



The ESCAPE Data Lake as the bridgehead for the EOSC: the DataLake-as-a-Service

Riccardo Di Maria

CERN

October 19th, 2021 - EGI Conference 2021



Science Projects



**EUROPEAN OPEN
SCIENCE CLOUD**

Horizon2020
European Union Funding
for Research & Innovation

Partners




rijksuniversiteit
groningen



Project Goals

- Prototype an infrastructure adapted to exabyte-scale **future needs** of large science projects
- Ensure sciences drive the development of EOSC
- Address FAIR data management principles



The Data Lake

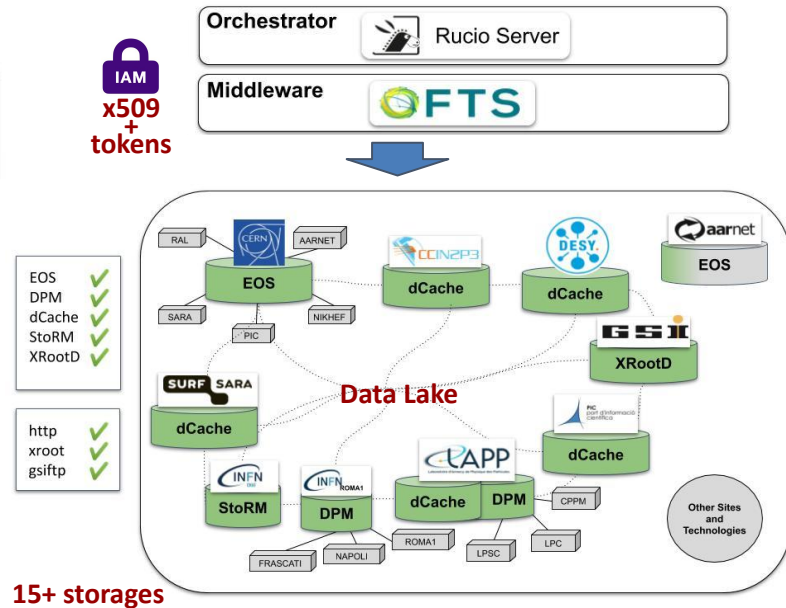
- Data Lake as modular ecosystem of services and tools shaped around the ESCAPE scientific communities
 - federated data management and access solution
 - heterogeneous resources
 - e.g. integration of HPC and **commercial Clouds**
- Hiding complexity and providing transparent access to data
 - layer for orchestration of resources as entry point for sciences
 - define data policies and rules
 - content delivery and caching layer
 - HTTP data access and Tokens awareness for future sustainability
 - latency hiding and file re-usability
 - facilitate ingress/egress with Clouds and HPC
- Storage and compute resources not necessarily colocated

9 sciences



NETWORK OPTIMIZATION

CACHING SOLUTION

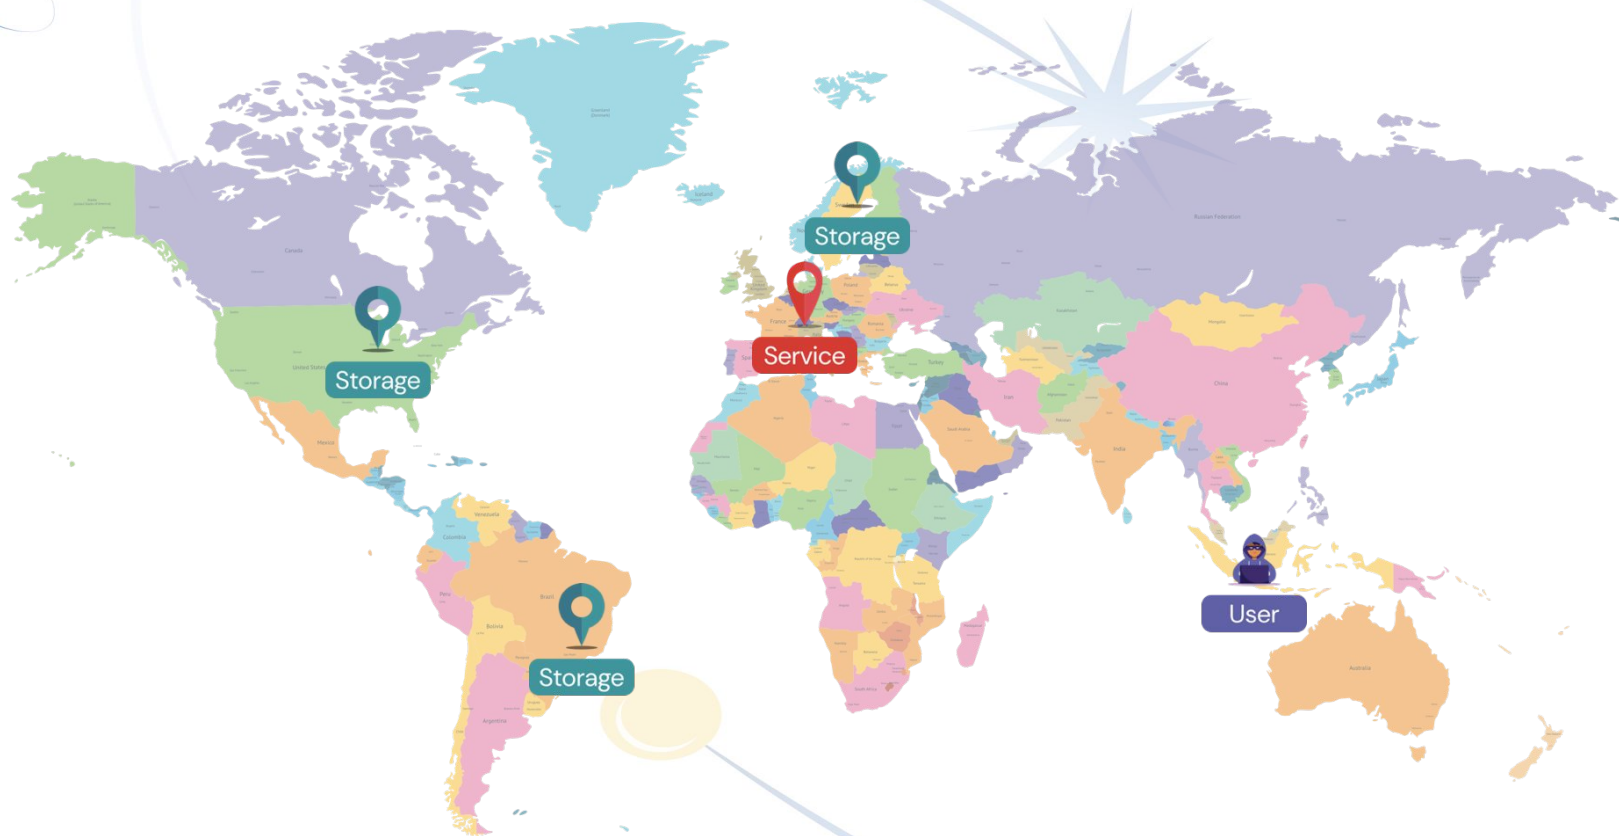


DataLake-as-a-Service for Open Science

- Challenging project as time-constrained (CERN OpenLab Summer Student programme)
- Goal: make **end-user comfortable** in embarking on a Data Lake experience
 - abstract the complexities of the Data Lake from the scientists
→ focus on doing science instead of data procurement
- An ever-increasing number of experiments are looking at Rucio Data Management system
 - **DLaaS** potentially interesting for both **aficionados** and **newcomers**
 - we are working for experiments and sciences → fresh feedback and ideas are very appreciated!
- Always looking for new collaborations and partnerships - *please step forward!*




DataLake-as-a-Service for Open Science



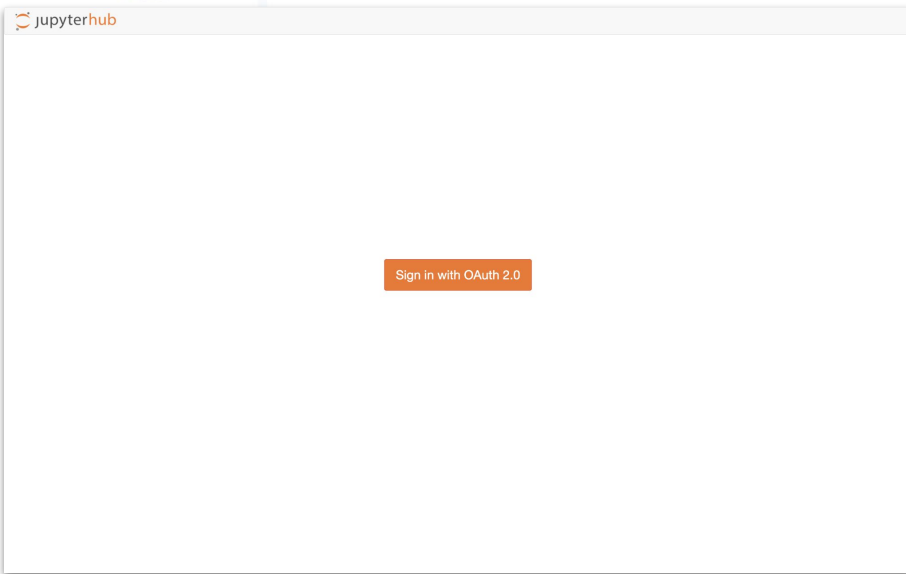
DataLake-as-a-Service for Open Science

→ Scientists Contact DataLake-as-a-Service

Requests are handled by Jupyter servers at CERN, Geneva



DataLake-as-a-Service for Open Science



<https://escape-notebook.cern.ch/>




DataLake-as-a-Service for Open Science


jupyterhub



Welcome to **escape**

Sign in with your escape credentials

 rdimaria

 *****

Sign in

[Forgot your password?](#)

Or sign in with

Your X.509 certificate

 Google

 eduGAIN

Not a member?

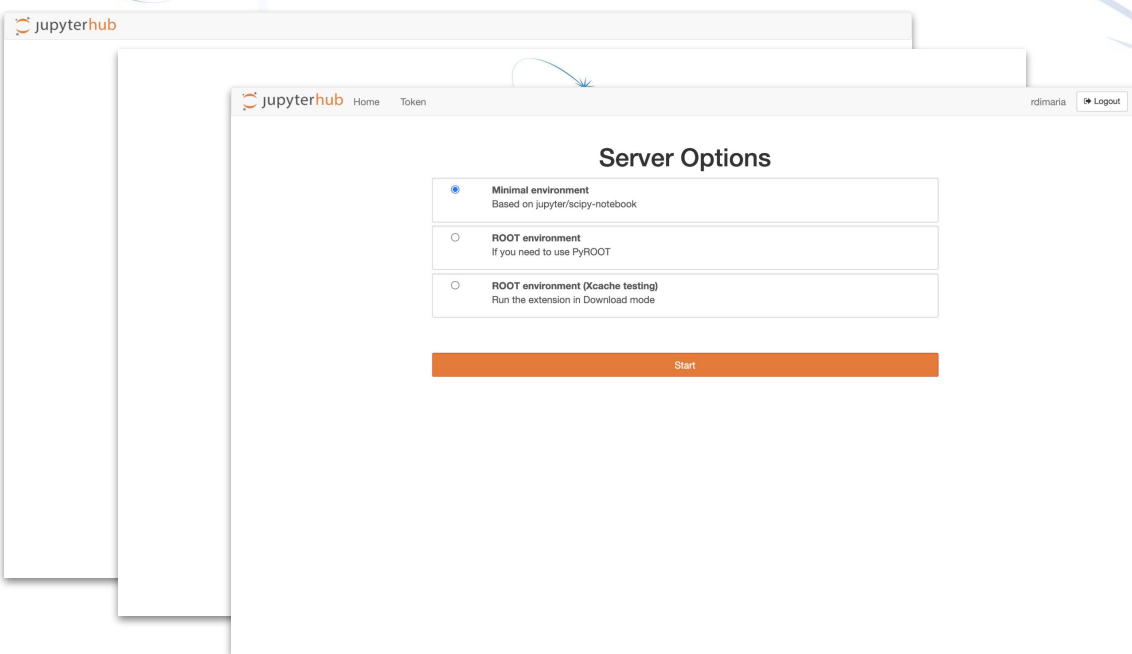
[Apply for an account](#)

[Privacy policy](#)

You have been successfully authenticated as
CN=Riccardo Di
Maria,CN=770219,CN=rdimaria,OU=Users,OU=Organic
Units,DC=cern,DC=ch



DataLake-as-a-Service for Open Science

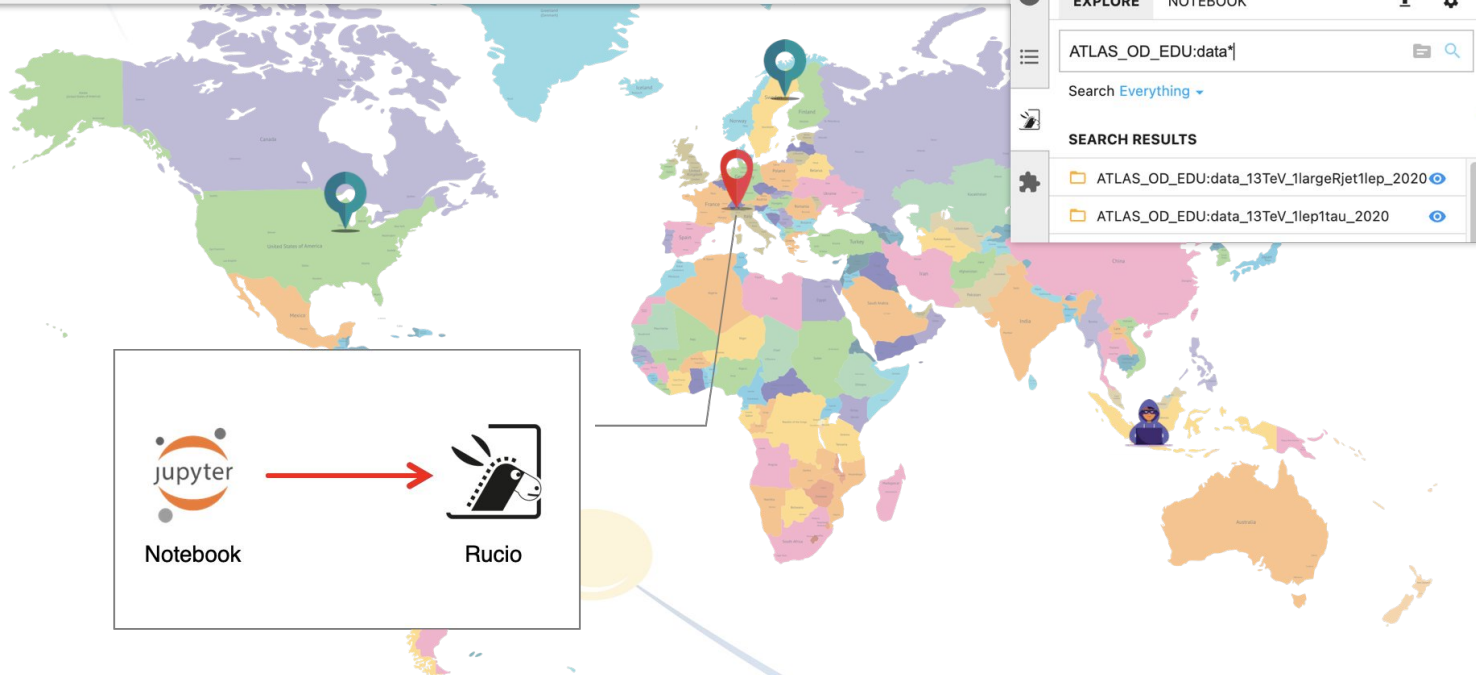




DataLake-as-a-Service for Open Science

→ Scientists Browse Data in the ESCAPE Data Lake

Requests are relayed to Rucio servers at CERN, Geneva



DataLake-as-a-Service for Open Science

→ Make Available

Rucio initiates transfers from worldwide storages to CERN RSE which is serving DLaaS



DataLake-as-a-Service for Open Science

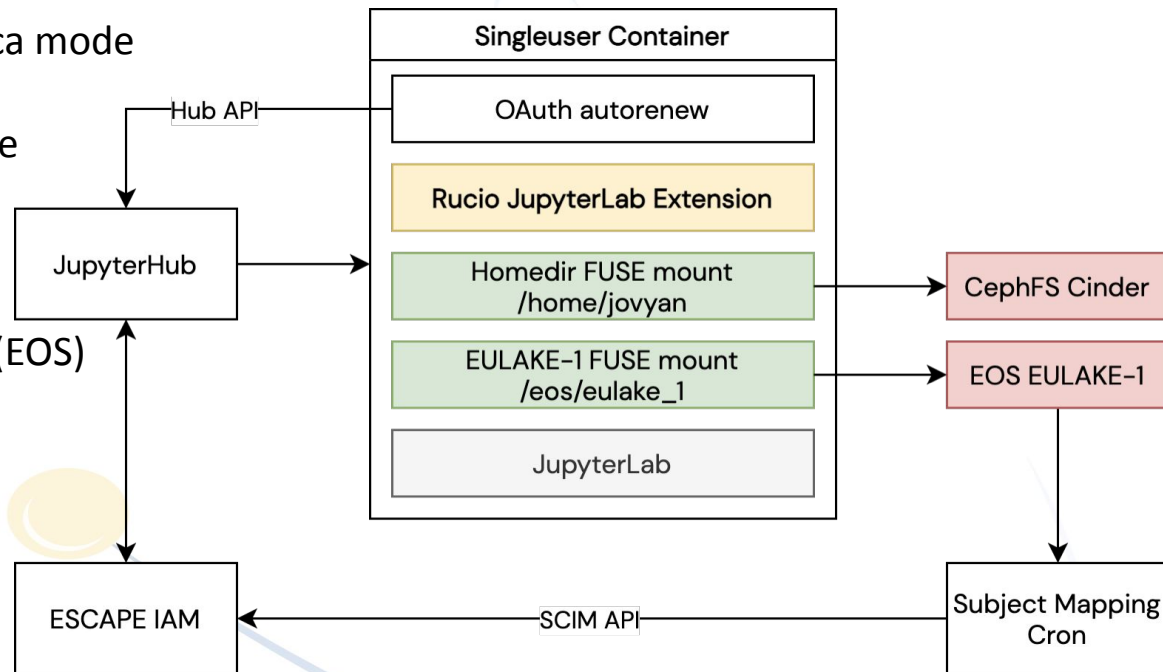
→ Scientists' Analysis Workflows Read Data

The code runs on the Notebook server at CERN, and the output is shown to the user



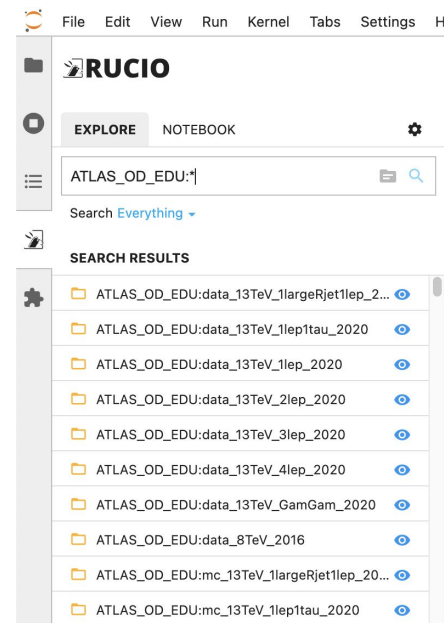
DLaaS Implementation

- Deployed in Kubernetes, using Zero-to-JupyterHub Helm chart → <https://escape-notebook.cern.ch>
- OAuth authentication using ESCAPE IAM (X509 still supported)
- [Rucio JupyterLab Extension](#) in Replica mode (i.e. TPC to local storage) used
 - connected to ESCAPE Data Lake (escape-rucio.cern.ch)
 - automatically pre-configured to use OIDC Auth
 - FUSE mount to EULAKE-1 RSE (EOS) *aka* creating a replication rule to move files to EULAKE-1
 - making files available
 - download mode still possible (if configured)



DLaaS Use Cases

- Data discovery and access
- Submitting jobs to external services (remote computing)
 - conveniently browse data in Rucio through the extension
 - access file PFN directly from the Notebook



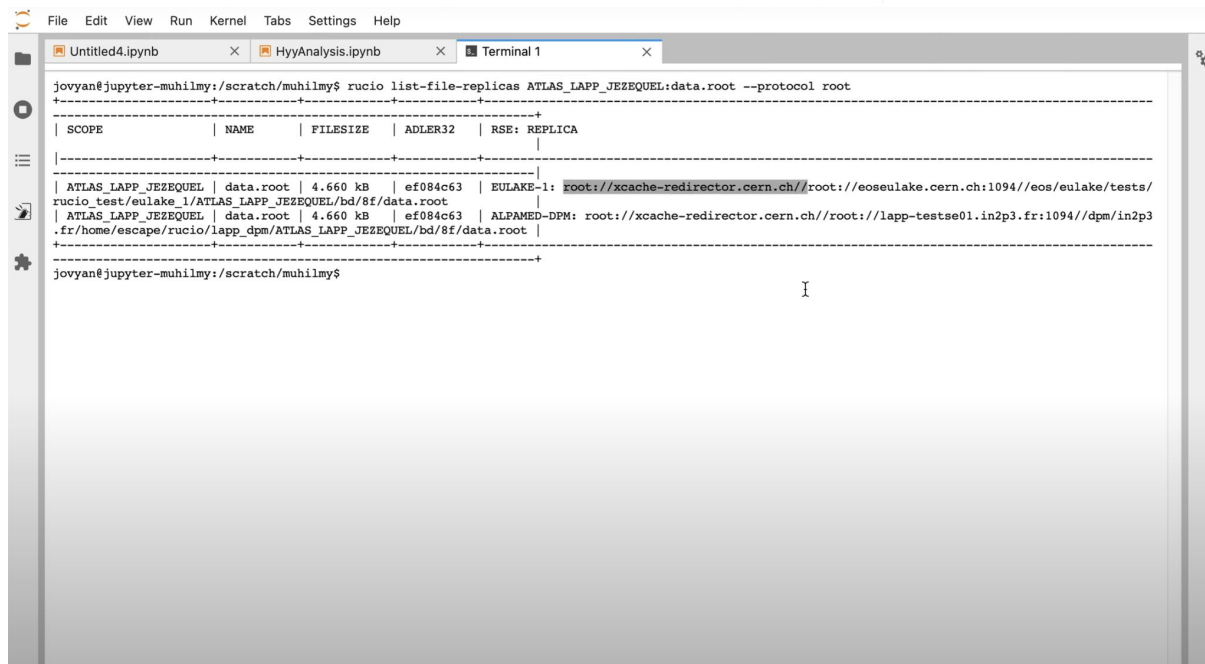
```
[3]: for item in hyy_20:
      print(item.pfn)
```

```
root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/ATLAS_LAPP_JEZEQUEL/6f/98/data_A.GamGam.root
root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/ATLAS_LAPP_JEZEQUEL/f1/3a/data_B.GamGam.root
root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/ATLAS_LAPP_JEZEQUEL/45/95/data_C.GamGam.root
root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/ATLAS_LAPP_JEZEQUEL/73/e3/data_D.GamGam.root
root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/ATLAS_LAPP_JEZEQUEL/6d/aa/mc_341081.tth125_gamgam.GamGam.root.1
root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/ATLAS_LAPP_JEZEQUEL/1b/95/mc_343981.ggH125_gamgam.GamGam.root.1
root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/ATLAS_LAPP_JEZEQUEL/ff/c7/mc_345041.VBFH125_gamgam.GamGam.root.1
root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/ATLAS_LAPP_JEZEQUEL/13/b8/mc_345318.WpH125J_Winc1_gamgam.GamGam.root.1
root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/ATLAS_LAPP_JEZEQUEL/76/fd/mc_345319.ZH125J_Zinc1_gamgam.GamGam.root
```



DLaaS Use Cases

- Data access
 - bonus track: content delivery and caching layer (*aka* XrootD-XCache)



```

jovyan@jupyter-muhilmy:/scratch/muhilmy$ rucio list-file-replicas ATLAS_LAPP_JEZEQUEL:data.root --protocol root
+-----+-----+-----+-----+-----+
| SCOPE          | NAME          | FILESIZE | ADLER32 | RSE: REPLICA |
+-----+-----+-----+-----+-----+
| ATLAS_LAPP_JEZEQUEL | data.root | 4.660 kB | ef084c63 | EULAKE-1: root://xcache-redirector.cern.ch//root://eosulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/ATLAS_LAPP_JEZEQUEL/bd/8f/data.root |
| ATLAS_LAPP_JEZEQUEL | data.root | 4.660 kB | ef084c63 | ALPAMED-DPM: root://xcache-redirector.cern.ch//root://lapp-testse01.in2p3.fr:1094//dpm/in2p3.fr/home/escape/rucio/lapp_dpm/ATLAS_LAPP_JEZEQUEL/bd/8f/data.root |
+-----+-----+-----+-----+-----+
jovyan@jupyter-muhilmy:/scratch/muhilmy$
  
```



DLaas Use Cases

- Data preparation and processing
 - prepare/process data and upload back to the Data Lake

EULAKE-1 storage FUSE-mounted

SCRATCH treated as Rucio storage element and FUSE-mounted

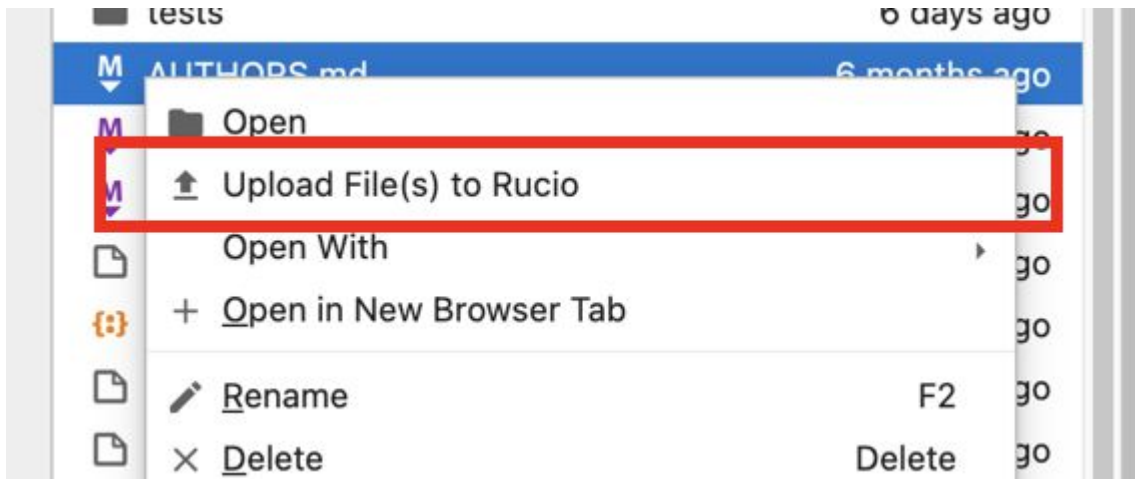
Rucio CLI embedded


using Rucio/FTS, hence TPC



DLaaS Use Cases

- Data preservation
 - produce data and upload to the Data Lake




RUCIO

Upload AUTHORS.md to Rucio

Please make sure that the necessary credentials are configured. You can see the upload status on the Rucio sidebar.

Destination RSE Expression:

Lifetime (in seconds):

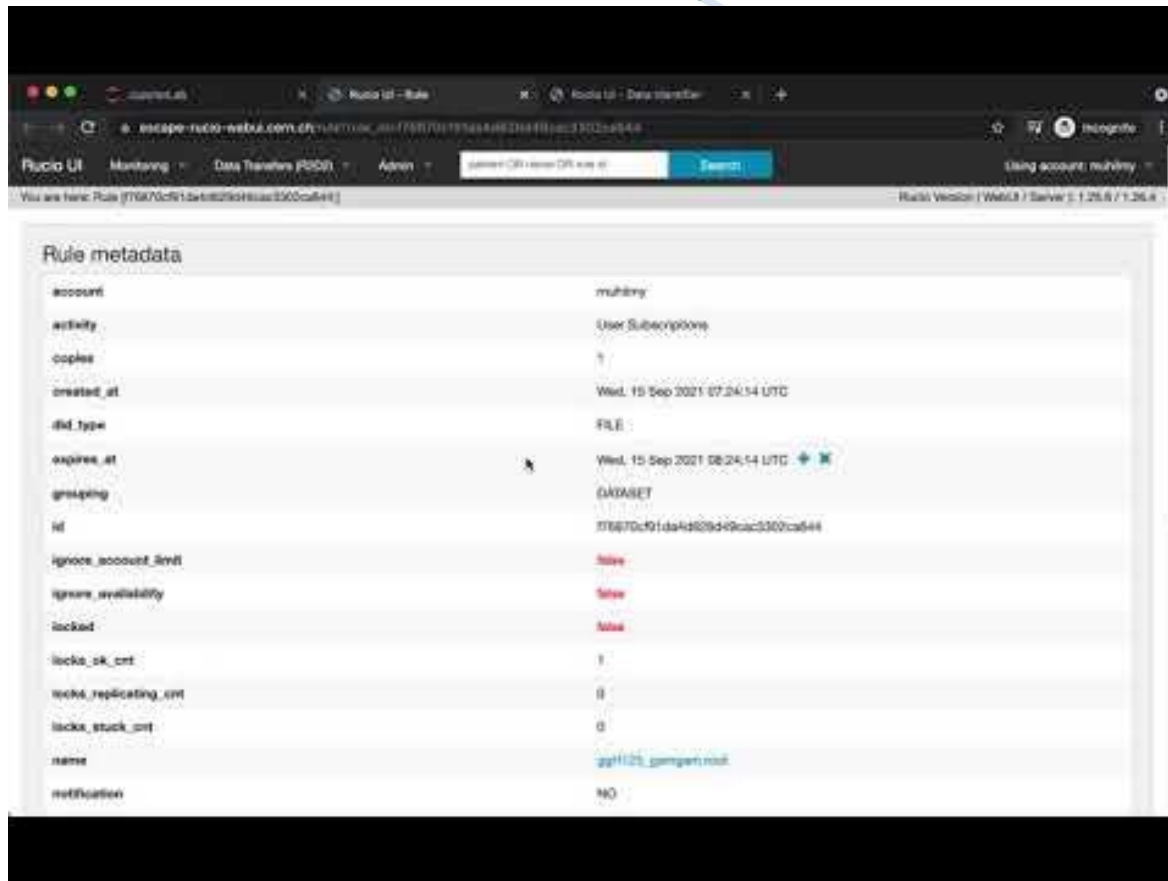
Scope:

ATLAS_OD_EDU

Cancel
 Upload



DataLake-as-a-Service for Open Science



The screenshot shows the Rucio UI web interface. The browser address bar displays the URL: `escape-rucio-webui.com.cn/rucio/ui/Rule/176670c9f1da4d829d49cac3301ca644`. The page title is 'Rule metadata'. The interface shows a table of rule metadata with the following fields and values:

Field	Value
account	ruhling
activity	User Subscriptions
copies	1
created_at	Wed, 15 Sep 2021 07:24:54 UTC
id	176670c9f1da4d829d49cac3301ca644
id_type	FILE
expires_at	Wed, 15 Sep 2021 08:24:54 UTC
grouping	DATASET
ignore_account_limit	False
ignore_availability	False
locked	False
locks_ok_cnt	1
locks_replicating_cnt	0
locks_stuck_cnt	0
name	ggfH125_gamgert.mad
notification	NO



Conclusion and Next Steps

- **DLaaS** hides the complexities of the Data Lake from the end users (so that they are happy and productive!)
 - interesting for both **aficionados** and **newcomers** of **Rucio**
 - community-driven development, hence driven by the needs of different experiments and sciences
- Successful assessment of the pilot Data Lake, pivotal to test model and concepts for several communities
 - Astro-particle Physics, Electromagnetic and Gravitational-Wave Astronomy, Particle Physics, and Nuclear Physics **pursuing together** FAIR and open-access data principles
- ESCAPE at prototype phase → full scale exercise for November 2021
 - close-to-production test on data management and data processing
 - complementing existing efforts in WLCG → ESCAPE as perfect environment to test new models/concepts
 - exploring non-HEP-standard scenarios
 - collaboration with other communities, e.g. PaNOSC, ExPaNDS, CS3MESH4EOSC
 - sharing common challenges, potentially address-able with common solutions



Backup Slides

- *25th International Conference on Computing in High-Energy and Nuclear Physics (vCHEP2021) talks and [proceedings](#)*
 - [ESCAPE Data Lake: Next-generation management of cross-discipline Exabyte-scale scientific data](#)
 - [The ESCAPE Data Lake: The machinery behind testing, monitoring and supporting a unified federated storage infrastructure of the exabyte-scale](#)
 - [Experience with Rucio in the wider HEP community](#)



Architecture

- Data Storage Services and Usage

- proven technologies from WLCG
- perfect playground for the implementation of storage Quality of Service intelligence



CERN
Tape Archive

- Data Orchestration

- Rucio customised to the needs of ESCAPE community



- File Transfer Service → WLCG FTS

- continuous testing for both FTS and file access libraries



- Networking → exploiting [perfSONAR infrastructure](#)



- Information System → WLCG CRIC

- ESCAPE instance containing services information and configuration for Rucio
- synchronisation services designed to pull necessary information



- AuthN/Z, and Identity layer → INDIGO Identity and Access Management (IAM) service

- Probes and monitoring exploiting open source solutions

- access availability and reliability of the services
- ESCAPE dashboards as multi-VO



Deployment Model and Techniques

- **Goal:** common infrastructure and core services for ESCAPE sciences that operate services and tools according to their scientific needs
 - improved project sustainability and successful exportability adoption by sciences
- Beyond ESCAPE term: ability to deploy, customise, and manage services at convenience
 - infrastructure should be resource-aware - [Kubernetes & GitOps-Flux](#)
 - properly-scaled, according to needs/necessities of partners
 - flexible, portable, plug-able, and easy-sizable
 - deploy a manifold system utilising a lightweight but complete implementation
 - sciences at different scale and trying to address multiple future use case
→ ESCAPE includes experiments with different data management requirements than LHC ones
- Contribution to Rucio on functionalities relevant for ESCAPE partners
 - expand Rucio scope and capabilities e.g. to address million-file rule injection use case
- Synergies and collaborations within and beyond ESCAPE
 - [Rucio/JupyterLab Integration Project](#) within CERN-HSF GSoC, exploited by LOFAR to analyse data during the FDR

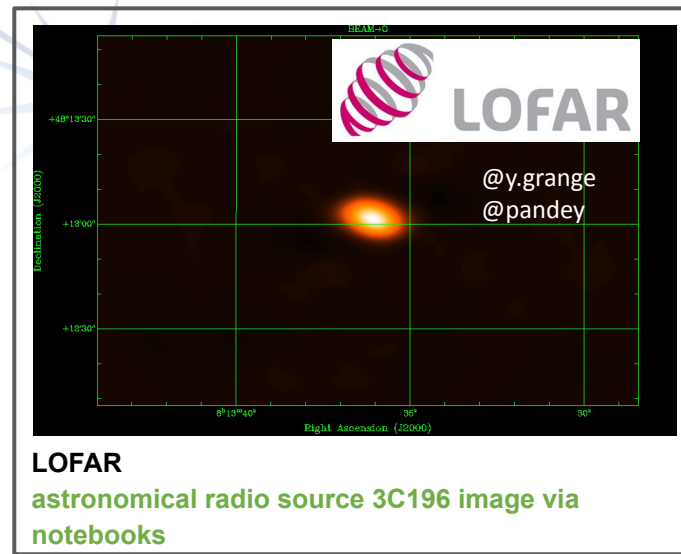






Data Lake 24-hour Full Dress Rehearsal

- Aim: demonstrate fulfilment of communities needs
 - functionalities and use cases specific per experiment
- ESCAPE Data Lake not a production service → specific testing-focused time-window
 - even out differences in knowledge and acquiring know-how on management and utilisation
 - workflow and lifecycle of experiment-specific data
 - assessment of robustness of the various Data Lake components, tools, and services
- Experiments injected data as during runs
 - effort to mimic realistic conditions in terms of object sizes and data organisation
 - demonstrator categories
 - **injection** - **rules and policies** - **data lifecycle** - **access** → *Backup slides*




DL 24-hour Full Dress Rehearsal Takeaway → [Workshop](#)



LSST
realistic LSST data rate (20TB/24h)
confirm integrity and accessibility of the data via
notebooks




DL 24-hour Full Dress Rehearsal Takeaway → [Workshop](#)




MAGIC

mimics a real MAGIC observation use case
remote storage next to the telescope as a buffer for injection
local deletion




FAIR

mock data ingestion from CBM (FAIR experiment)
request 2 replicas in different QoS



EGO/VIRGO

4h VIRGO public data sampled at 4 kHz
download to CNAF-STORM and assess data quality



SKA

Pulsar Observations injection
files fall into two containers, representing different SKA Projects
24h test moving data on basis of QoS class



Future Developments

- Additional kernel compatibility
 - currently, only Python supported
- Token-support for direct download and upload
 - OIDC integration ongoing to all remaining ESCAPE RSEs
- Integration with content delivery and caching layer
 - XCache can be integrated to allow faster file download
 - completely transparent from the user PoV
 - successfully tested at small scale



References

- ESCAPE, <https://projectescape.eu/> - <https://doi.org/10.1051/epjconf/202024504019>
- ESCAPE Data Lake, https://wiki.escape2020.de/index.php/WP2_-_DIOS#Datalake_Status
- ESCAPE Data Lake deployment, <https://gitlab.cern.ch/escape-wp2/flux-rucio>
- ESCAPE Data Lake tools, <https://github.com/ESCAPE-WP2>
- CERN-HSF GSoC project: [Rucio-SWAN Integration](#)
- Data Lake 24-hour Full Dress Rehearsal (FDR) results, <https://indico.in2p3.fr/event/22693/>
- Rucio, <https://rucio.cern.ch/>
- FTS, <https://fts.web.cern.ch/fts/>
- GFAL, <https://dmc-docs.web.cern.ch/dmc-docs/gfal2/gfal2.html>
- PerfSONAR, <https://www.perfsonar.net/>
- PaNOSC, <https://www.panosc.eu/>
- ExPaNDS, <https://www.panosc.eu/>
- CS3MESH4EOSC, <https://cs3mesh4eosc.eu/>

