

# Data Lake as a Service

It's a lake.. but for data.. presented as a service

Muhammad Aditya Hilmy  
CERN Openlab Summer Student 2021

Rucio Development Meeting, 26 Aug 2021


# A bit of context

- We will have HL-LHC and other experiments coming online.
- Data volume expected to increase to exabyte scale.
- We need to think about how to store and manage the data.
- The Data Lake is a place where experiments can 'dump' their data.
- ...and scientists can 'fish' data from.
- The challenge: making sure the scientists can 'fish' easily.



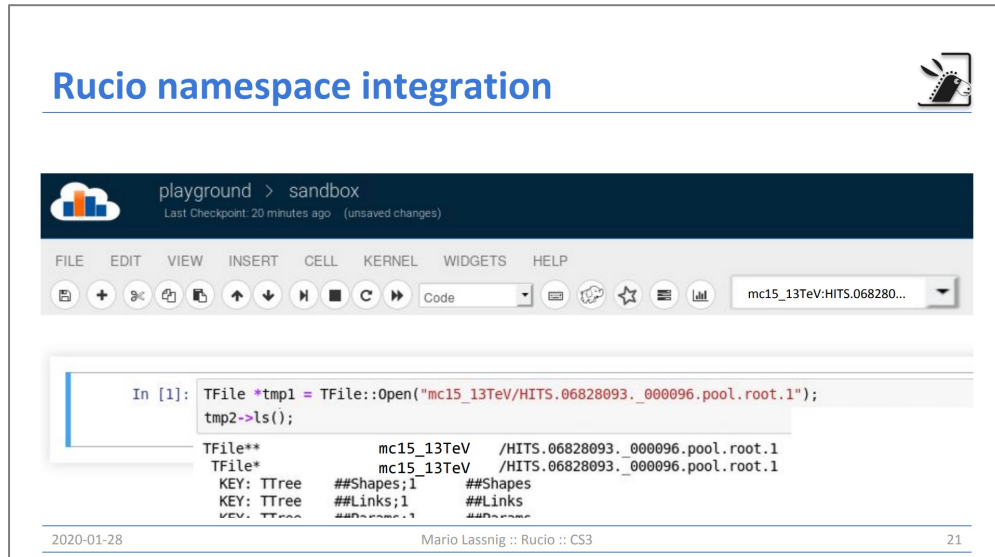
# Making 'data fishing' easier

- The Data Lake has a lot of moving parts.
- The goal of the service is to abstract the complexities of the Data Lake from the scientists.
- This way, scientists can focus their time on data analysis instead of procurement.



# A humble beginning

- Started as an idea presented on CS3 2020 by the Rucio team [1]
- Developed “Rucio JupyterLab Extension” as a part of Google Summer of Code 2020 [2]
- Deployed the extension as “Data Lake as a Service” as a part of Openlab Online (🙄) Summer Students Programme 2021



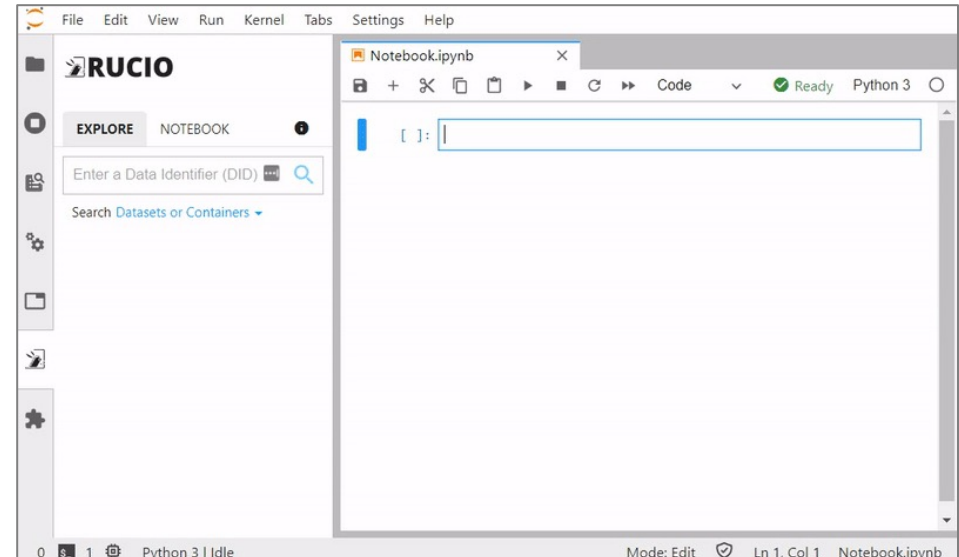
The screenshot shows the Rucio namespace integration interface in JupyterLab. The title bar reads "Rucio namespace integration". Below it, the breadcrumb "playground > sandbox" is visible. The interface includes a menu bar (FILE, EDIT, VIEW, INSERT, CELL, KERNEL, WIDGETS, HELP) and a toolbar with various icons. A dropdown menu shows "mc15\_13TeV:HITS.068280...". The main area contains a code cell with the following code:

```
In [1]: TFile *tmp1 = TFile::Open("mc15_13TeV/HITS.06828093._000096.pool.root.1");
tmp2->ls();
```

The output of the code cell is:

```
TFile**      mc15_13TeV /HITS.06828093._000096.pool.root.1
TFile*      mc15_13TeV /HITS.06828093._000096.pool.root.1
KEY: TTree  ##Shapes;1    ##Shapes
KEY: TTree  ##Links;1    ##Links
KEY: TTree  ##Links;1    ##Links
```

At the bottom, the footer indicates "2020-01-28" and "Mario Lassnig :: Rucio :: CS3" with a page number "21".



The screenshot shows the Rucio JupyterLab extension interface. The title bar reads "RUCIO". The interface includes a menu bar (File, Edit, View, Run, Kernel, Tabs, Settings, Help) and a toolbar with various icons. The main area is divided into two sections: "EXPLORE" and "NOTEBOOK". The "EXPLORE" section has a search bar labeled "Enter a Data Identifier (DID)" and a "Search Datasets or Containers" dropdown. The "NOTEBOOK" section shows a code cell with the following code:

```
[ ]:
```

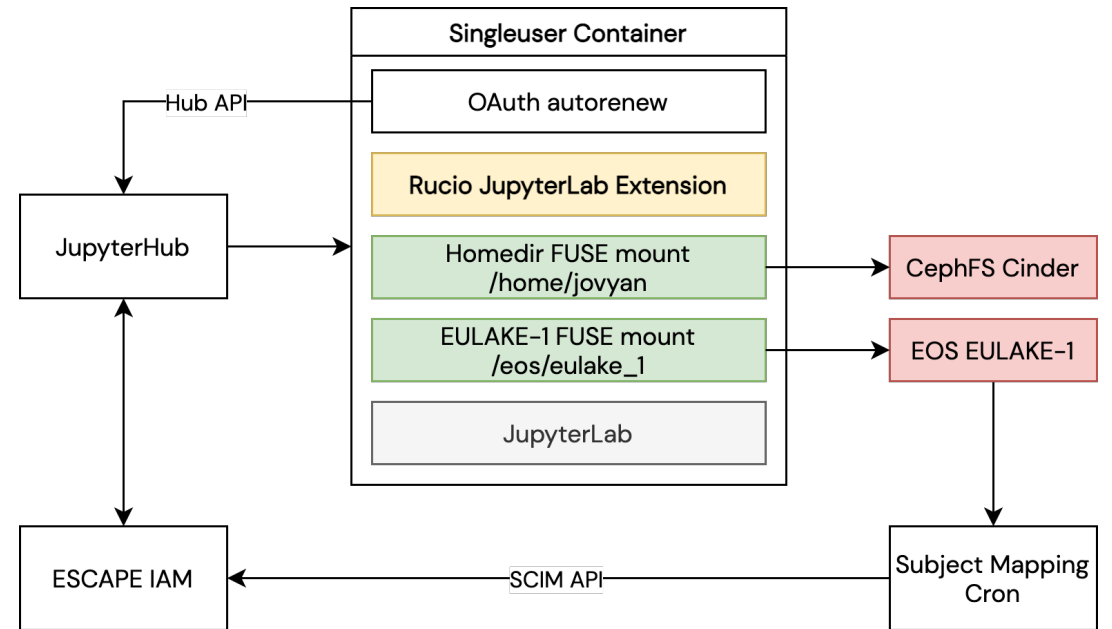
At the bottom, the footer indicates "Python 3 | Idle" and "Mode: Edit Ln 1, Col 1 Notebook.ipynb".

[1] <https://indico.cern.ch/event/854707/contributions/3680520/>

[2] [https://hepsoftwarefoundation.org/gsoc/2020/proposal\\_SWAN\\_RUCIO\\_integration.html](https://hepsoftwarefoundation.org/gsoc/2020/proposal_SWAN_RUCIO_integration.html)

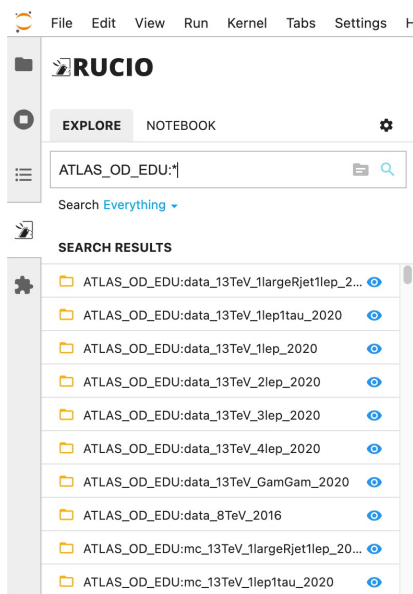
# Implementation

- Deployed in Kubernetes, using Zero-to-JupyterHub Helm chart.
  - <https://escape-notebook.cern.ch>
- OAuth authentication using ESCAPE IAM.
  - X509 and Userpass are still supported
- Uses [Rucio JupyterLab Extension](#) in Replica mode (i.e. TPC to local storage)
  - Connected to ESCAPE Data Lake (escape-rucio.cern.ch)
  - Automatically preconfigured to use OIDC authentication
  - Has a FUSE mount to EULAKE-1 RSE (EOS)
  - Making files available means creating a replication rule to move files to EULAKE-1
  - Download mode is still possible, if configured



# Use Cases

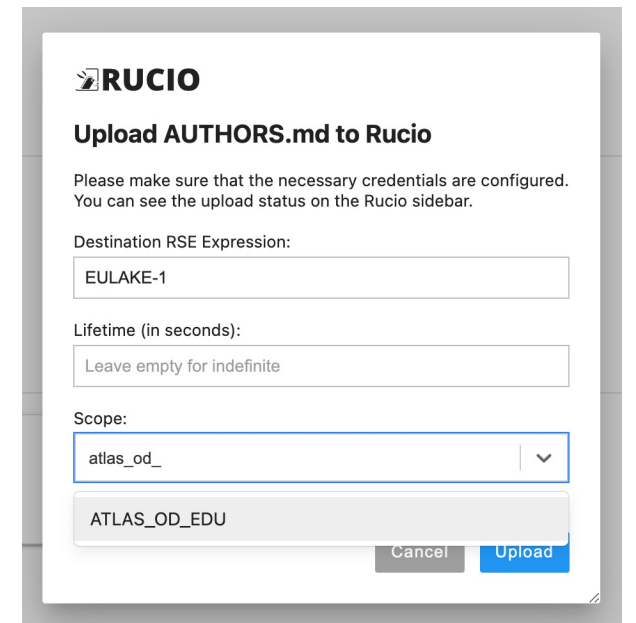
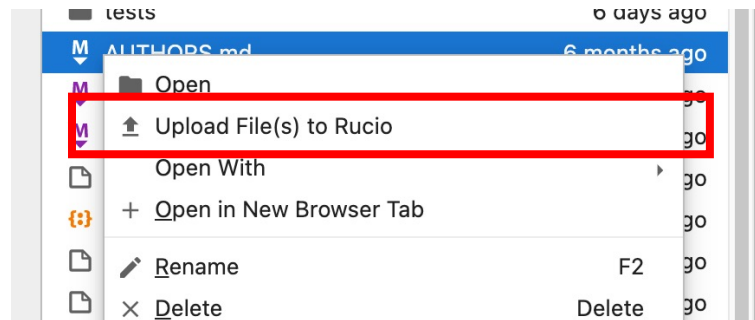
- Data discovery and access
- Submitting jobs to external service (remote computing)
  - Users can use the convenience of the extension to browse data in Rucio and access the file PFN directly from the notebook code.



```
[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.tH125_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_Wincl_gamgam.GamGam.root.1  
root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/ATLAS_LAPP_JEZEQUEL/76/fd/mc_345319.ZH125J_Zincl_gamgam.GamGam.root
```

# Use Cases (2)

- Data preparation and processing
  - Use the service to preprocess data, and once done, upload it back to the Data Lake.
- Data preservation
  - Use the service to produce data and reupload them 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

# Future Developments


- More kernel compatibility
  - Currently, only Python is supported
- Token-support for direct download and upload
  - For now, the only universally acceptable authentication for direct download/upload is x509
  - Token support for download/upload is limited to certain protocols and storages.
- Integration with content delivery and caching layer
  - XCache can be integrated to allow faster file download
  - Will be completely transparent from the PoV of the users
  - Successfully tested at small scale





# Thank you.

...and now, a short demo.

 Muhammad Aditya Hilmy

 mhilmy@hey.com

 didithilmy