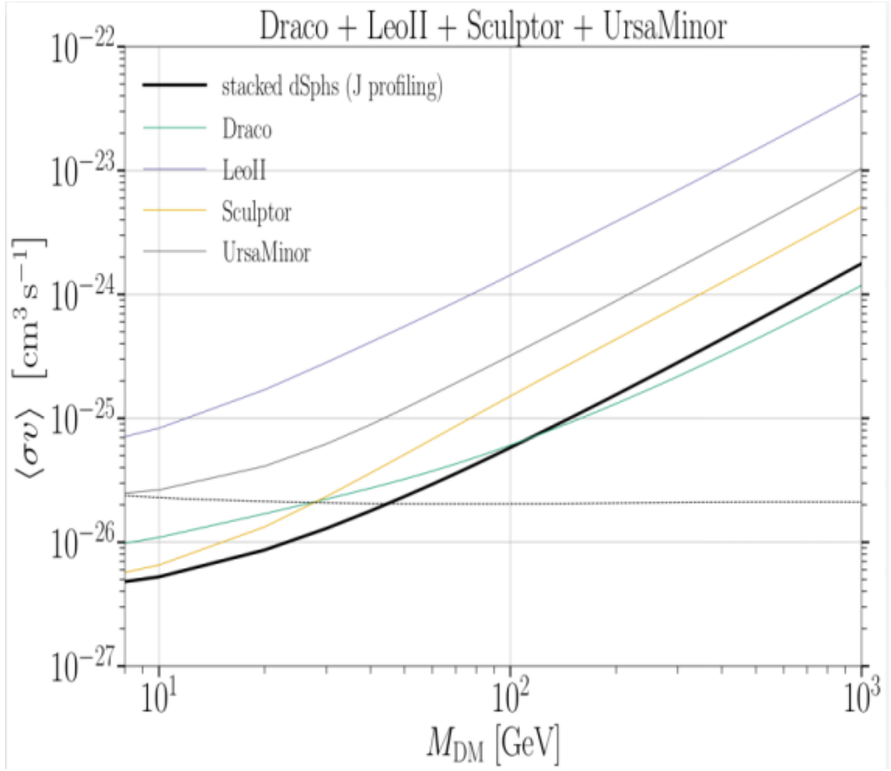

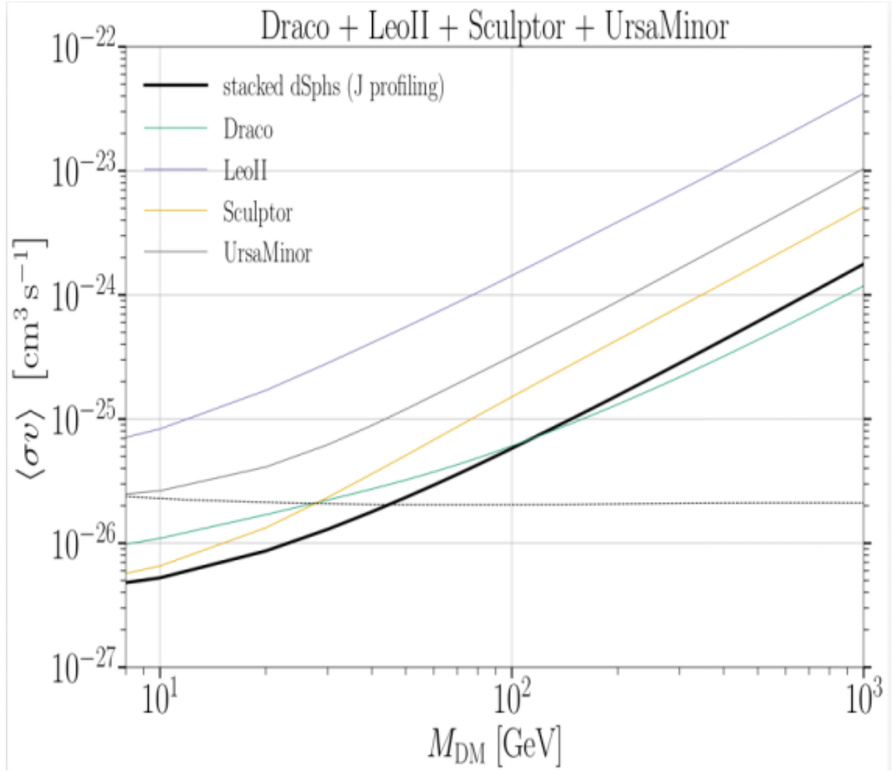

MLFermiDwarfs code is accessible at https://gitlab.in2p3.fr/escape2020/virtualenvironment/mlfermilatdwarfs

Brown Dwarf Analysis

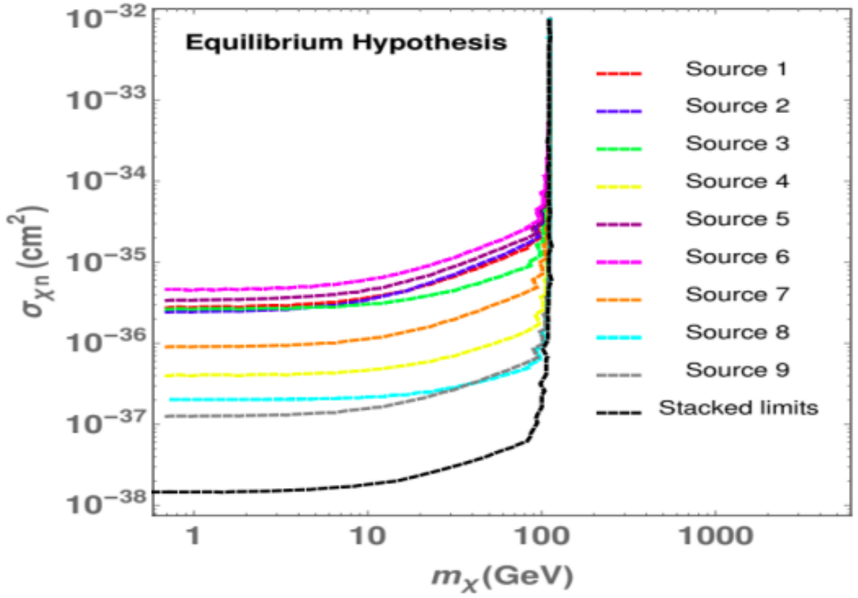

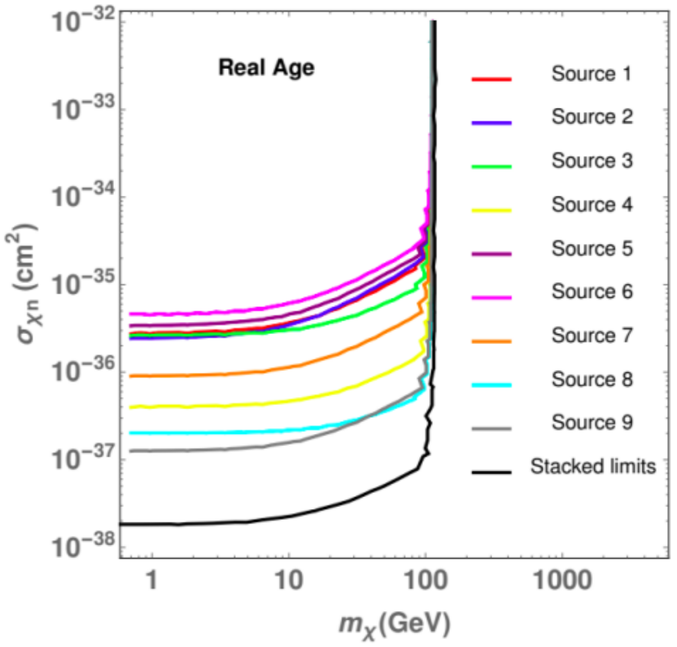

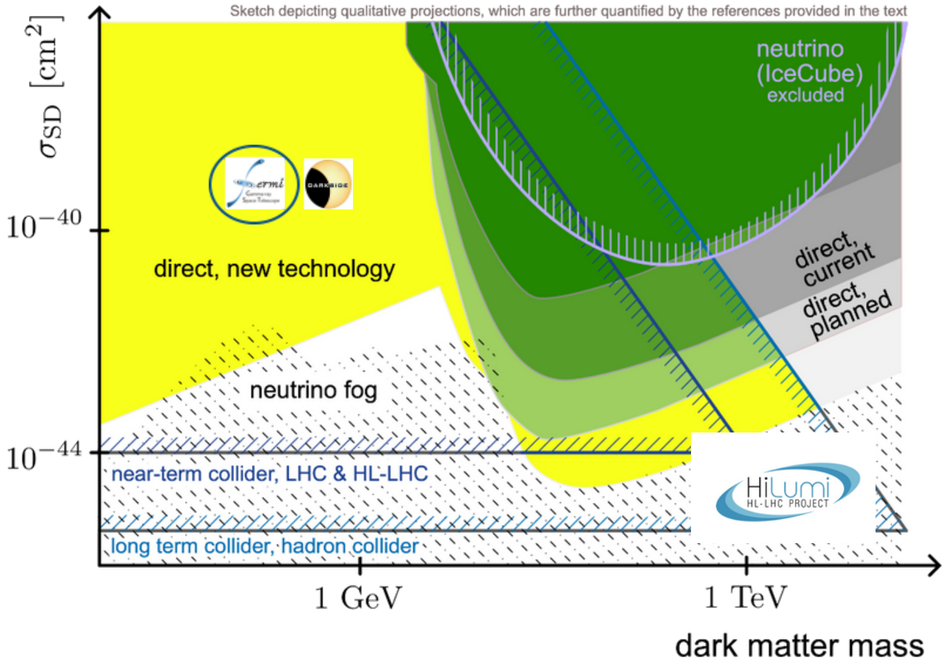

Based on the recent Published paper on **Bhattacharjee et.al, PRD,107, 043012, 2023.**

Code is accessible at https://gitlab.in2p3.fr/escape2020/virtual-environment/brown-dwarfs-gamma

Instrument Response Function of KM3NeT for point-source analysis

Mikhail Smirnov
(Friedrich-Alexander University FAU-ECAP)

Supervised by:
Kay Graf
Friedrich-Alexander University FAU-ECAP

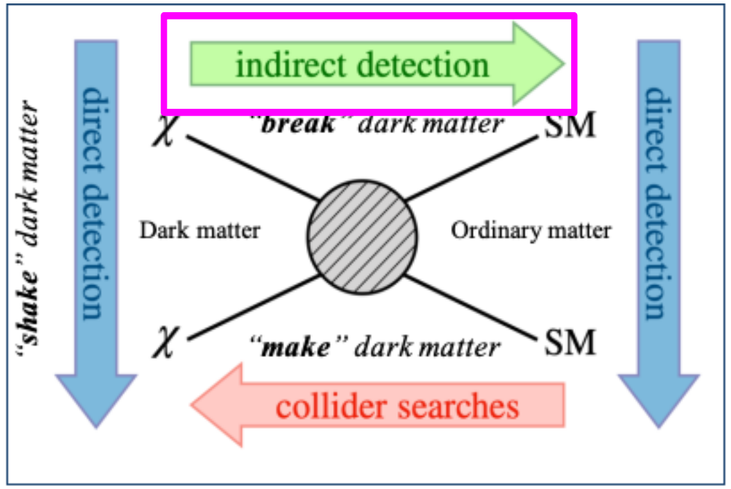

KM3NeT Detector

KM3NeT (cubic kilometer neutrino telescope) **J.Phys. G43 (2016) 084001**

KM3NeT/ARCA (**A**stroparticle **R**esearch with **C**osmics in the **A**byss)
discovery and observation of HE cosmic neutrino sources
(Eν ~ GeV-PeV) high energy neutrinos
Depth – 3500 m – offshore Sicily (Italy)

KM3NeT/ORCA (**O**scillation **R**esearch with **C**osmics in the **A**byss)
determination of the neutrino mass hierarchy
(Eν ~ MeV - GeV) low energy neutrinos
Depth – 2500 m – offshore Toulon (France)

Parameter	ARCA	ORCA
DU distance	90 m	20 m
DOM spacing	36 m	9 m
DU height	~ 800 m	~ 200 m
Instrumented mass	2*500 Mton	7 Mton
Amount of DUs	115*2	115

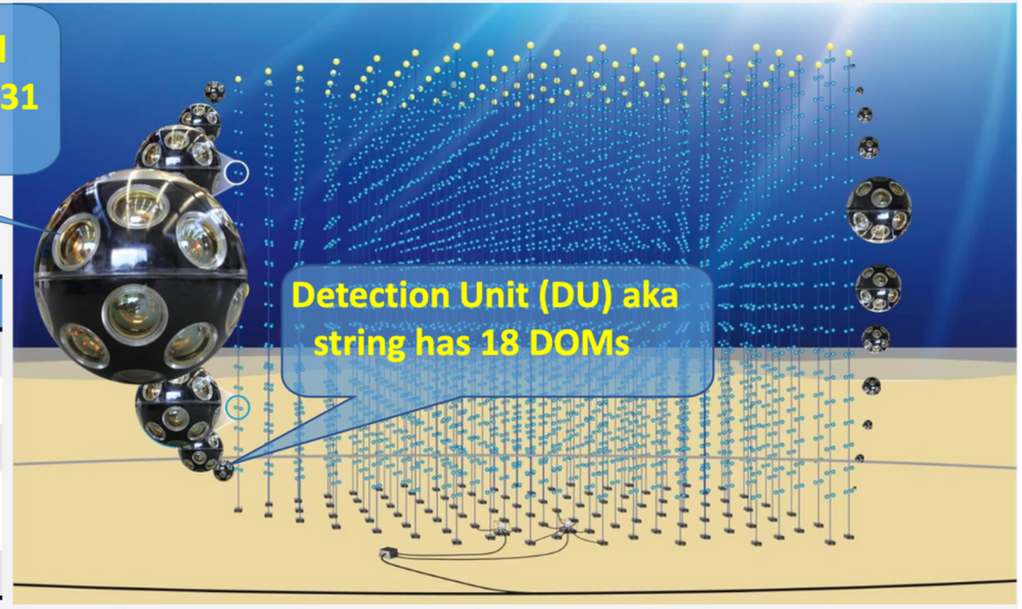

KM3NeT – IRF Concept

Instrument Response Function of neutrino telescope provides a quantitative estimation of the event and background rates.

- Contains physical characteristics of the detector.
- It allows to avoid extensive MC simulations each time for a new configuration of neutrino source.
- It supports different configurations of neutrino sources:
  - Point source with power law E^-a
  - Diffuse source
  - Extended source

- Compatibility with **gammapy** will give an easy combination with other gamma experiments like CTA.
- Active development of the **km3irf** python package.

```
pip install km3irf
from km3irf import utils
new_plot = utils.DrawAeff()
new_plot.peek()
```

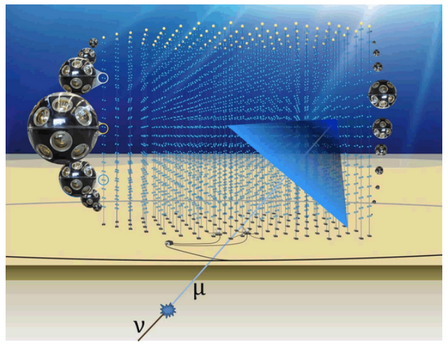

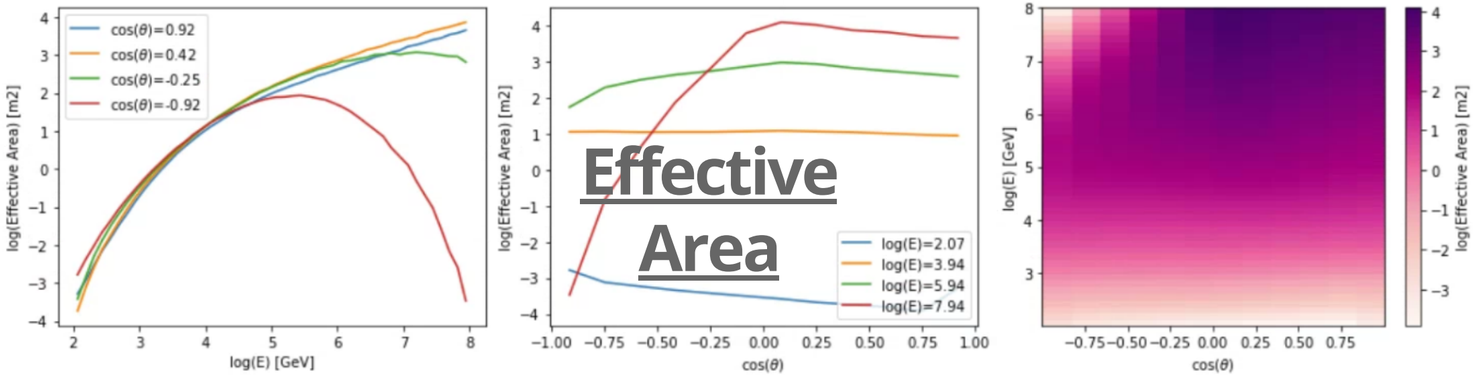

Common tools: Machine learning for big data compression

Axel Gallén, Alexander Ekman
(Lund University)

Supervised by:
Caterina Doglioni
University of Manchester and Lund University

Baler: data compression using ML

Idea behind the Baler compression tool:

- Train autoencoder on scientific (e.g. HEP) data
- Compress/decompress data by storing model + autoencoder's latent space (fewer dimensions)

Idea behind its inclusion in European Open Science Cloud / EOSC Software Catalogue:

- Provide "off the shelf" algorithms/tools that everyone can use

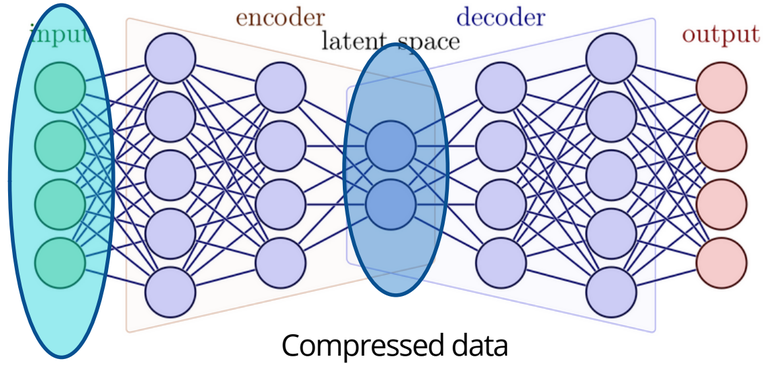

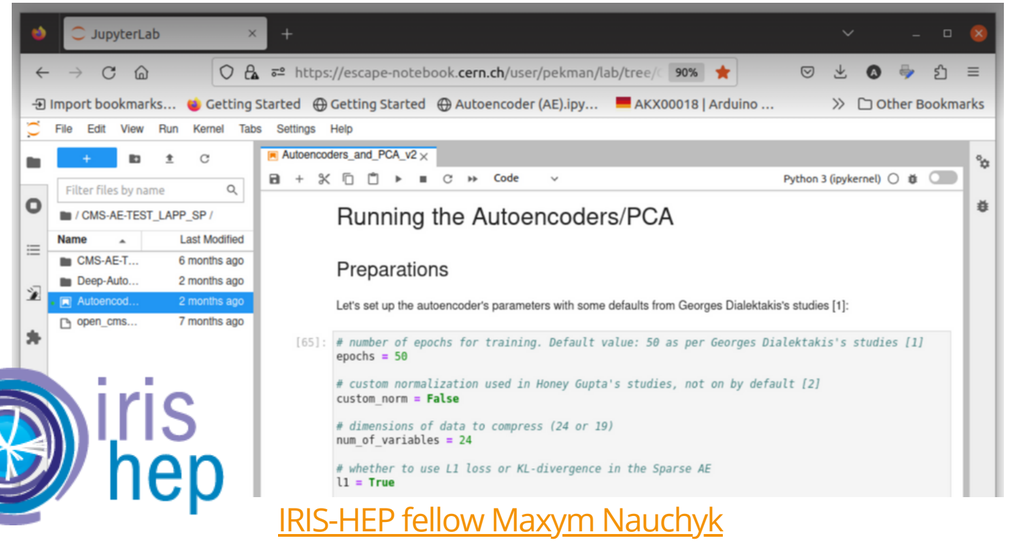

IRIS-HEP fellow Maxym Nauchyk

Status:

- Jupyter notebook containing Baler prototype available on Virtual Research Environment
- Many improvements since (Zenodo release), these will be ported on VRE as well

Demo

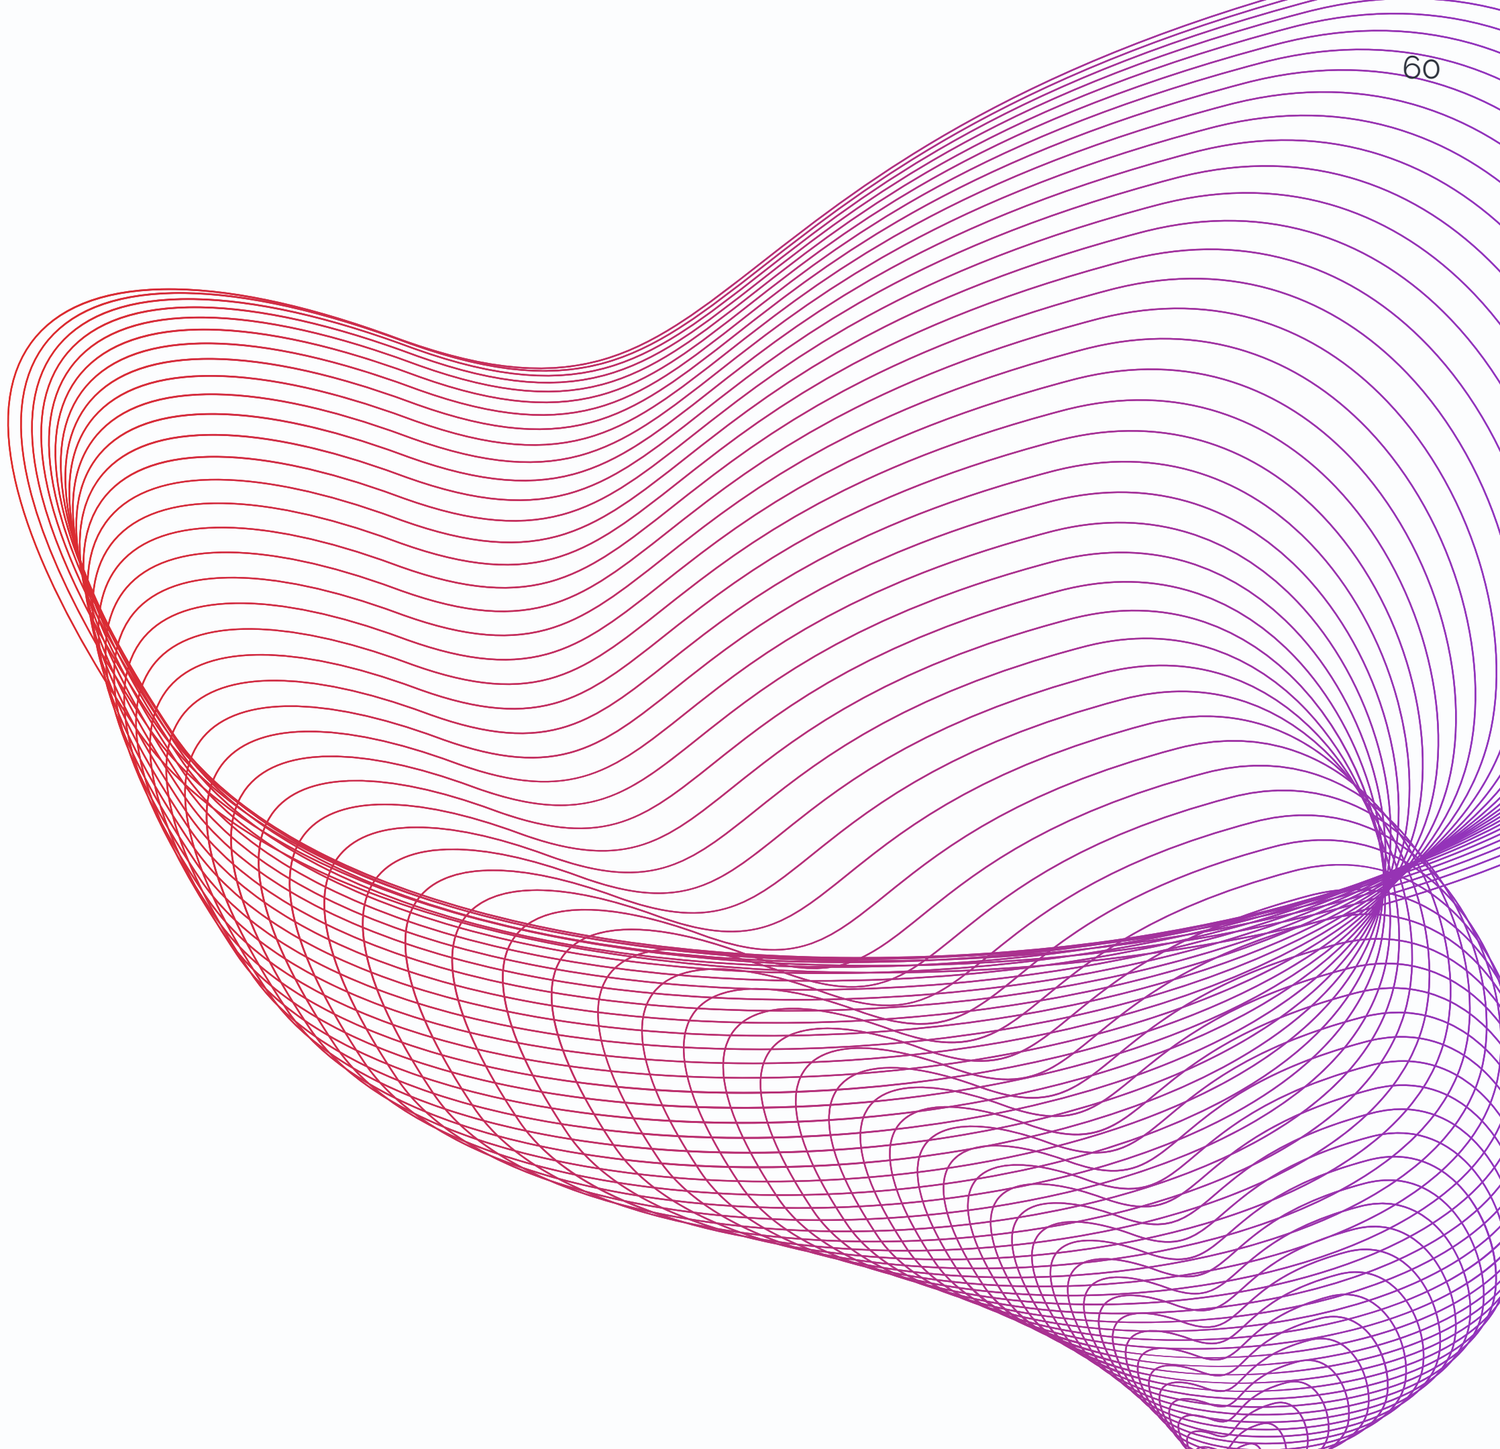

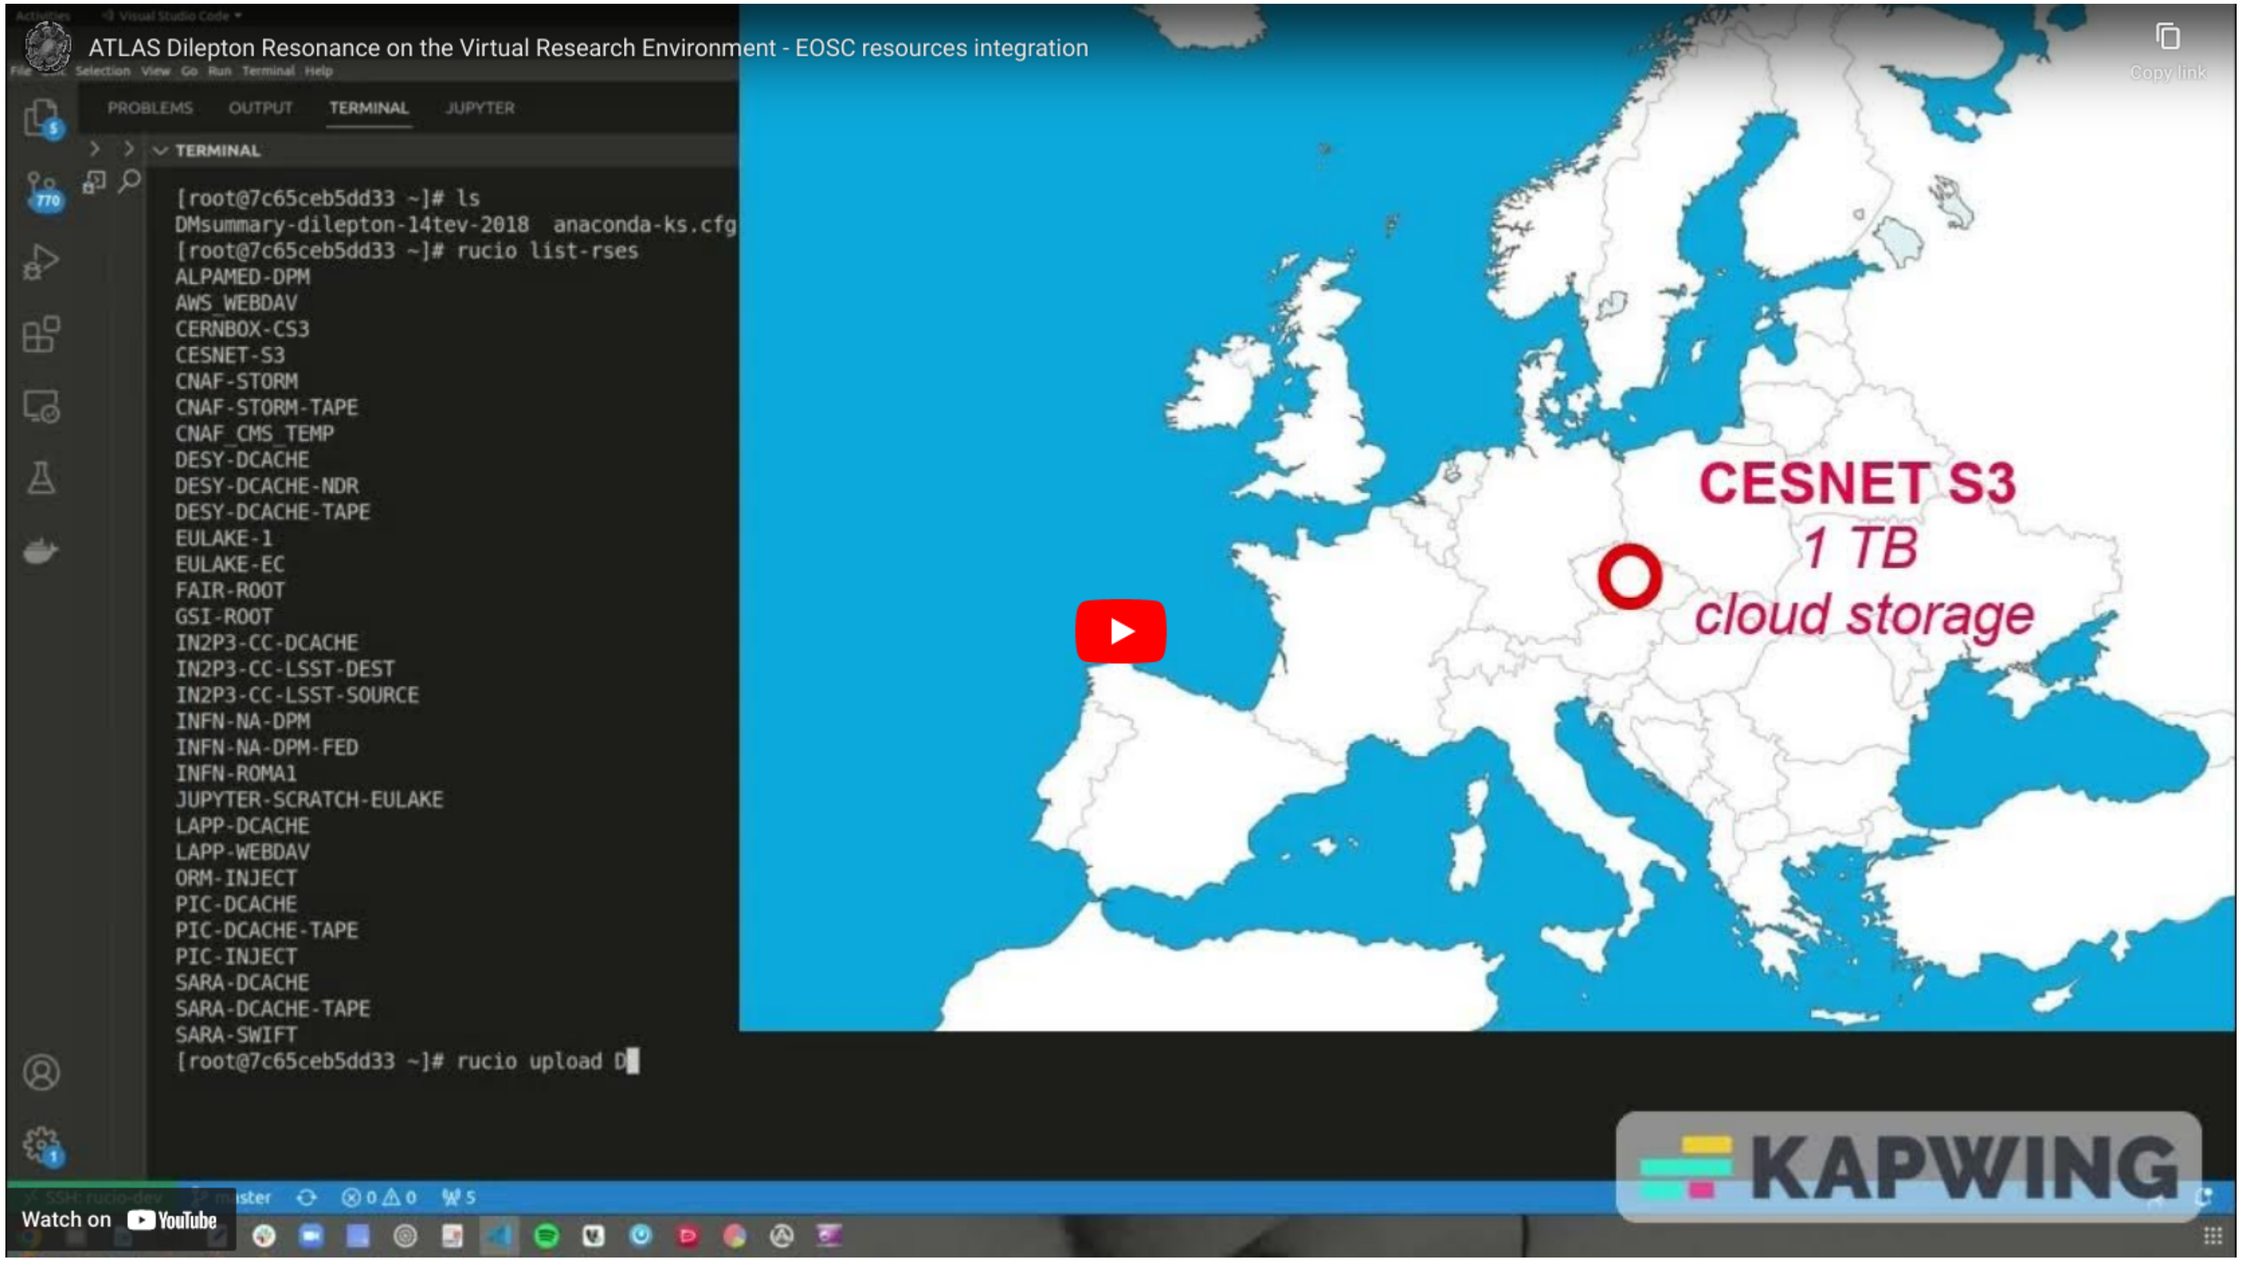

https://www.youtube.com/watch?v=nYp_wsXhKSo&ab_channel=ElenaGazzarrini

Conclusions

It has been a successful journey!

- DM Science Project's analyses and tools on the VRE are providing:
  - **new scientific results** discovering or constraining dark matter hypotheses
  - multiple communities with the necessary understanding to reproduce the analysis
  - possibility to **comparing and combining** results from different experiments
  - **FAIR data and interoperable workflows** as an example for the community
  - **working prototype cell for the European Open Science Cloud**
  - testing ground for software & computing that can be explored by future experiments
- Escape Open Collaboration Agreement ensures the collaboration and joint common activities across scientific communities in the development of VREs
- EOSC resources have been successfully integrated
- 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 outlook

Future plans include

- Consolidation of EOSC Future Science Projects
  - widening participation of scientists to Open Science tools
  - onboard new analyses requiring o(TB) data
  - guarantee restricted data access until embargos lifted
  - expand use cases to real-time analysis and more complex workflows on constrained infrastructure
  - ensure the sustainability of the VRE infrastructure
  - strengthening cooperation and sharing experience across scientists
  - publish all software and pipelines on ESCAPE Software Catalogue
  - use Gambit software for combination of results
- 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)
- Caching data on distributed storage on the VRE for faster data access

Thank you!

Questions?

e-mail

elena.gazzarrini@cern.ch

jared.little@cern.ch

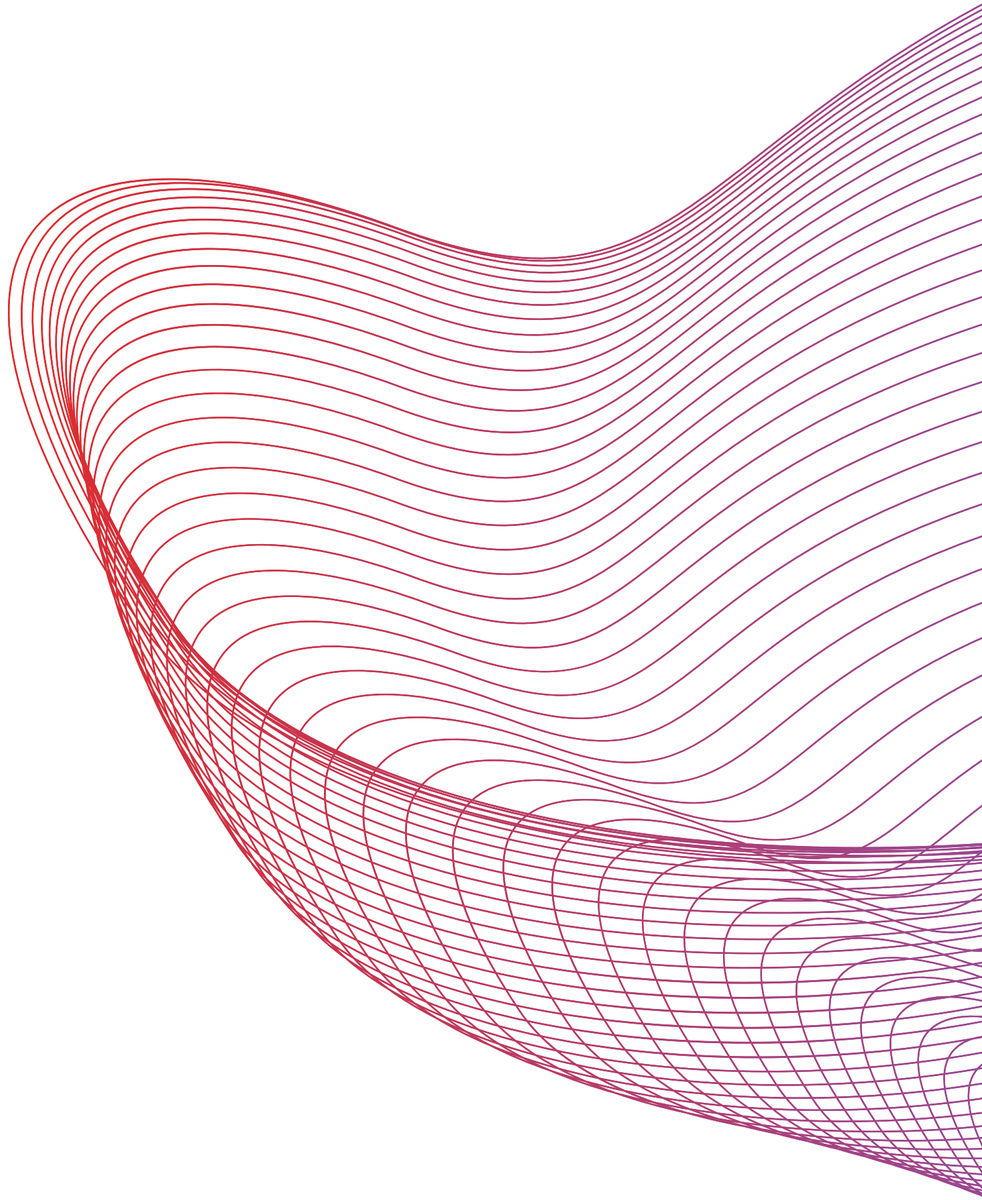

Behind the scenes...

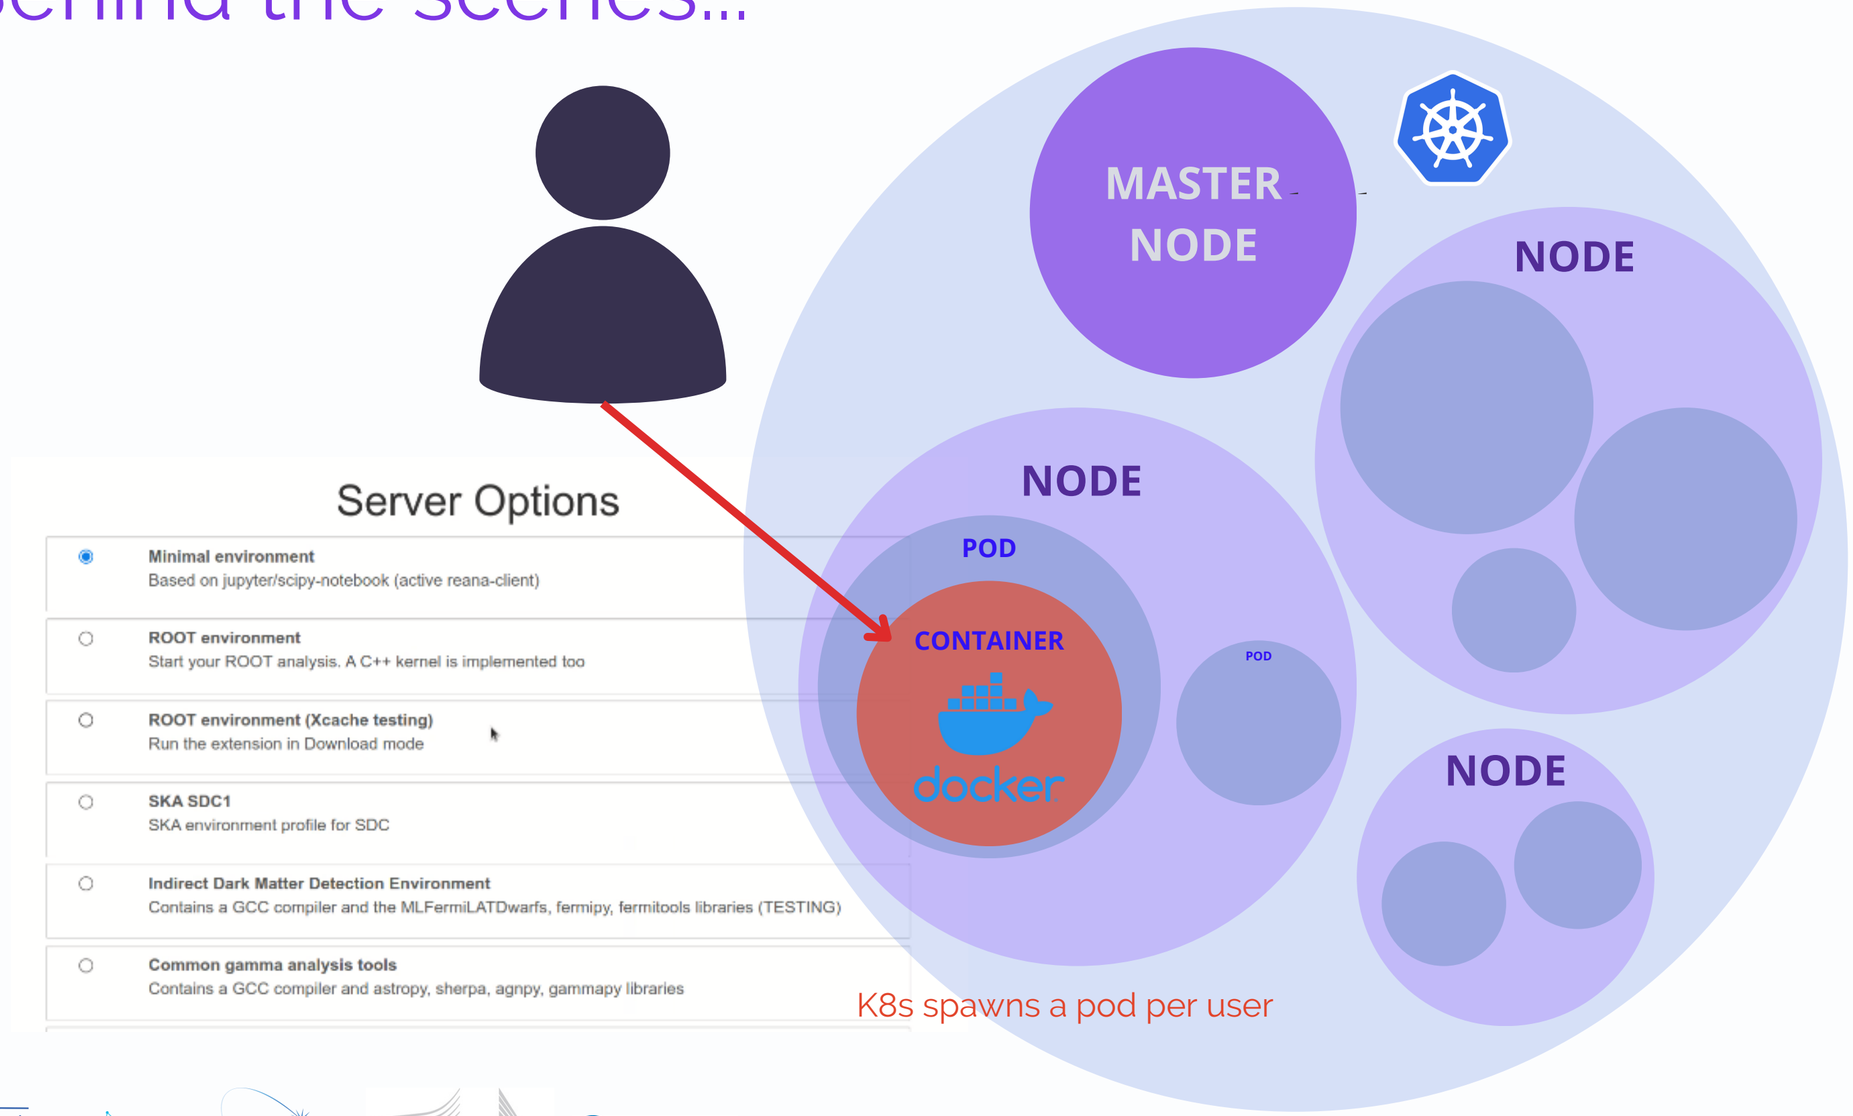

Behind the scenes...

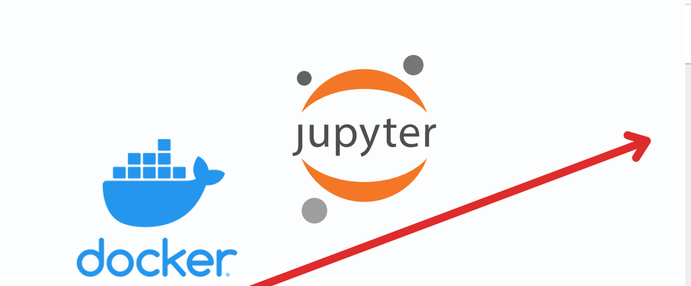

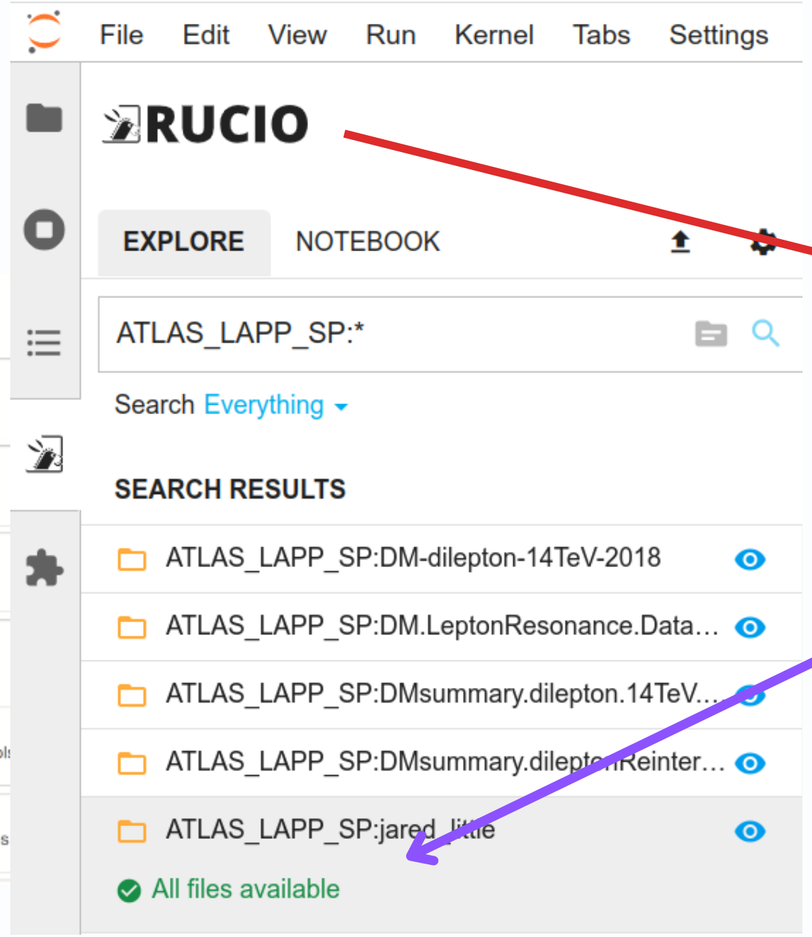

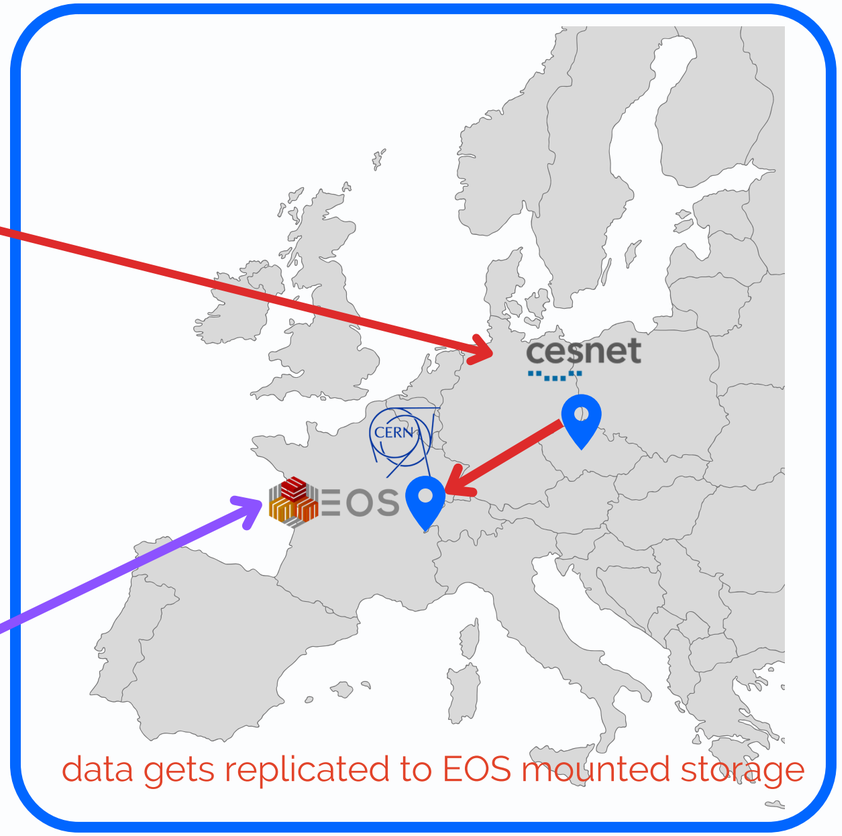

Behind the scenes...

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/
```

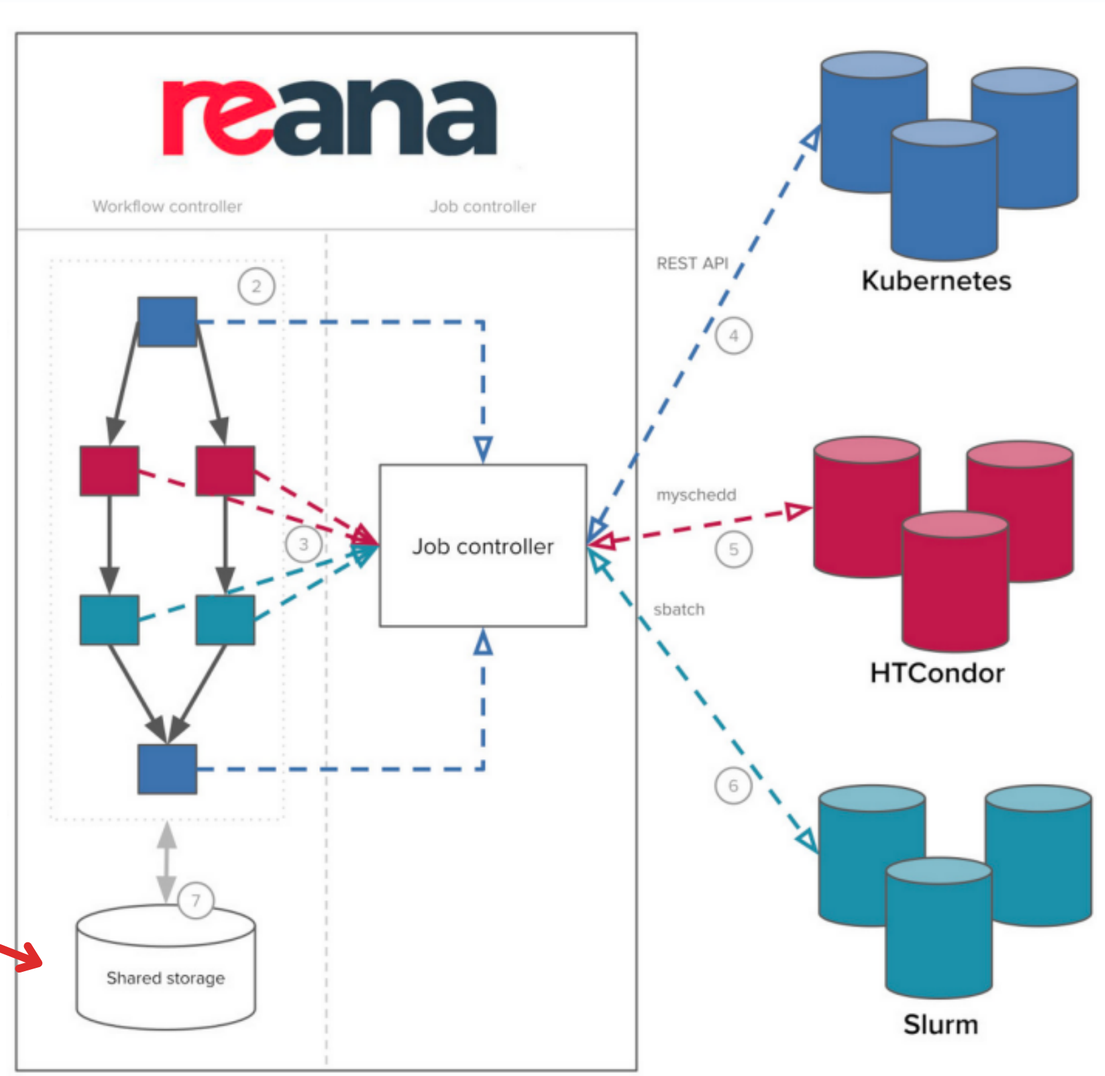

data gets downloaded on Reana storage

Back up

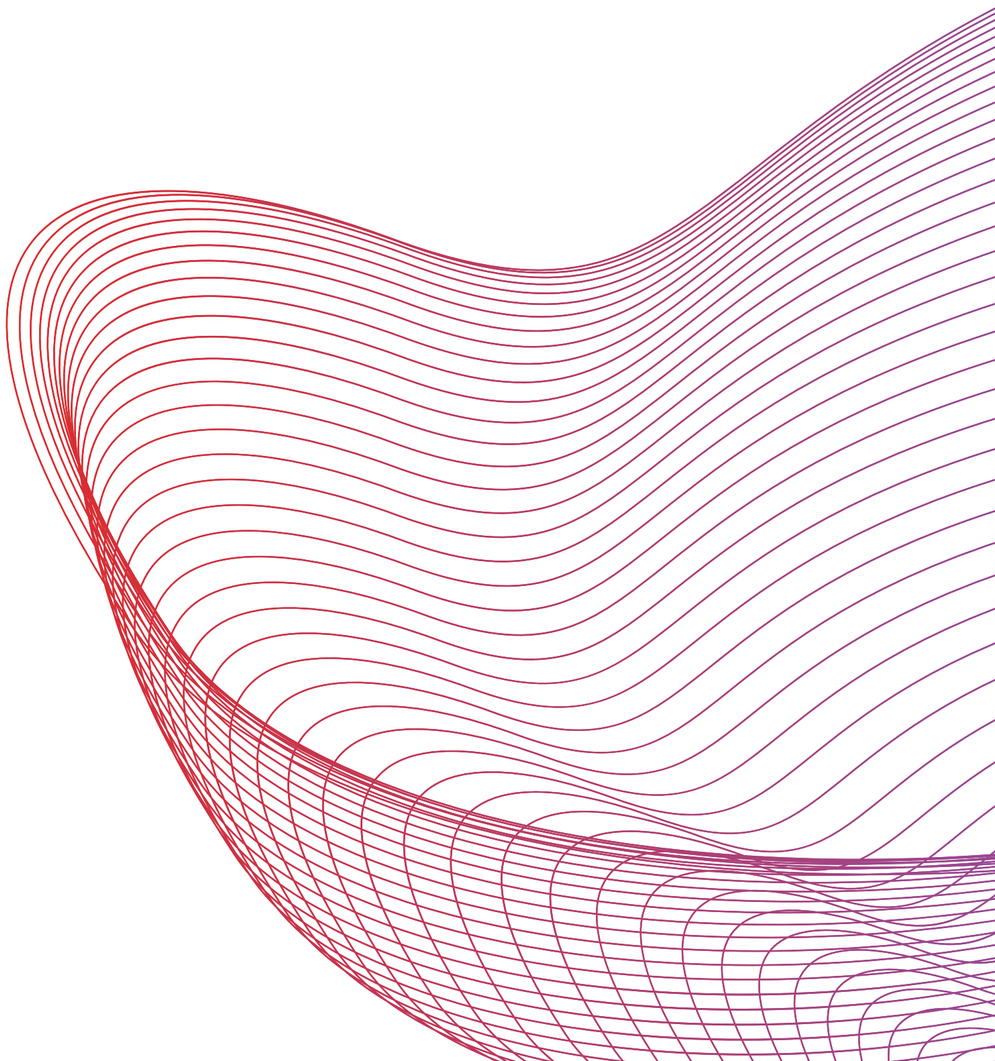

Status

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