

# Current CERN platforms for reproducible and interactive scientific analysis

D. Castro, A. Delgado, E. García, **G. Guerrieri**, B. Jones, A. Sciabà, T. Simko, M. Schulz, E. Tejedor (CERN IT)

February 26<sup>th</sup> 2025

# What is an analysis facility?

## Analysis Facilities White Paper

D. Ciangottini<sup>1,b</sup>, A. Forti<sup>2,b</sup>, L. Heinrich<sup>3,b</sup>, N. Skidmore<sup>4,b</sup>,  
C. Alpigiani<sup>5</sup>, M. Aly,<sup>2</sup> D. Benjamin<sup>6</sup>, B. Bockelman<sup>7</sup>, L. Bryant<sup>8</sup>, J. Catmore<sup>9</sup>, M. D'Alfonso<sup>35</sup>, A.  
Delgado Peris<sup>17</sup>, C. Doglioni<sup>2</sup>, G. Duckeck<sup>10</sup>, P. Elmer<sup>11</sup>, J. Eschle<sup>12</sup>, M. Feickert<sup>13</sup>, J. Frost<sup>14</sup>, R.  
Gardner<sup>8</sup>, V. Garonne<sup>6</sup>, M. Giffels<sup>36</sup>, J. Gooding<sup>15</sup>, E. Gramstad<sup>9</sup>, L. Gray<sup>16</sup>, B. Hegner<sup>31</sup>, A.  
Held<sup>13</sup>, J. Hernández<sup>17</sup>, B. Holzman<sup>16</sup>, F. Hu<sup>8</sup>, B. K. Jashal<sup>18,19</sup>, D. Kondratyev<sup>20</sup>, E. Kourlitis<sup>3</sup>, L.  
Kreczko<sup>21</sup>, I. Krommydas<sup>22</sup>, T. Kuhr<sup>10</sup>, E. Lancon<sup>6</sup>, C. Lange<sup>23</sup>, D. Lange<sup>11</sup>, J. Lange<sup>34</sup>, P. Lenzi<sup>1</sup>,  
T. Linden<sup>24</sup>, V. Martinez Outschoorn<sup>27</sup>, S. McKee<sup>25</sup>, J. F. Molina<sup>17</sup>, M. Neubauer<sup>26</sup>, A. Novak<sup>35</sup>, I.  
Osborne<sup>11</sup>, F. Ould-Saada<sup>9</sup>, A. P. Pages<sup>28</sup>, K. Pedro<sup>16</sup>, A. Perez-Calero Yzquierdo<sup>17</sup>, S. Piperov<sup>20</sup>, J.  
Pivarski<sup>11</sup>, E. Rodrigues<sup>29</sup>, N. Sahoo<sup>30</sup>, A. Sciaba<sup>31</sup>, M. Schulz<sup>31</sup>, L. Sexton-Kennedy<sup>16</sup>, O.  
Shadura<sup>32</sup>, T. Šimko<sup>31</sup>, N. Smith<sup>16</sup>, D. Spiga<sup>1</sup>, G. Stark<sup>33</sup>, G. Stewart<sup>31</sup>, I. Vukotic<sup>8</sup>, G. Watts<sup>5</sup>,

*<sup>b</sup> Editor, <sup>1</sup> INFN, <sup>2</sup> University of Manchester, <sup>3</sup> Technische Universität München, <sup>4</sup> University of Warwick <sup>5</sup> University of Washington <sup>6</sup> Brookhaven National Laboratory <sup>7</sup> Morgridge Institute for Research <sup>8</sup> University of Chicago <sup>9</sup> University of Oslo <sup>10</sup> Ludwigs-Maximilians Universität München <sup>11</sup> Princeton University <sup>12</sup> Syracuse University <sup>13</sup> University of Wisconsin-Madison <sup>14</sup> University of Oxford <sup>15</sup> Technische Universität Dortmund <sup>16</sup> FNAL <sup>17</sup> CIEMAT <sup>18</sup> IFIC <sup>19</sup> TIFR <sup>20</sup> Purdue University <sup>21</sup> University of Bristol <sup>22</sup> Rice University <sup>23</sup> Paul Scherrer Institute <sup>24</sup> Helsinki Institute of Physics <sup>25</sup> University of Michigan <sup>26</sup> University of Illinois at Urbana-Champaign <sup>27</sup> University of Massachusetts Amherst <sup>28</sup> IFAE <sup>29</sup> University of Liverpool <sup>30</sup> University of Birmingham <sup>31</sup> CERN <sup>32</sup> University of Nebraska-Lincoln <sup>33</sup> SCIPP UC Santa Cruz <sup>34</sup> University of Hamburg <sup>35</sup> Massachusetts Institute of Technology <sup>36</sup> Karlsruhe Institute of Technology*

# What is an analysis facility?

**Analysis facility (AF):** infrastructure and services that provide integrated data, software and computational resources to execute one or more elements of an analysis workflow.

- Perform fast research iterations on large datasets interactively.
- Convert interactive to batch-schedulable workloads.
- Interact with the WLCG and scale outside of the facility on occasion.
- Reproducibly instantiate desired software stack.
- Collaborate in a multi-organisational team.
- Efficiently train machine learning models for HEP.

# The CERN ecosystem

## A facility of Analysis Facilities



LxBatch



reana

LxPlus



HTCondor

Kubeflow



CVMFS

# SWAN at CERN

## Service for Web based ANalysis

### Jupyter notebooks technology

- Interactive analysis with a web browser
  - No local installation is needed
  - Based on Jupyter Notebooks
  - Calculations, input data and results “in the Cloud”
- Good for data analysis and exploration, but also for teaching
- Easy sharing of scientific results: plots, data, code
- Added value: integration with CERN infrastructure and services!



# SWAN at CERN

## Spinning up the service

Home Token

### Configure Environment

Specify the parameters that will be used to contextualise the container which is created for you. See [SWAN service website](#) for more details and contact to administrators.

Try out our new experimental interface based on **JupyterLab** and let us know your feedback!

User Interface [more...](#)

Try the new JupyterLab interface (experimental)

Software stack [more...](#)

106a

Use Python packages installed on CERNBox

Platform [more...](#)


AlmaLinux 9 (gcc13)

Environment script [more...](#)

e.g. `$CERNBOX_HOME/MySWAN/myscript.sh`

Number of cores [more...](#)

Start my Session



File Edit View Run Kernel Git Tabs Settings Help

Launcher

SWAN\_projects

Notebook

Python 3 julia 1.10.4 Octave R ROOT C++

Console

Python 3 julia 1.10.4 Octave R ROOT C++

Other

\$ \_

Simple 0 Mem: 125.63 / 8192.00 MB

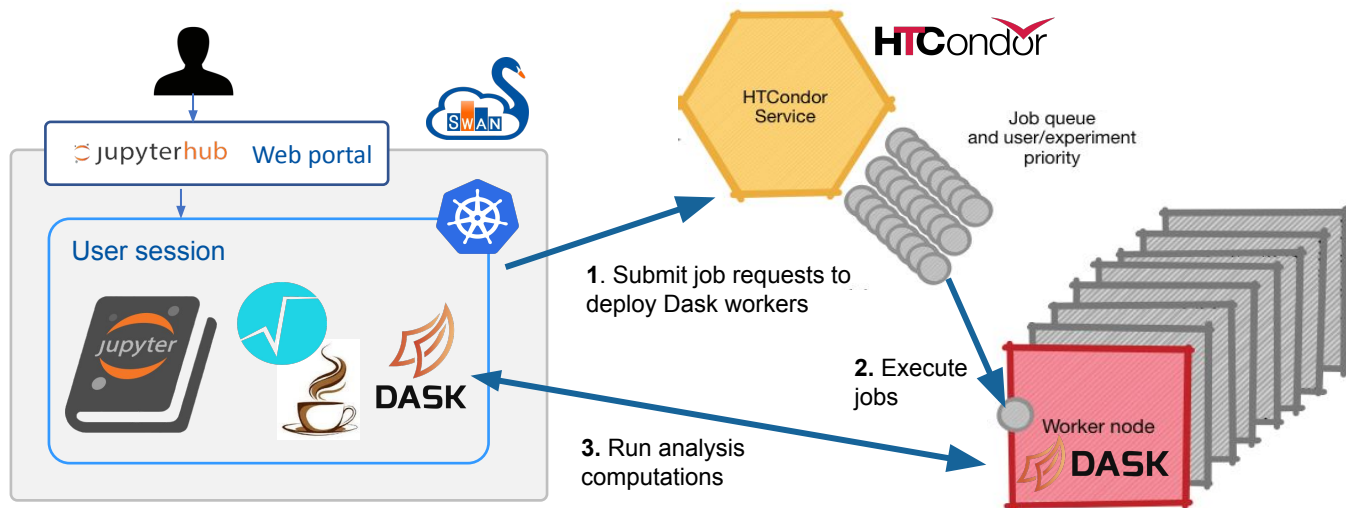
# The CERN pilot Analysis Facility

## Scaling columnar interactive analysis to batch

ref.

Leverages CERN resources (HTCondor)

- Via RDataFrame / coffea + Dask




# The CERN pilot Analysis Facility

## Scaling interactive analysis to batch

Pilot's interactive jobs benefit from dedicated resources in the CERN HTCondor pool

- Dedicated buffer of ~1.2k cores for quick (number) allocation of resources
- Extra resources from common pool, job start time **subject to experiment quotas**





# Best test users for the pilot

## User's profile

- Have a columnar-based analysis (**RDataFrame** or **Coffea**) that leverages Dask.
- Be able to submit condor batch jobs on lxplus.
- Have data stored on eos.

## What feedback we would like to have

- Clarity of documentation.
- What **you like**, what you **don't like** and **what you would like**.
- Signal any possible improvements so that it suits better the needs of an analyser.
- Report issues with interface, resource allocation, analysis execution, data access, software, ...

More information on SWAN and the pilot at  
<https://swan.docs.cern.ch>







Do you have a use case for the pilot?  
Join the [Mattermost Channel](#)

*(First join the [IT-dep mattermost team](#))*

# What is an analysis facility?

## A quick recap

**Analysis facility:** infrastructure and services that provide integrated data, software and computational resources to execute one or more elements of an analysis workflow.

-  Perform fast research iterations on large datasets interactively.
-  Convert interactive to batch-schedulable workloads.
-  Interact with the WLCG and scale outside of the facility on occasion.
-  Execute desired analysis workflows reproducibly.
-  Collaborate in a multi-organisational team.
-  Efficiently train machine learning models for HEP.

## Reproducible research data analysis platform

REANA is a reproducible analysis platform allowing scientists to run **declarative computational data analysis** pipelines on containerised compute clouds.

REANA was built with the goal of **fostering computational reproducibility**.

It has installations and code distributions in Europe (e.g. [AIP](#)) and US (e.g. [UChicago](#))

### Flexible

Run many computational workflow engines.



### Scalable

Support for remote compute clouds.



### Reusable

Containerise once, reuse elsewhere. Cloud-native.

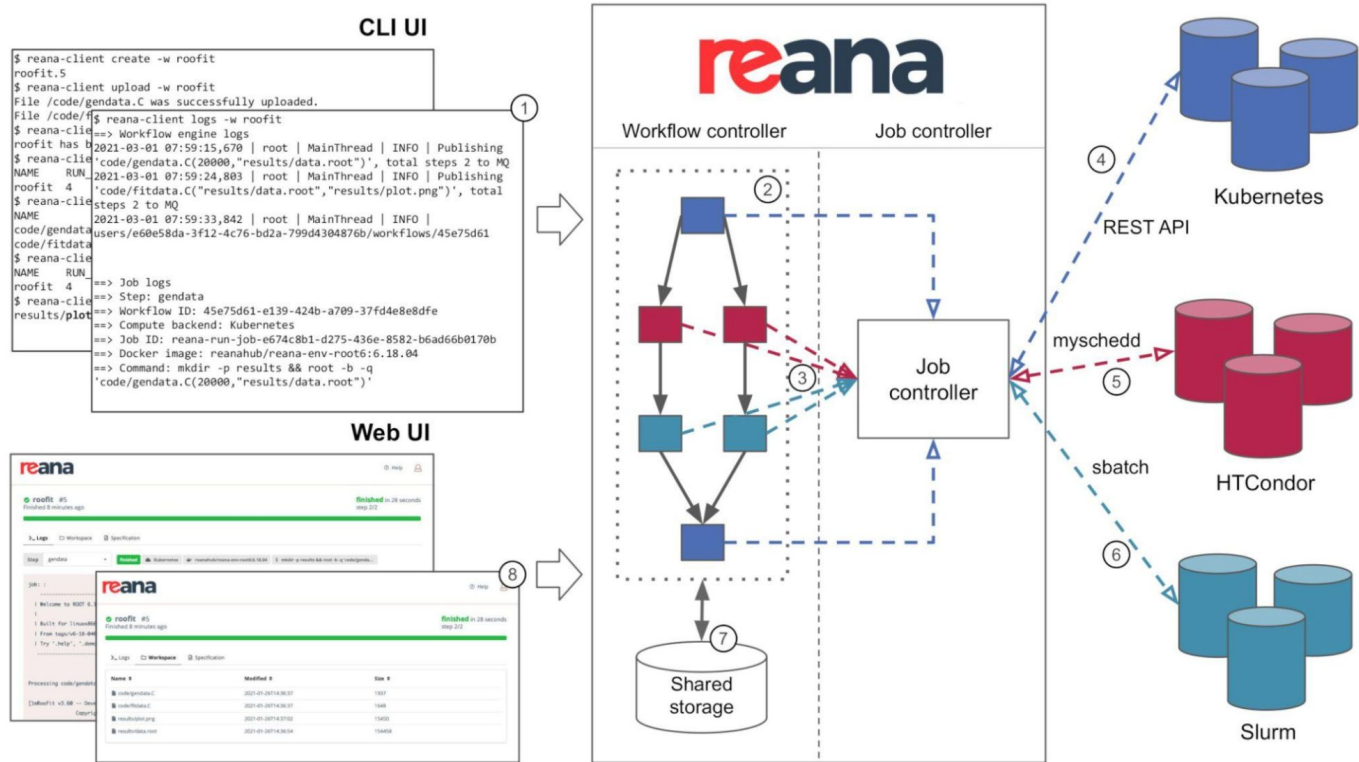


### Free

Free Software. MIT licence. Made with ❤️ at CERN.



## Reproducible research data analysis platform



### Multiple compute backends:

- Kubernetes
- HTCondor
- Slurm

### Multiple workflow languages:

- CWL
- Serial
- Snakemake
- Yadage

### Multiple means of use:

- Command-line client
- Web UI

# REANA

## A quickstart

ref.

```
# workflow/snakemake/Snakefile
```

```
rule all:
  input:
    "results/data.root",
    "results/plot.png"
```

```
rule gendata:
  output:
    "results/data.root"
  params:
    events=20000
  container:
    "docker://docker.io/reanahub/reana-env-root6:6.18.04"
```

```
shell:
  """
  mkdir -p results
  root -b -q 'code/gendata.C({params.events},{output})'
  """
```

```
rule fitdata:
  input:
    data="results/data.root"
  output:
    "results/plot.png"
  container:
    "docker://docker.io/reanahub/reana-env-root6:6.18.04"
  shell:
    "root -b -q 'code/fitdata.C({input.data},{output})'"
```

1

### Structure your analysis

```
inputs:
  files:
    - code/gendata.C
    - code/fitdata.C
  directories:
    - workflow/snakemake
  parameters:
    input: workflow/snakemake/inputs.yaml
  workflow:
    type: snakemake
    file: workflow/snakemake/Snakefile
  outputs:
    files:
      - results/plot.png
```

more

2

### Select a REANA cluster...

```
$ export REANA_SERVER_URL=https://
reana.cern.ch/
$ export REANA_ACCESS_TOKEN=XXXXXX
```

### ...or install your own

```
# install kubect1 1.19+, kind 0.9+ and
helm 3.0+
$ sudo dpkg -i kubect1*.deb kind*.deb
kubernetes-helm*.deb
# create Kubernetes cluster (or use your
own!)
$ wget https://raw.githubusercontent.com/
reanahub/reana/maint-0.9/etc/kind-
localhost-30443.yaml
$ kind create cluster --config kind-
localhost-30443.yaml
$ wget https://raw.githubusercontent.com/
reanahub/reana/maint-0.9/scripts/prefetch-
images.sh
$ sh prefetch-images.sh
# deploy REANA using Helm
$ helm repo add reanahub https://
reanahub.github.io/reana
$ helm repo update
$ helm install reana reanahub/reana --wait
# create an admin user
$ wget https://raw.githubusercontent.com/
reanahub/reana/maint-0.9/scripts/create-
admin-user.sh
$ sh create-admin-user.sh default reana
john.doe@example.org mysecretpassword
```

more

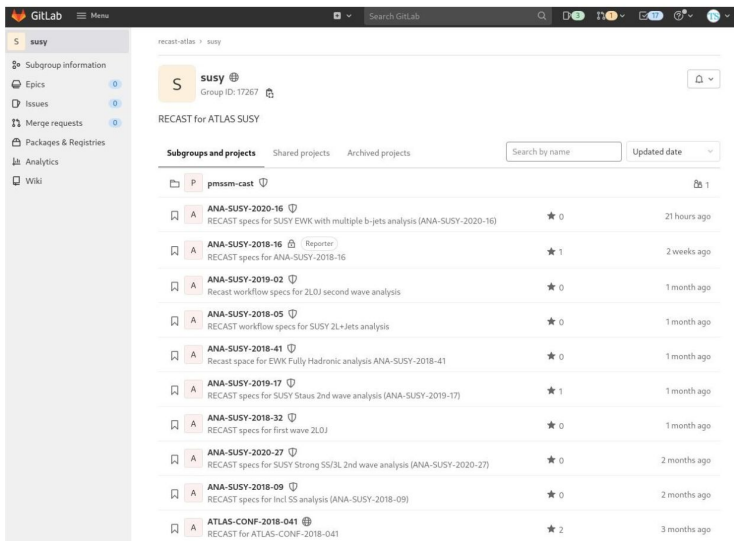
3

### Run your analysis

```
# create new virtual environment
$ virtualenv ~/.virtualenvs/myreana
$ source ~/.virtualenvs/myreana/bin/
activate
# install REANA client
$ pip install reana-client
# create new workflow
$ reana-client create -n my-analysis
$ export REANA_WORKON=my-analysis
# upload input code and data to workspace
$ reana-client upload ./code ./data
# start computational workflow
$ reana-client start
# check its progress
$ reana-client status
# list workspace files
$ reana-client ls
# open interactive notebook session
$ reana-client open jupyter
# download output results
$ reana-client download results/plot.png
```


more

## Usage example: pMSSM reinterpretation of ATLAS Run-2 searches




**Figure 1.** A screenshot of the ATLAS SUSY group analyses preserved on GitLab. Each repository is labeled with the internal ATLAS analysis identifier and contains both workflow files and additional data files needed for the computational processing.

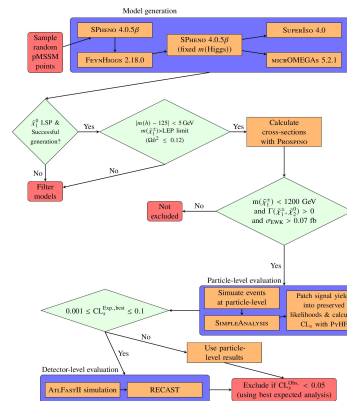
<https://arxiv.org/abs/2403.03494>



**Figure 2.** A typical pMSSM workflow. The computational runtime is about 10 minutes without systematics (test payload) and about 10 hours with all systematics (real payload).



<https://arxiv.org/abs/2402.01392>









More information on **REANA** at  
<https://www.reanahub.io/>



# What is an analysis facility?

## A quick recap

**Analysis facility:** infrastructure and services that provide integrated data, software and computational resources to execute one or more elements of an analysis workflow.

-  Perform fast research iterations on large datasets interactively.
-  Convert interactive to batch-schedulable workloads.
-  Interact with the WLCG and scale outside of the facility on occasion.
-  Execute desired analysis workflows reproducibly.
-  Collaborate in a multi-organisational team.
-  Efficiently train machine learning models for HEP.

I don't have a CERN account 😞

# ESCAPE: European Science Cluster of Astronomy and Particle Physics



Consortium of 31 members, including:

- 10 [ESFRI](#) projects & landmarks: CTA, EST, FAIR, HL-LHC, KM3NeT, SKA, LSST, VIRGO, ESO, JIVE
- 2 pan-European International Organizations: CERN and ESO
- 2 European Research Infrastructures: EGO and JIV-ERIC
- 4 supporting European consortia: APPEC, ASTRONET, ECFA and NuPECC


Budget: **15.98 M€**

Duration: **48 months** (1/2/2019 -31/1/2023)

ESCAPE has received funding from the European Union's Horizon 2020 research and innovation programme under **Grant Agreement no. 824064**.

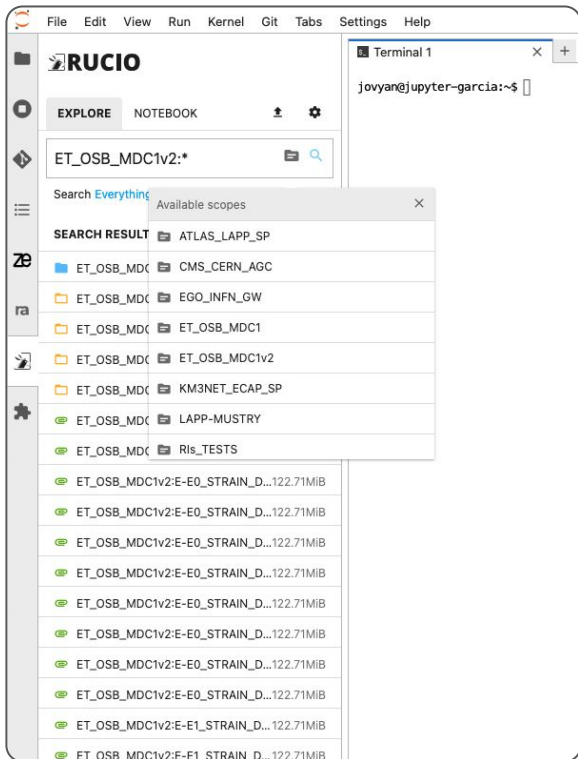
# The CERN Virtual Research Environment

A CERN-agnostic version of SWAN

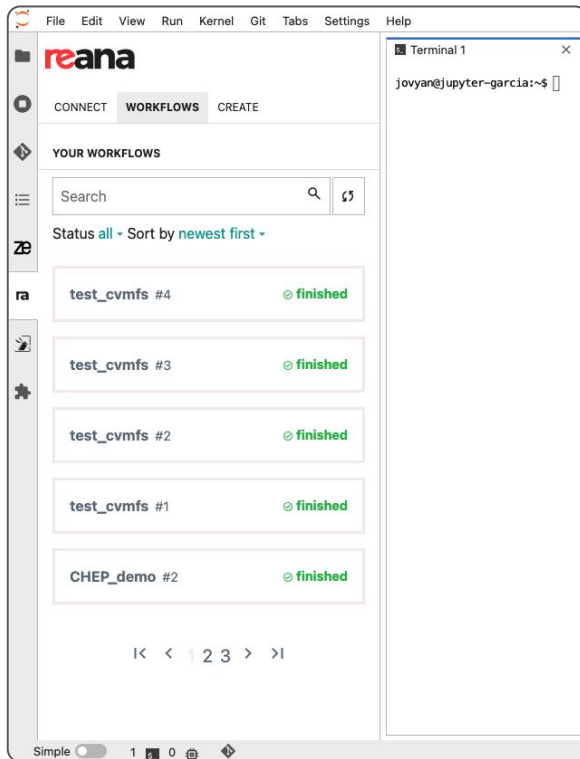


# User's interactions with the VRE

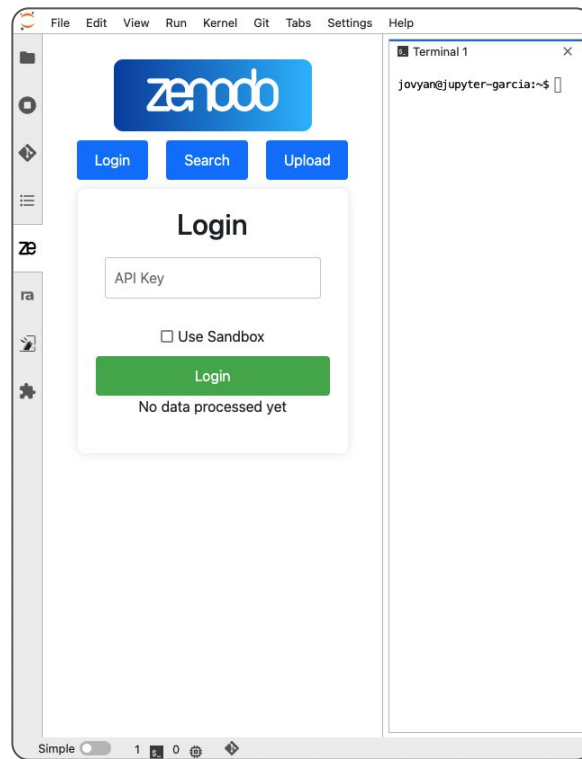
See more in [Sukanya's presentation](#)



The screenshot shows the RUCIO web interface. The main content area displays a search for 'ET\_OSB\_MDC1v2:\*'. A search results dropdown is open, showing a list of search results. The results include folders like 'ATLAS\_LAPP\_SP', 'CMS\_CERN\_AGC', 'EGO\_INFN\_GW', 'ET\_OSB\_MDC1', and 'ET\_OSB\_MDC1v2'. Below these, there is a list of specific file paths, such as 'ET\_OSB\_MDC1v2:E-E0\_STRAIN\_D...122.71MIB'.



The screenshot shows the reana web interface. The main content area displays a list of workflows under the heading 'YOUR WORKFLOWS'. The workflows listed are 'test\_cvmsfs #4', 'test\_cvmsfs #3', 'test\_cvmsfs #2', 'test\_cvmsfs #1', and 'CHEP\_demo #2'. All of these workflows are marked as 'finished' with a green checkmark icon. The interface also includes a search bar and a status filter set to 'all'.



The screenshot shows the zenodo web interface. The main content area displays a login form with the heading 'Login'. The form includes a text input field for 'API Key', a checkbox for 'Use Sandbox', and a green 'Login' button. Below the button, it says 'No data processed yet'. The interface also includes a search bar and an 'Upload' button.

# Comparison

(is the thief of joy)



- Available at the click of a button ✓
- Data persistence ✓
- No account restrictions ✗
- Change easily the software stack ✓
- Access >200 cores ✓
- Access to batch ✓
- Access to GPUs ✓
- Integration with Data Lake tech. ✗
- Integration with REANA ✗

**Production service**  
Can sustain heavy workloads



- ✓ Available at the click of a button
- ✓ Data persistence
- ✓ No account restrictions
- ✓ Change easily the software stack
- ✗ Access >200 cores
- ✗ Access to batch
- ✗ Access to GPUs
- ✓ Integration with Data Lake tech.
- ✓ Integration with REANA







**Proof of concept platform**  
Features more functionalities

**More information on the VRE at**  
<https://vre-hub.github.io/>

# What is an analysis facility?

## A quick recap

**Analysis facility:** infrastructure and services that provide integrated data, software and computational resources to execute one or more elements of an analysis workflow.

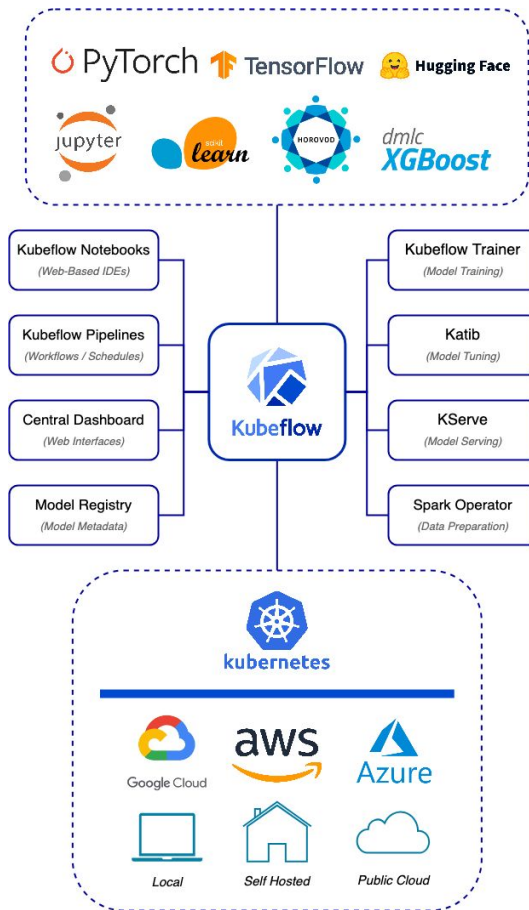
-  Perform fast research iterations on large datasets interactively.
-  Convert interactive to batch-schedulable workloads.
-  Interact with the WLCG and scale outside of the facility on occasion.
-  Execute desired analysis workflows reproducibly.
-  Collaborate in a multi-organisational team.
-  Efficiently train machine learning models for HEP.



### What is Kubeflow?

Kubeflow is a community and ecosystem of **open-source** projects to address each stage in the ML lifecycle







- **Community & Ecosystem:** Open-source projects addressing each stage of the ML lifecycle.
- **Purpose:** Simplifies, makes AI/ML on Kubernetes portable and scalable.
- **Audience:** Researchers, data scientists, ML engineers, developer teams.
- **Modular & Scalable Tools:** From building to deploying ML models for AI applications.



More information on ml.cern.ch at  
<https://ml.cern.ch/>

# Summary

**Analysis facility:** infrastructure and services that provide integrated data, software and computational resources to execute one or more elements of an analysis workflow.

-  Perform fast research iterations on large datasets interactively.
-  Convert interactive to batch-schedulable workloads.
-  Interact with the WLCG and scale outside of the facility on occasion.
-  Execute desired analysis workflows reproducibly.
-  Collaborate in a multi-organisational team.
-  Efficiently train machine learning models for HEP.