CERN openlab

ESCAPE
European Science Cluster of Astronomy & Particle physics ESFRI research Infrastructures

Data Lake as a Service

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

Muhammad Aditya Hilmy <mhilmy@hey.com>
Riccardo Di Maria <riccardo.di.maria@cern.ch>

ESCAPE EEB Extended Discussions Day – 28 Sep 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 doing science instead of data procurement.

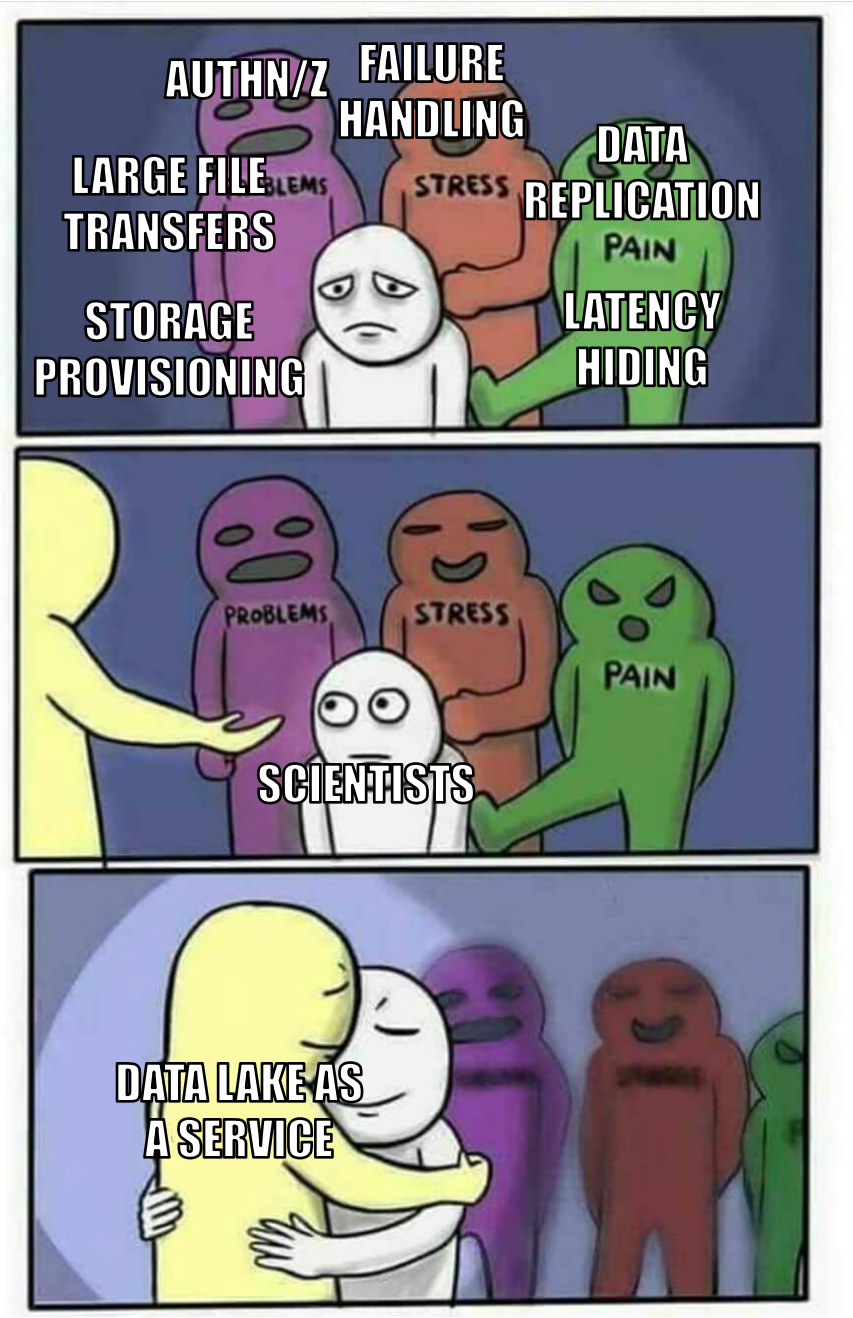

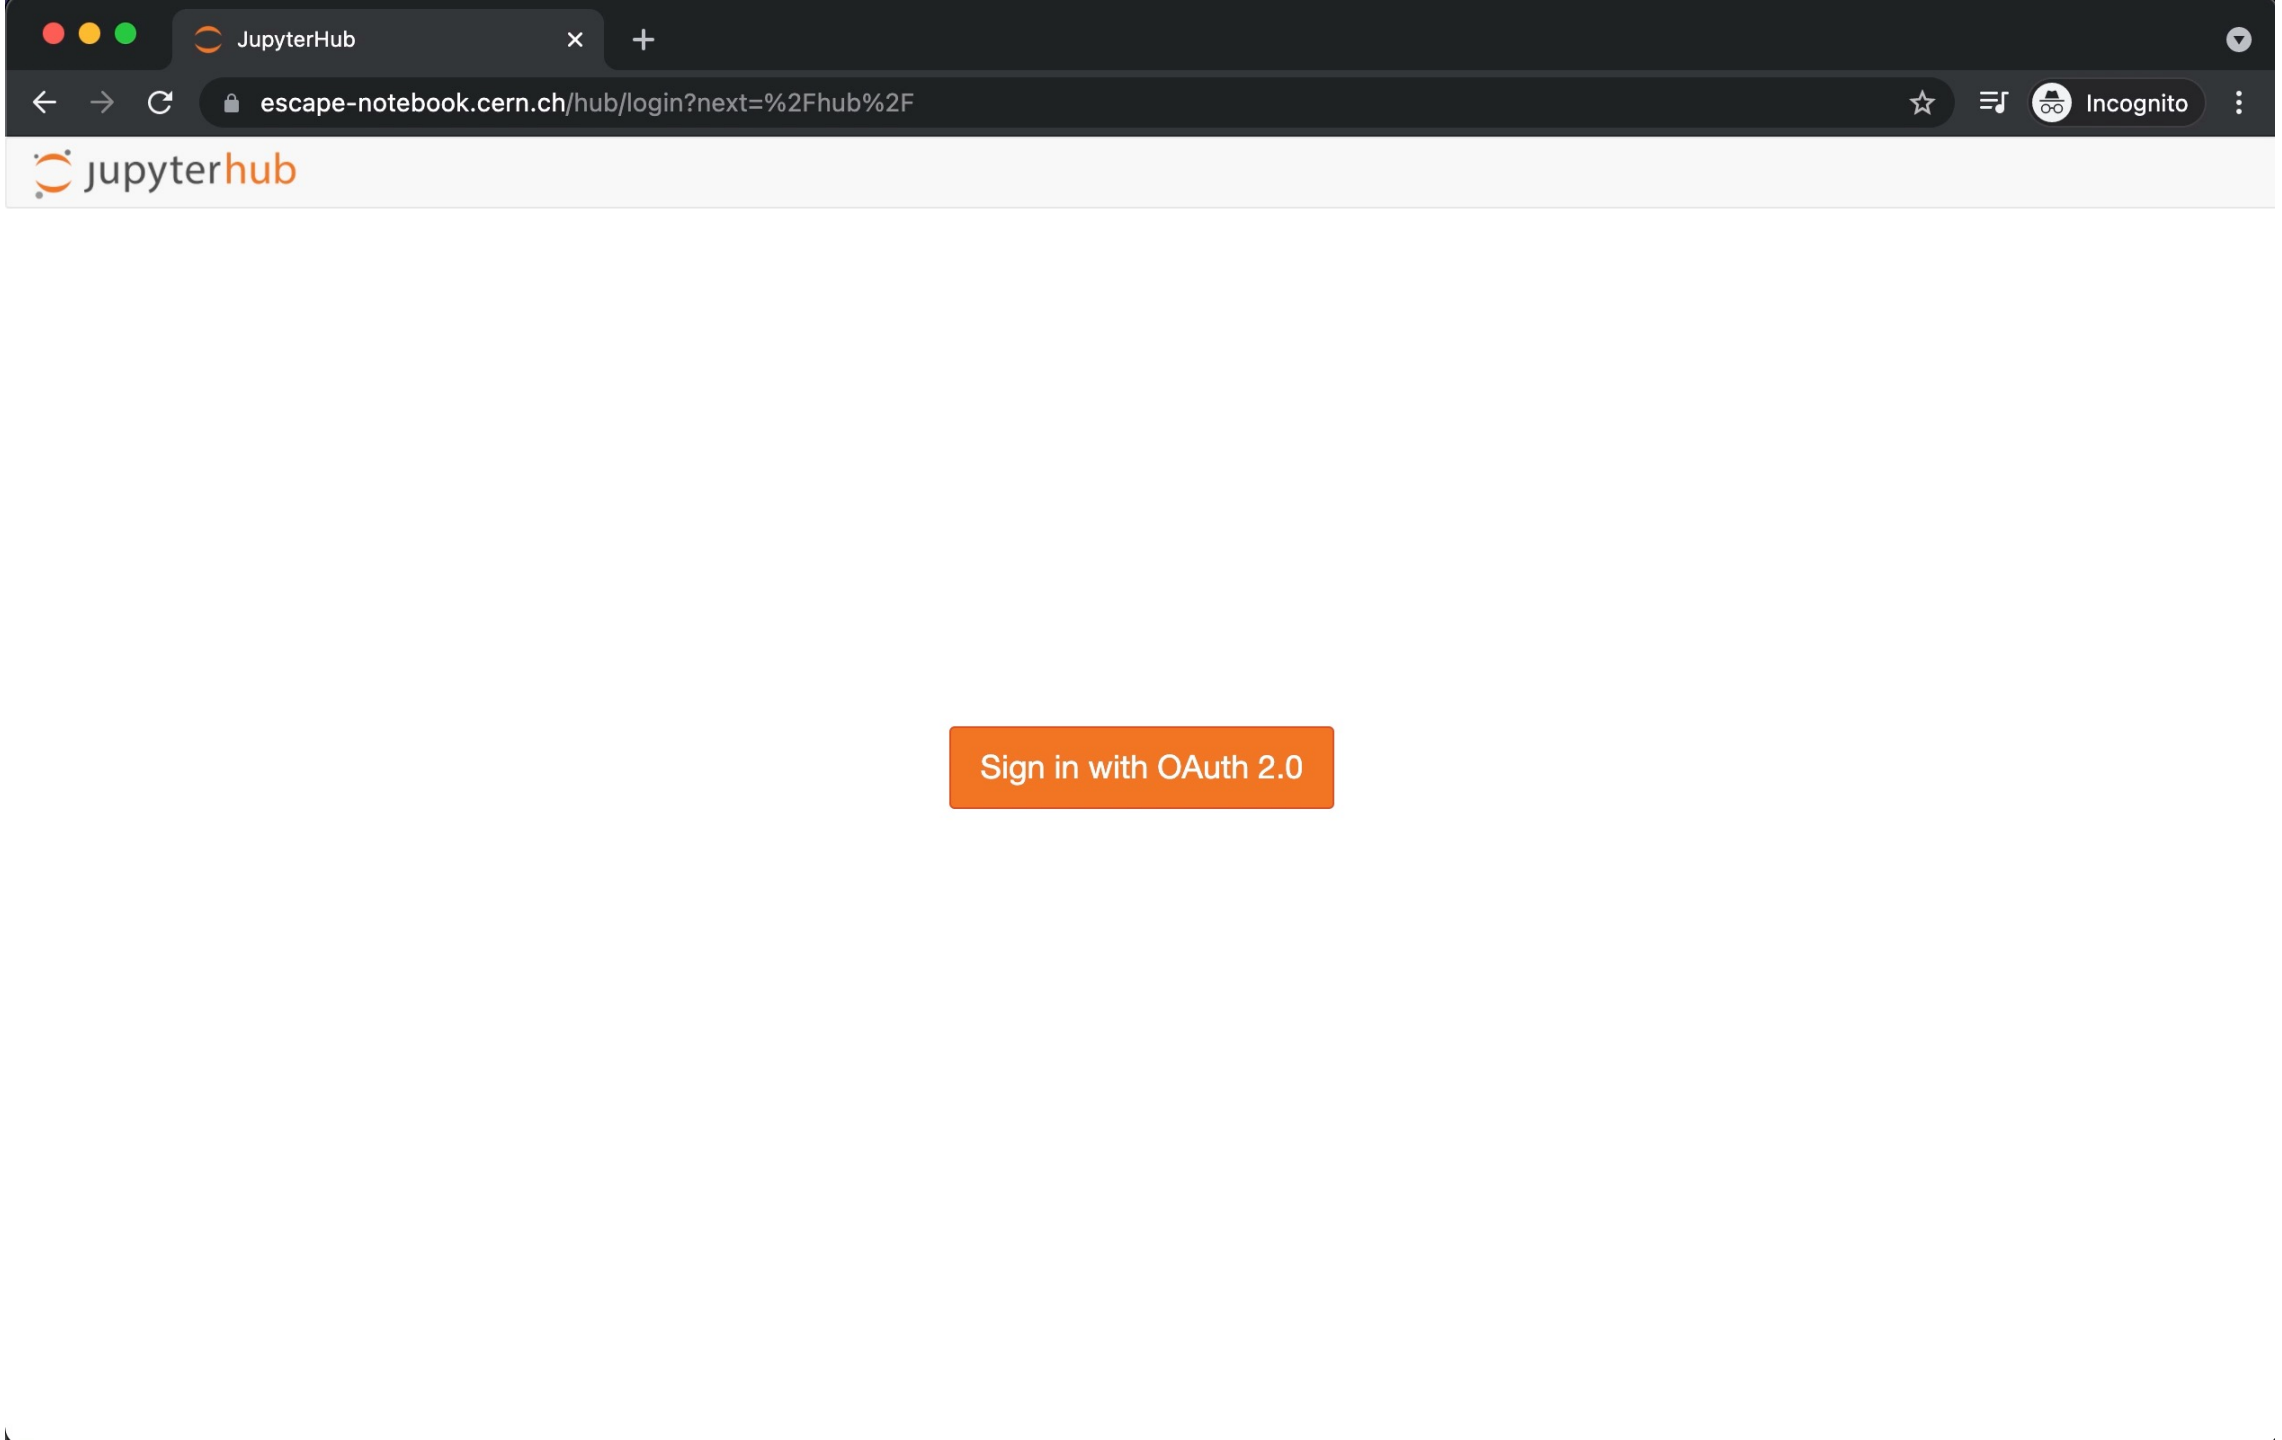

youtu.be/pRLII3fhuWc

So, what happens behind the curtains?

Don’t worry, this kind of curiosity doesn’t kill the cat.

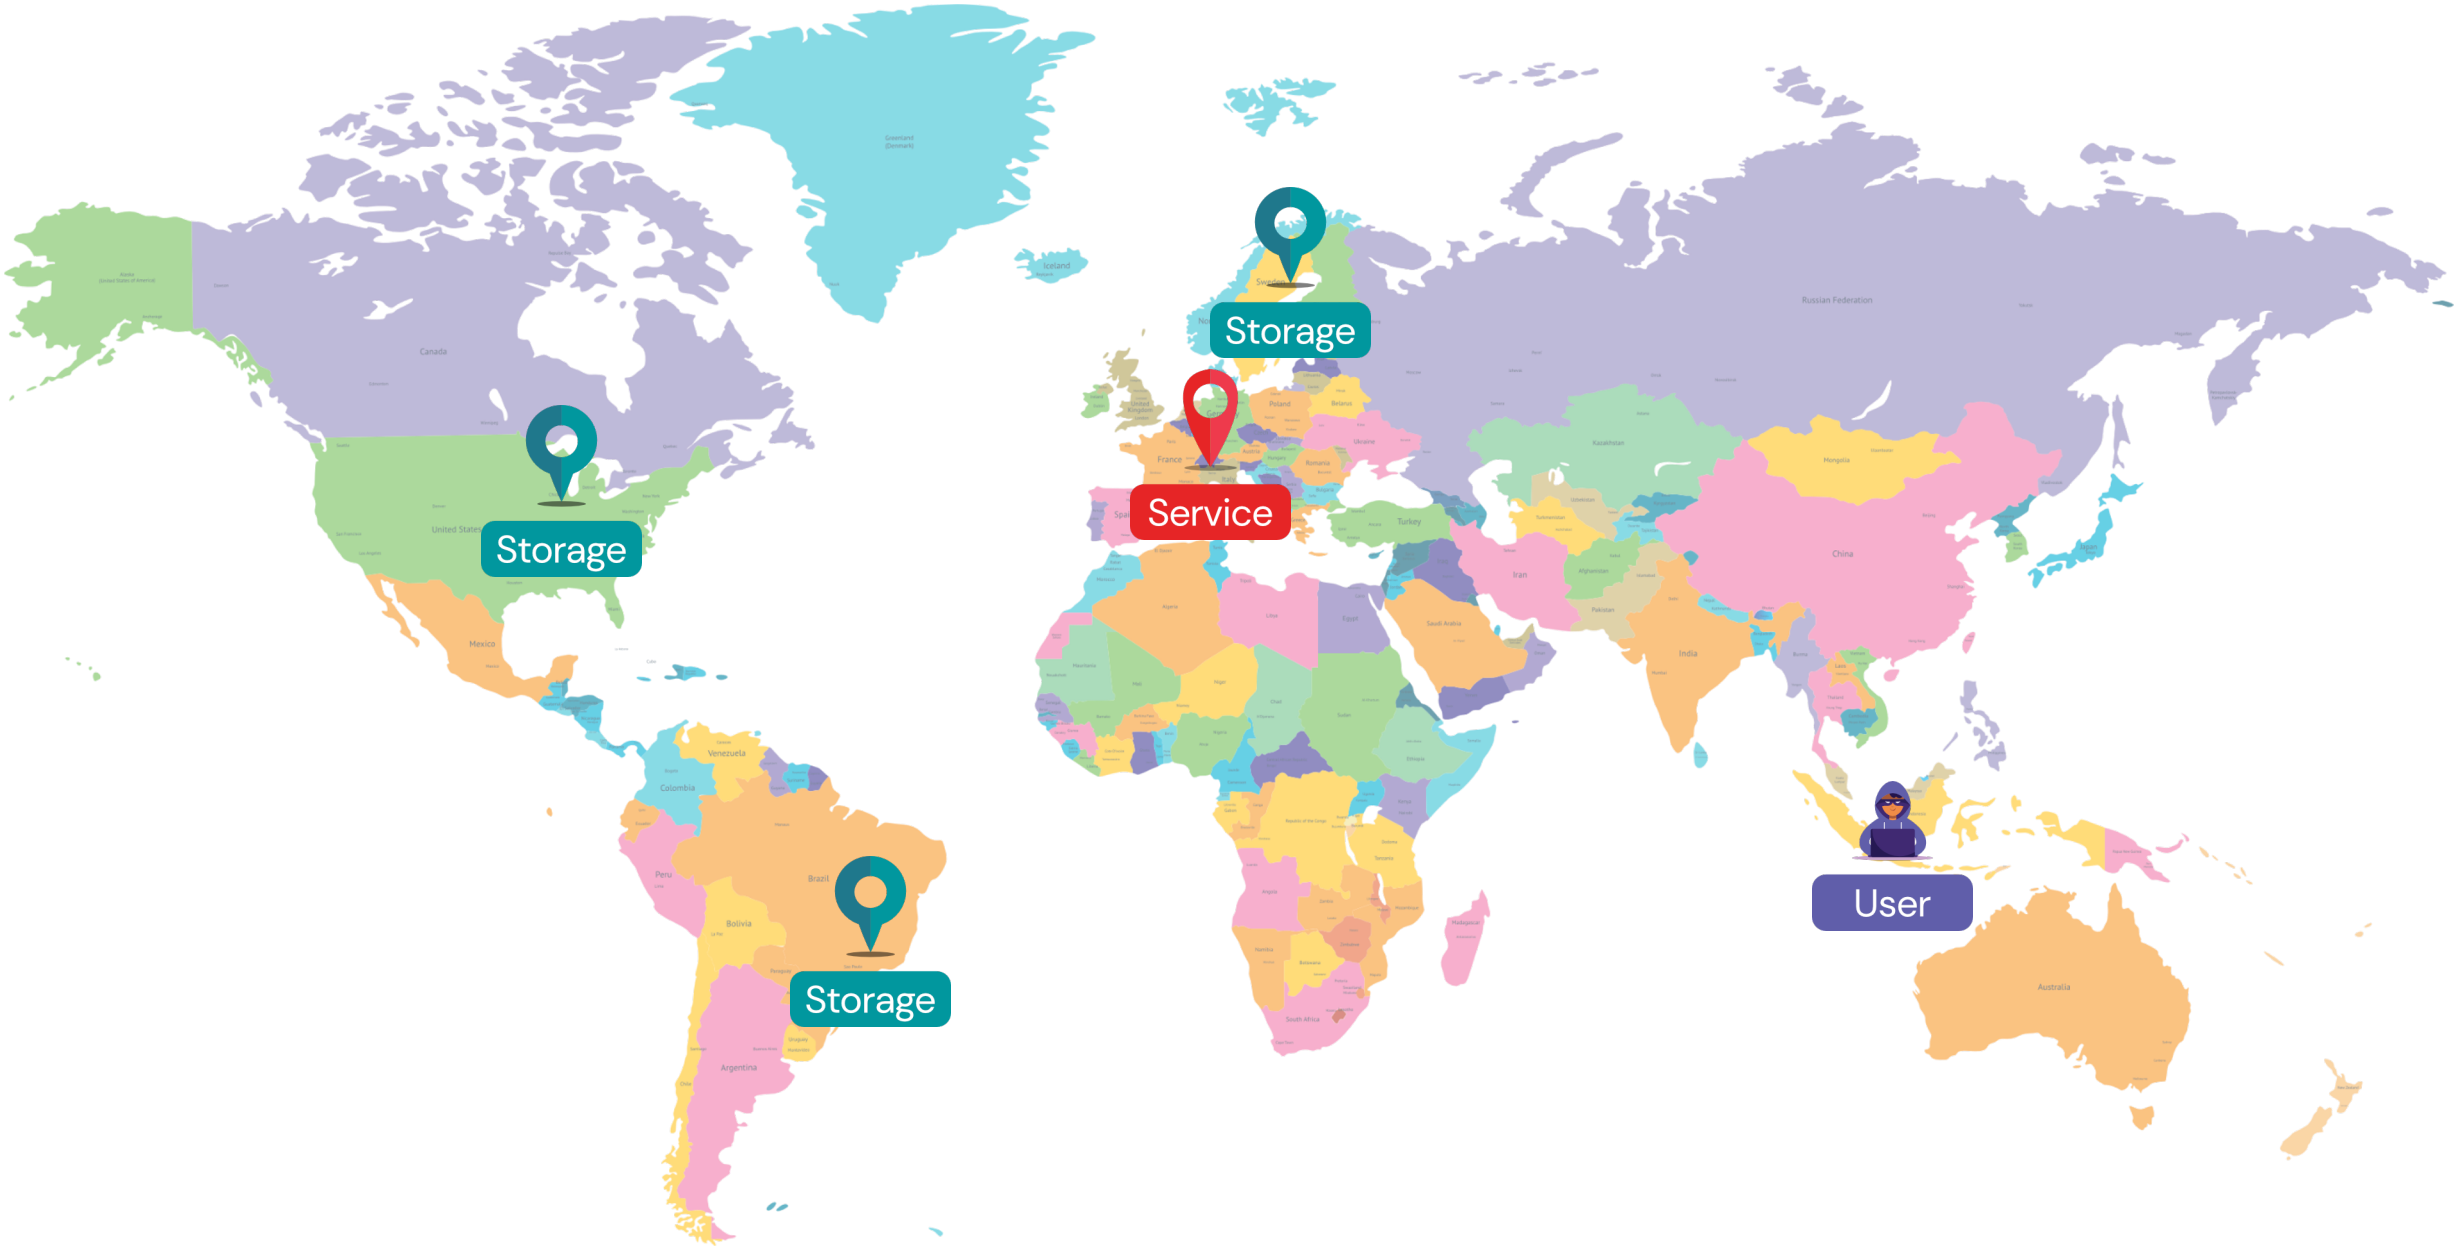
Storage
Storage
Service
Storage
User

→ User opens the Data Lake as a Service

Request is handled by Jupyter servers at CERN in Geneva

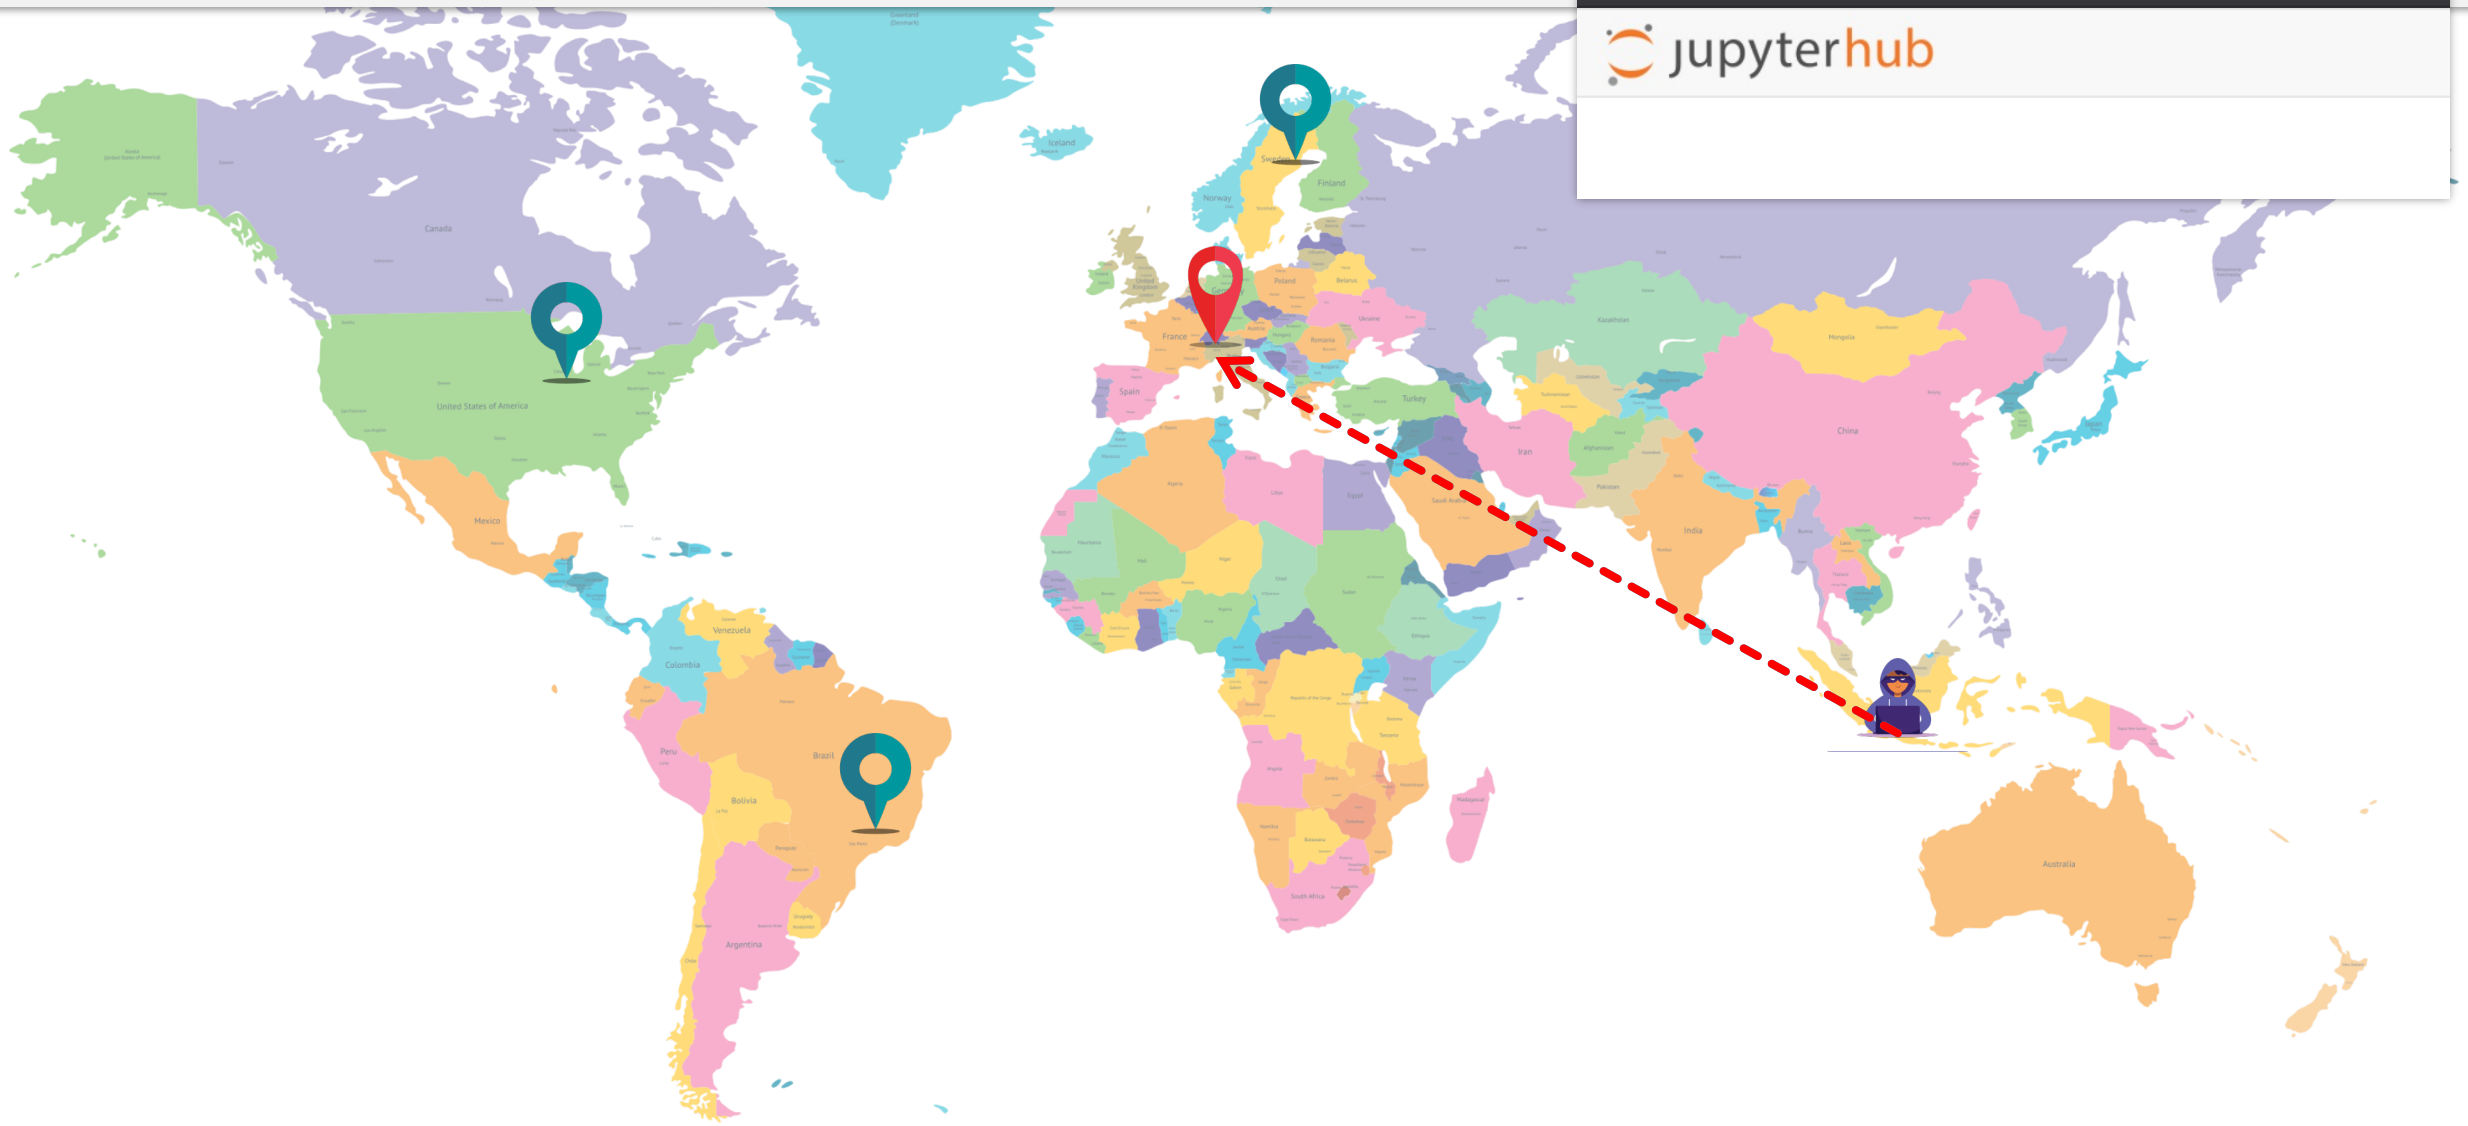

→ User browses data in the Data Lake

Request is relayed to Rucio servers at CERN in Geneva

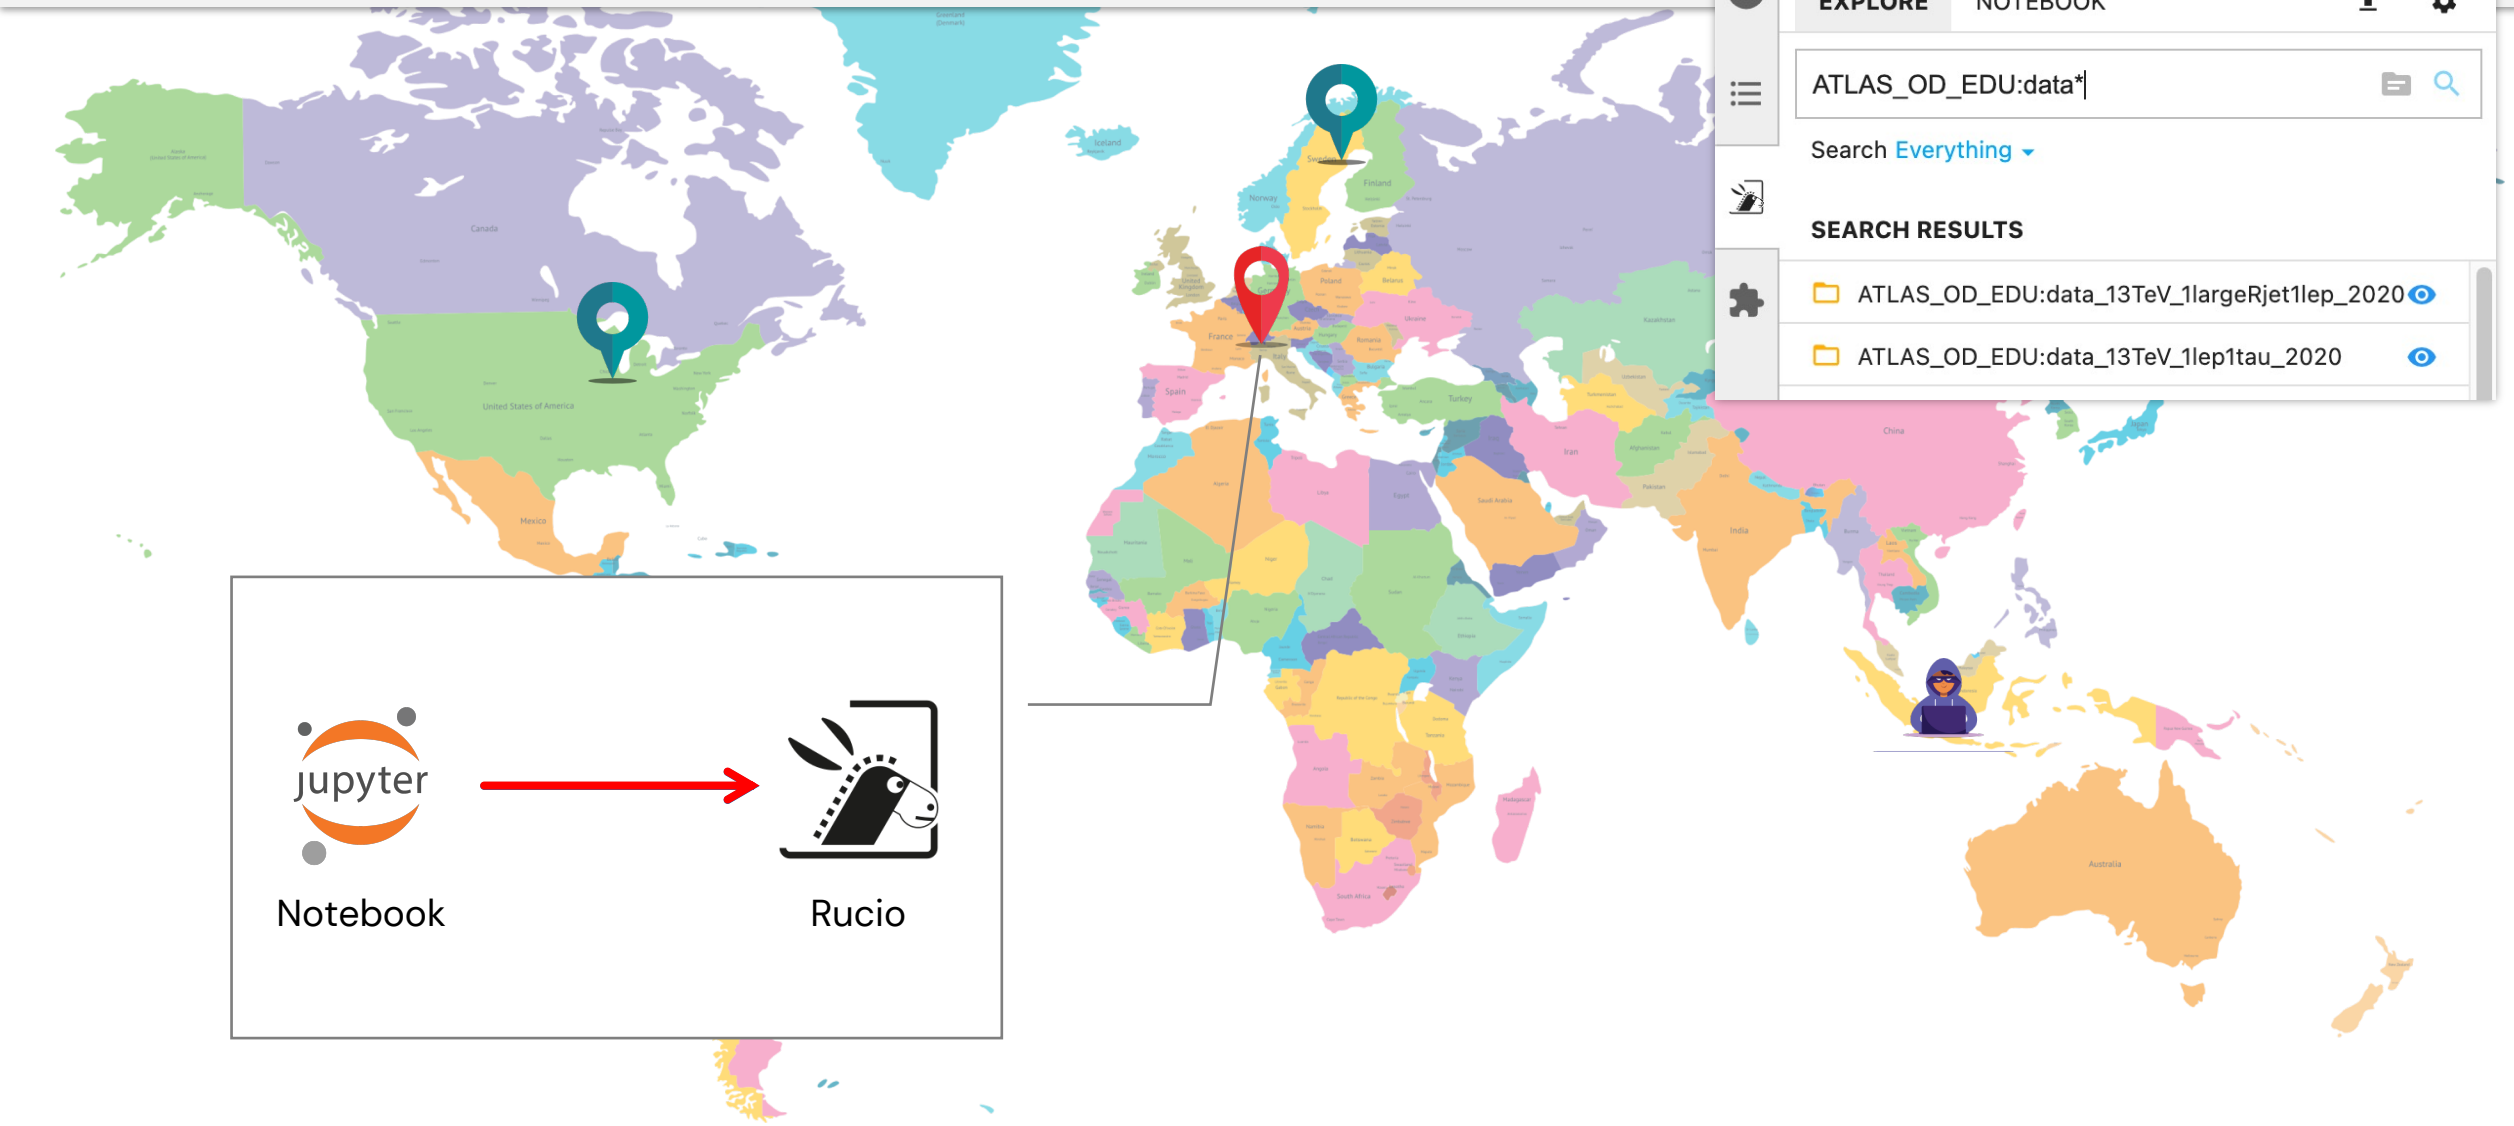

→ Make Available is clicked

Rucio initiates a transfer from storage in the USA, Brazil, and Sweden to a storage at CERN in Geneva

→ User runs a code that reads the files

The code runs on the notebook server at CERN in Geneva, and the output is shown to the user

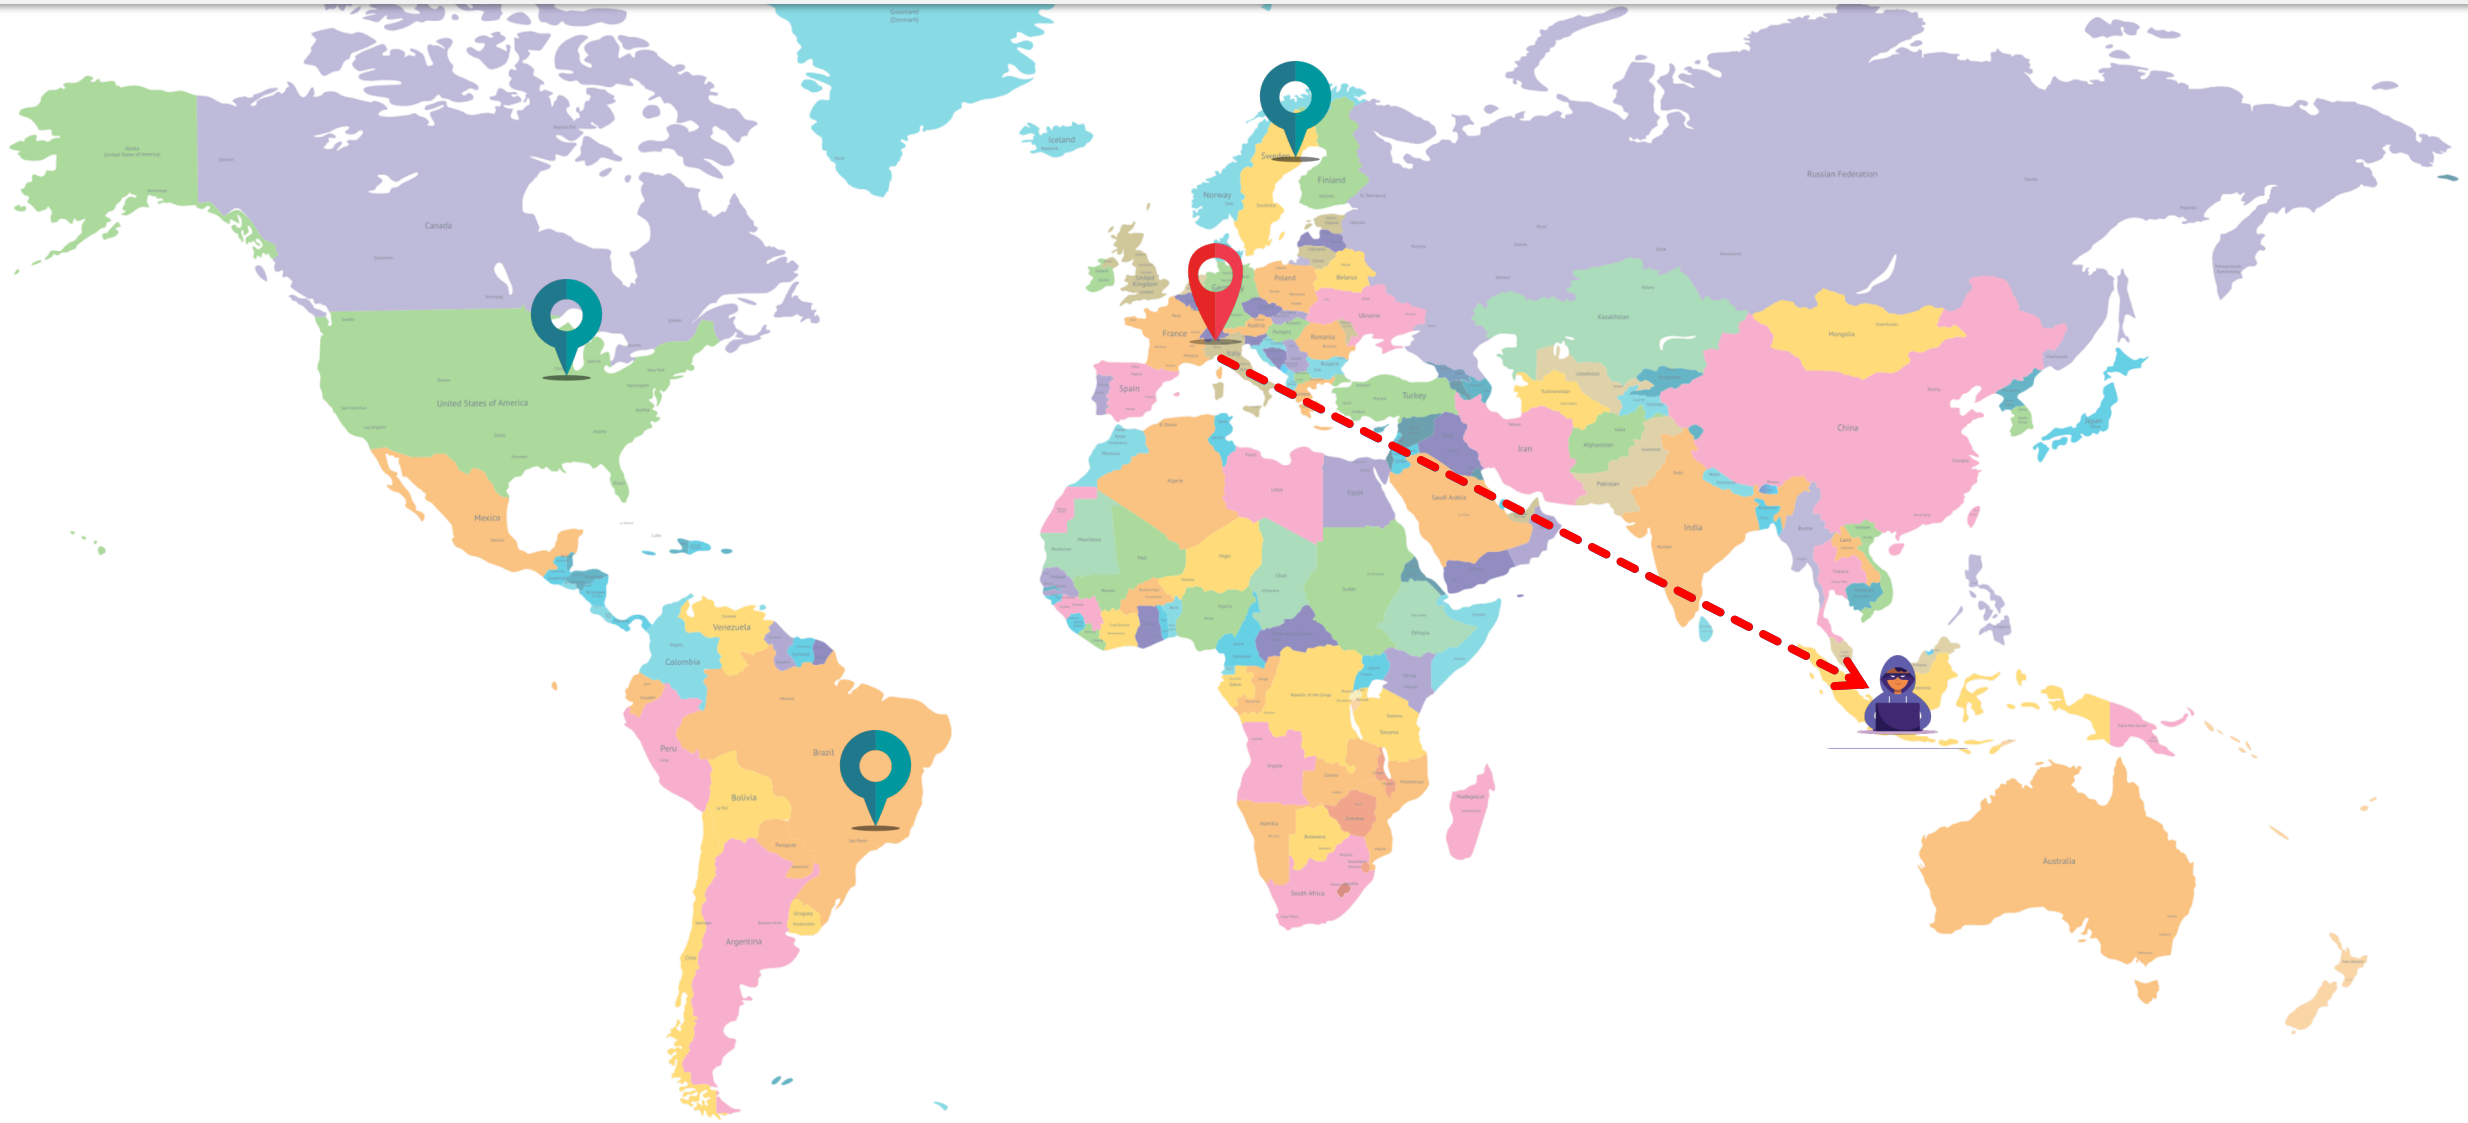

Implementation

- 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

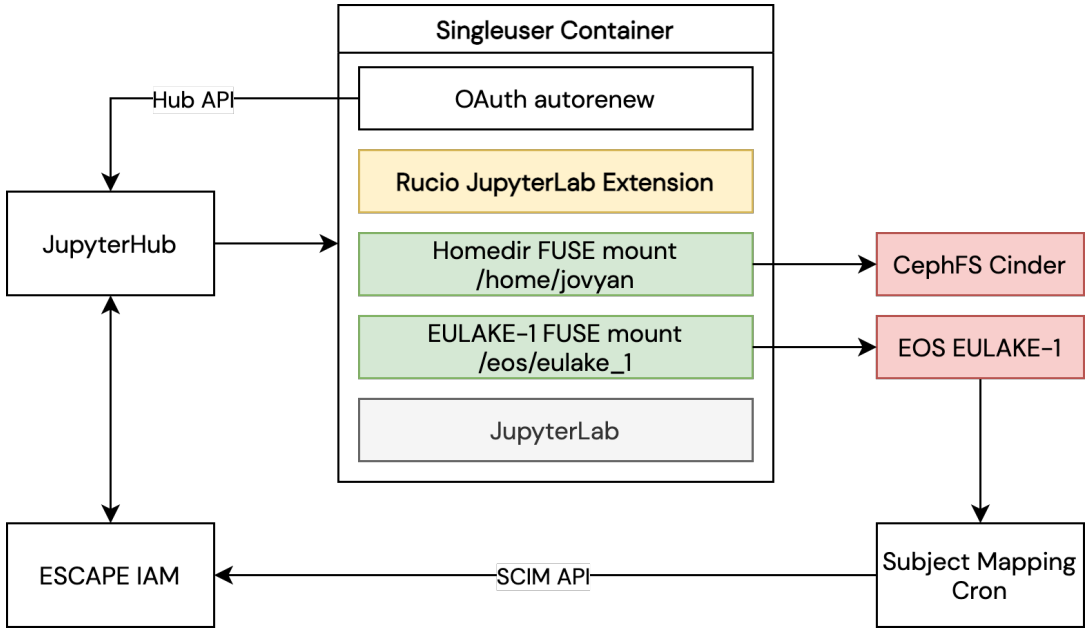

Use Cases

- 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.

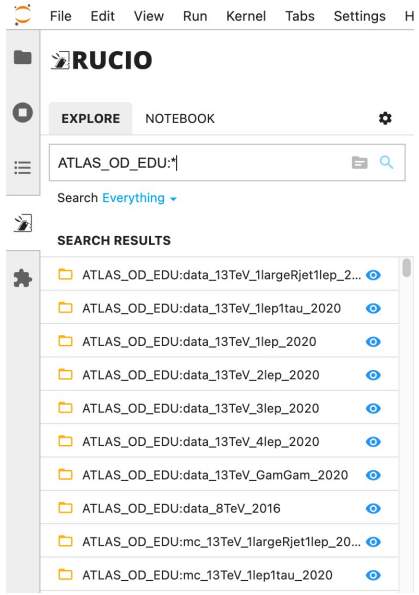


```
[3]: for item in hyy_20:
         print(item.pfn)
```

```
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.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_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

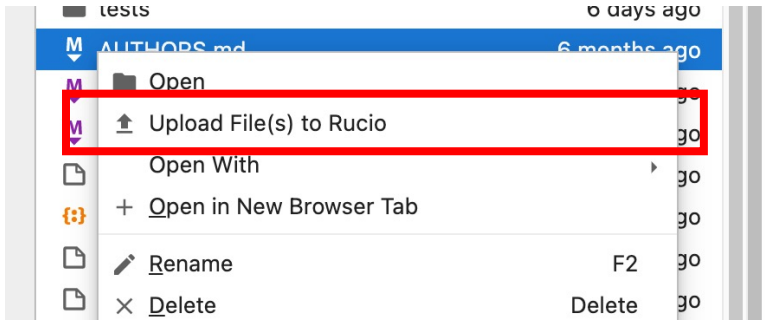

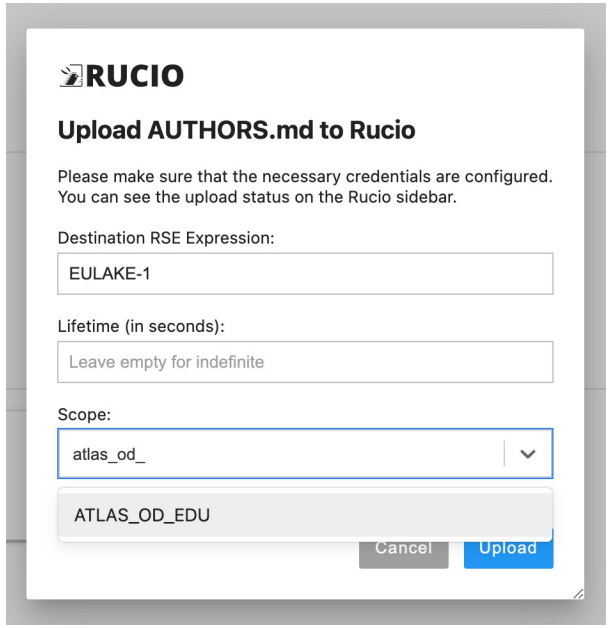

Conclusion

- The goal: hide the complexities of the Data Lake from the end users
  - We want to keep them happy and productive, after all ☺
- Lots of experiments are trying out Rucio Data Management system
  - Data Lake-as-a-Service is potentially interesting for both expert users and newbies.
- Community-driven development
  - Development of the Data Lake — and by extension the DLaaS — is driven by the needs of different experiments and sciences.

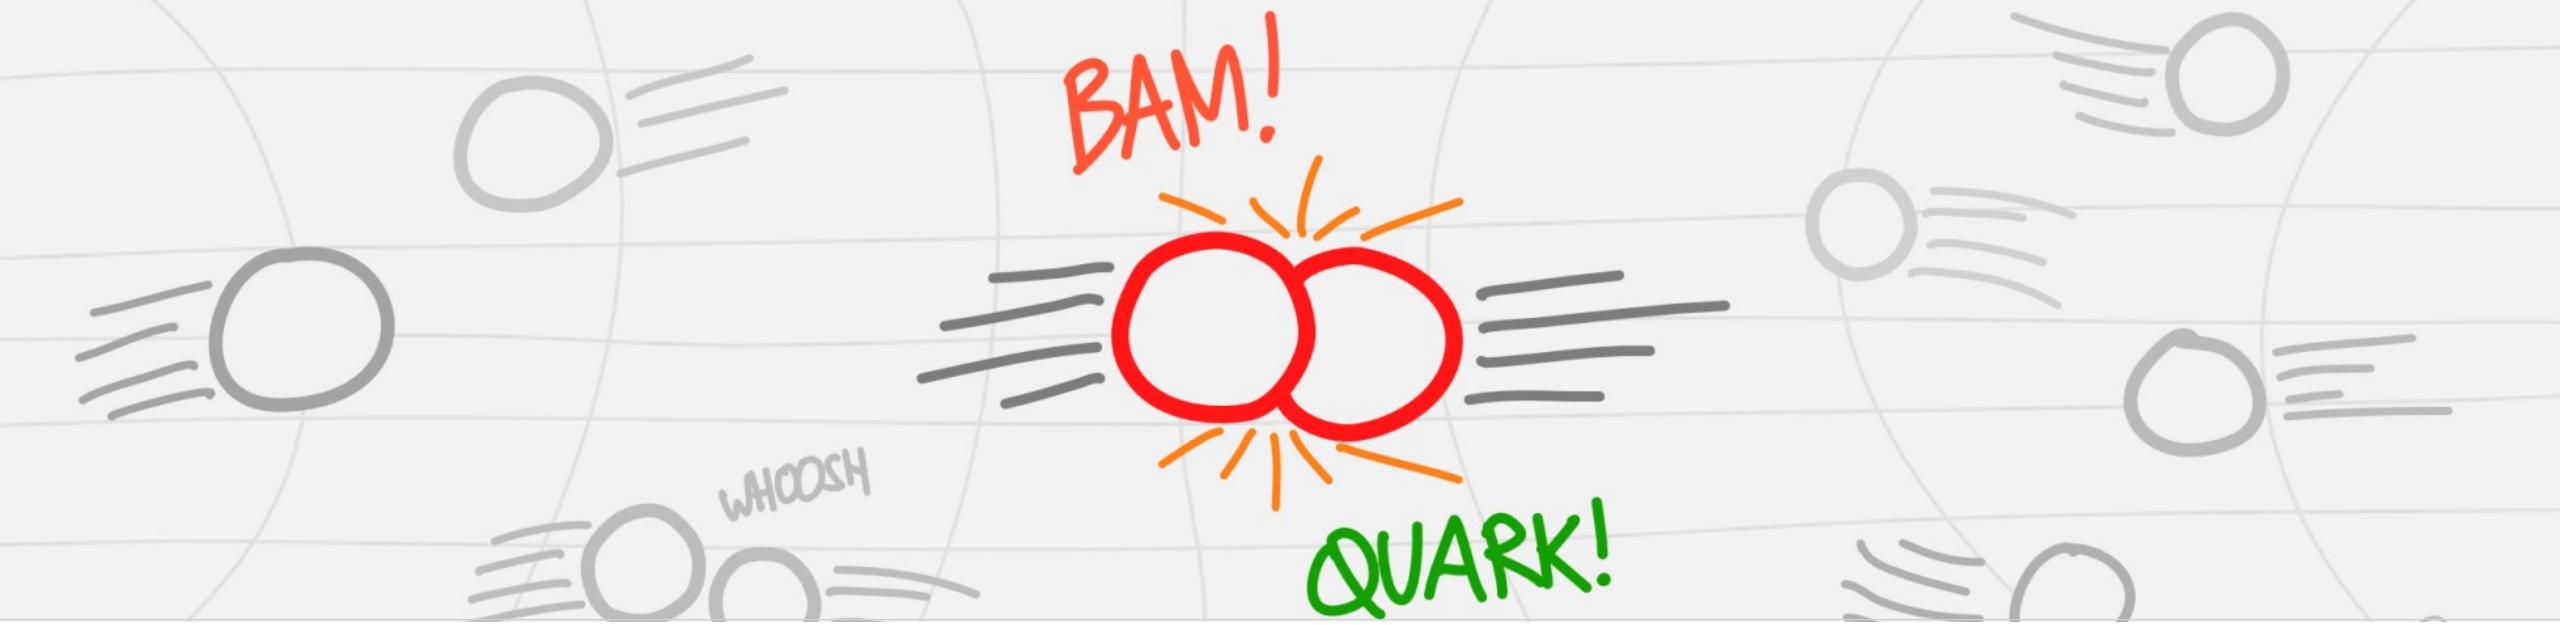

Thank you.

Attributions:
The SpongeBob GIF is a copyright of Viacom International
Cat illustration is taken from blush.design
World map and pin icons are taken from Freepik

Backup slides

Just in case you need more of these :)

Future Developments

- More kernel compatibility
  - Currently, only Python is 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
  - Will be completely transparent from the user PoV
  - Successfully tested at small scale

Desktop Data Lake as a Service

An installable package that gives the possibility to connect to the Data Lake seamlessly.

(A preconfigured JupyterLab installation in Anaconda Navigator could be an option)