

# Rucio CERN IT Community report

Hugo Gonzalez Labrador  
Storage Engineer  
CERN IT

7th Rucio workshop, SDSC, 30/09/2024





# Outline

Where did we start?

Where are we?

Where are we going?

How was so far?

Future

# Where did we start?

## Have you failed to reproduce a result?



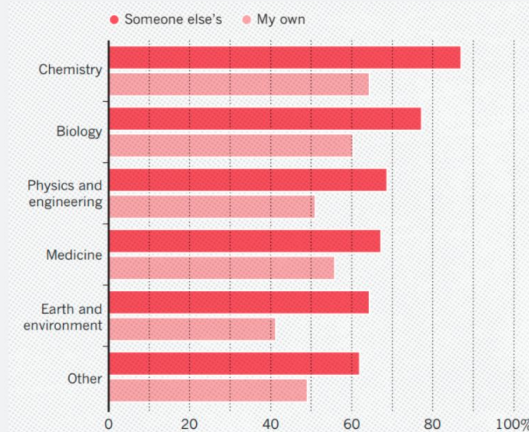
**70%**

researchers tried and failed to reproduce others' results



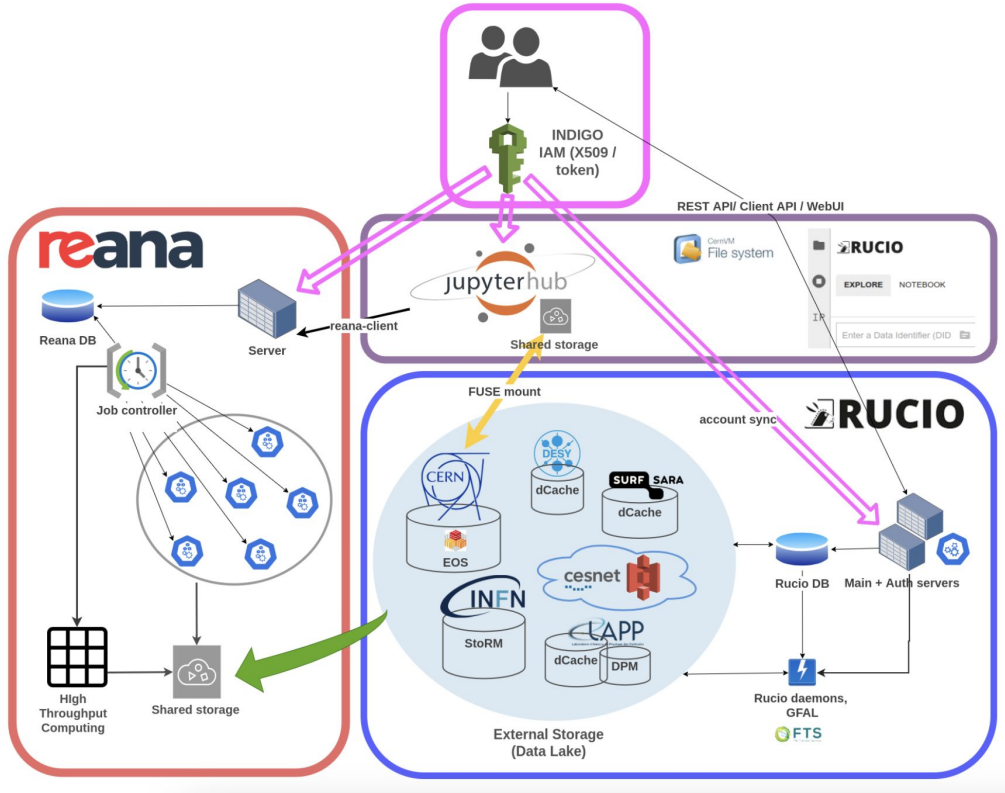
**> 50%**

researchers failed to reproduce own results



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

# Where did we start? VRE hub



A Virtual Research Environment composed of

- AAI
- Rucio Data Lake
- Reana cluster
- Notebook service

Rucio is the chosen technology to facilitate data management across the data lake

VRE very much alive today and used by new members willing to try out!



# Where are we?

The screenshot shows the RUCIO web interface. On the left, there's a sidebar with 'RUCIO' branding, 'EXPLORE' and 'NOTEBOOK' tabs, and configuration options for 'Active Instance' (ESCAPE), 'Rucio Authentication' (X.509 User Certificate), and 'X.509 USER CERTIFICATE' (Certificate file path and Key file path). The main area displays a Jupyter Notebook with the following code and output:

```
[12]: print(test_zoom)
      a = open(test_zoom)
      a.read()

/home/jovyan/rucio/ESCAPE/downloads/orsxg5djnzttu5dfon2f6ztjnrsv6ztpojpwk43boa/testing/test_file_for_esap

[12]: 'Hello zoom!\n\n'

[2]: atlas_gamgam2

[2]: /home/jovyan/rucio/ESCAPE/downloads/mf2gyllthjwqxztgq2tgmjyfxlksbrgi2uu
x2nfxgg3c7m5qwz2z3bnuxeoylni5qw2ltsn5xxilrr/atlas/mc_345318.WpH125J_Wincl
_gamgam.GamGam.root.1

[3]: mariotest


[3]: /home/jovyan/rucio/ESCAPE/downloads/mf2gyllthjwqxzrgeydsmbtfznfa4tjnvstc
mbqaxhe33poq/atlas/mc_110903.ZPrime1000.root

[10]: !rm -rf ~/rucio

[ ]:
```

CERN IT is co-developing the integration of the SWAN/JupyterHub integration with Rucio

Bringing data sets closer to where analysis is done (in a nice WebUI)



Source: Frank, this morning  
<https://indico.cern.ch/event/1343110/contributions/6136463/attachments/2937843/5160592/Rucio%20Welcome%202024.pdf>



# Where are we?

Run and Maintain the Rucio ESCAPE cluster

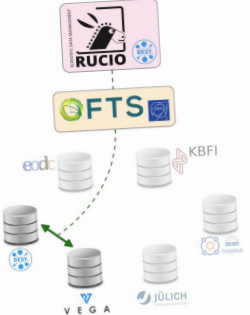
- European Science Cluster of Astronomy & Particle Physics

Contribute to InterTwin EU project

## Task 5.2: Federated data infrastructure

### Establishing the interTwin DataLake testbed:

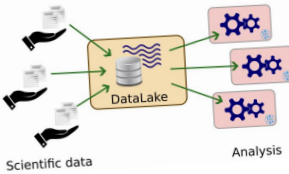
- Based on ESCAPE DataLake model
- Core service (Rucio) deployed at DESY
- Integration of dCache (DESY and VEGA)
- Work in progress @ FZ/JSC, KBFI, EODC and EGI DataHub



The diagram illustrates the data infrastructure. At the top, RUCIO and FTS are shown as core services. Below them, data is collected from various sources: eodc, VEGA, JÜLICH, and KBFI. These sources feed into a central DataLake, which is represented by a stack of cylinders. The DataLake is then used for analysis, as shown by the 'Analysis' section.

### Started integration with science use-cases


- Collected exemplar datasets for five (of seven) DT.
- The datasets are injected in the Datalake building science driven playground



The diagram shows 'Scientific data' being injected into the 'DataLake' (represented by a stack of cylinders). The DataLake then feeds into 'Analysis' (represented by three gear icons).



# Where are we?



Since March 2024, CERN IT participates in the **Rucio DevOps support for ATLAS & CMS** (around ~1 FTE)

Focus mostly on development/issues related to core storage technologies at CERN (FTS, EOS, CTA) and their integration with Rucio

- Examples:
  - Remove the most common transfer error from CMS deployment
  - Migrate some in-house ATLAS Rucio monitoring to central IT services

What do we bring?

- Expertise in core storage and less ping-pong across teams



# Where are we?

Establishing a **Rucio reference data management service at CERN** (For SMEs and other sciences)

- Started with one SME: COMPASS
- COMPASS prototype ready since Q3 2024, on hold for now





# Where are we?

Rucio challenges for the SME COMPASS prototype:

- Steep learning curve:
  - Expertise in compiled languages (Go) -> newbie in Python
  - Expertise in bare metal deployments with Puppet -> newbie K8s + Flux + Charts
  - Expertise in MySQL -> newbie PostgreSQL
  - In-place OS upgrades -> re-creating K8s clusters
  - And ... learning Rucio on the way
- Rucio with local accounts is easy, x509 is not (many configuration options and downstream dependencies (FTS delegation, RSE x509 config, ...))
- Started with MySQL but failed due to PK on SSH authentication for the bootstrapping code:

```
raise exc_class(errno, errval),\nsqlalchemy.exc.OperationalError: (pymysql.err.OperationalError) (1071, 'Specified key was too long; max key length is 3072 bytes')\n[SQL]
```

<https://github.com/rucio/rucio/issues/6755>



# Where are we?

Rucio challenges for the SME COMPASS prototype:

- Not easy to just take a production deployment (ATLAS, CMS) and just run it
  - Differences in K8s config, policies, integration with services (CRIC, Gitlab), etc ...
  - Help from ATLAS and CMS folks was invaluable
- SMEs they already have data somewhere, they are not a new experiment
  - There is no way to inject existing data with the Rucio CLI, one has to write a script (not user friendly)
  - Got help from Armin@ATLAS for their T0 data registration into Rucio
- Missing documentation, many expert knowledge not reflected in the documentation
  - Found the formula (Martin + Rucio paper + code browsing) the best source of knowledge



# Where are we heading?

Identify and engage with other SMEs

- Many SMEs develop their own data management solutions, sometime this knowledge is only in person, big bus factor
- They will benefit from a Rucio managed service, like we already do for FTS ([fts3-public.cern.ch](https://fts3-public.cern.ch)) and RSEs ([eospublic.cern.ch](https://eospublic.cern.ch), [eosctapublic.cern.ch](https://eosctapublic.cern.ch))
- We start with model of 1xSME -> 1xRucio cluster to ease the learning curve, but this will not scale with dozens of SMEs
  - Rucio multi-VO will be the only way forward in medium future
  - Testing not so thorough as single-VO instances used in production deployments



# Where are we heading?

Submitted grant to OSCARS (CERN IT/EP) for Open Data support in Rucio

- Granted, waiting selection of DevOps Engineer to start in January 2025
- More on this on Thursday talk

Continue supporting ATLAS and CMS communities with development and operational effort (> more than the ~ 1 FTE right now)



# How was so far?

Very warm welcome by the Rucio folks and the ATLAS and CMS DM teams

Learning from constructive feedback:

- PR changing 10 lines of code -> Ricardo writing 20 lines of review comments ;)

CERN IT trusted by Rucio project management to take co-responsibility of components

- Transfers, Deployment and Deletions

Precise responsibilities from experiments:

- From ATLAS: help in Rucio project development
- From CMS: help in Rucio DevOps + data management on storage technologies

Glue between the CERN IT Storage team and the ATLAS and CMS experiments



## Future

The amount of required knowledge to contribute to ATLAS and CMS Rucio DevOps and create Rucio reference management platform needs a consolidated team and long-term support