


# Cold Storage support on CERN Open Data Portal

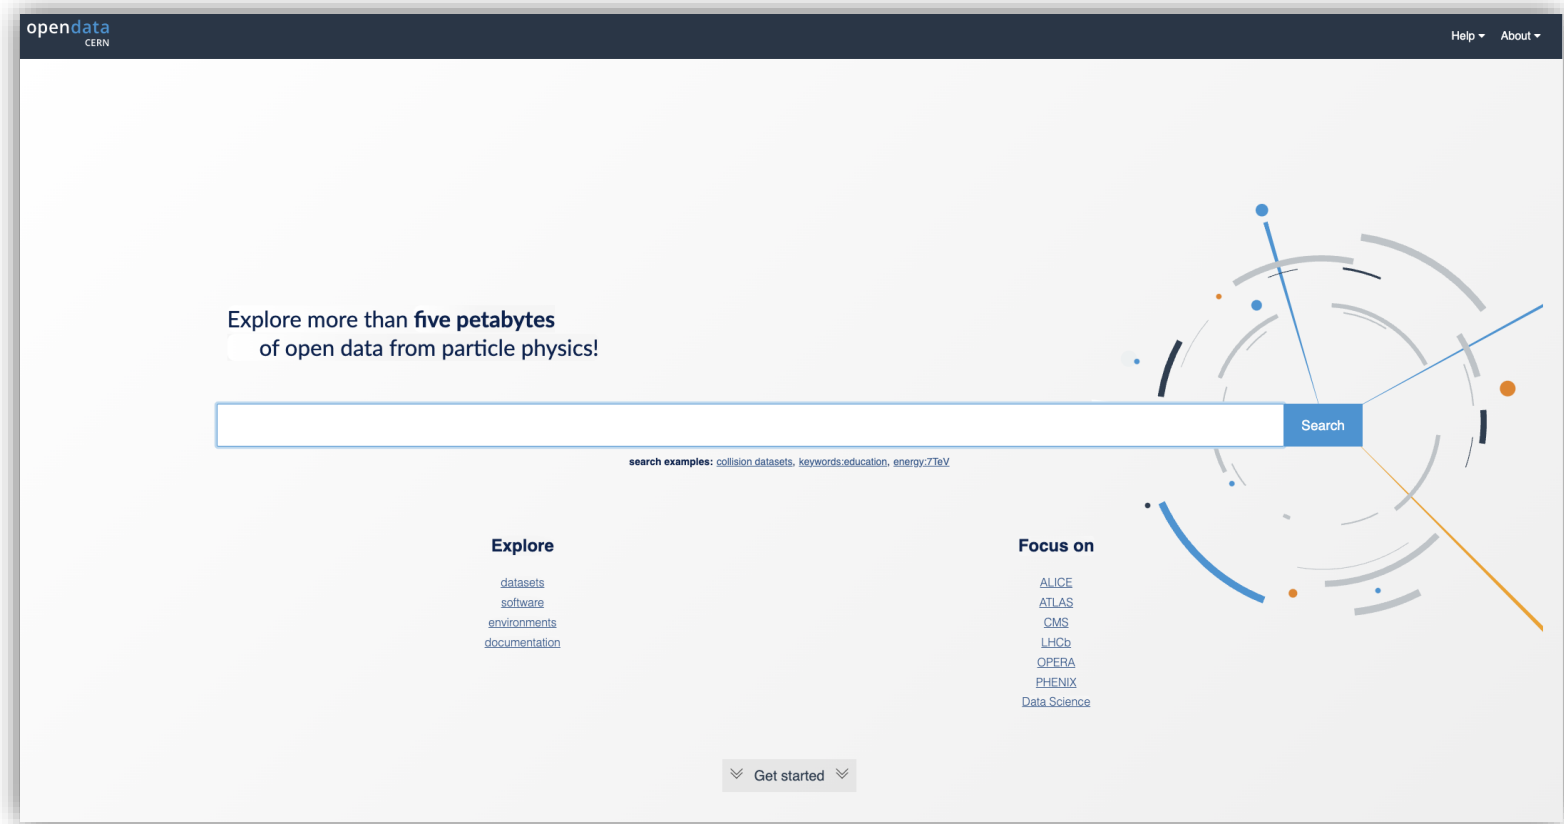
Pablo Saiz

21<sup>st</sup> Oct 2024



# CERN Open Data Portal

- **Repository of data**
  - Launched in November 2014
- **Plenty of content**
  - Dataset
    - Collision, simulated and derived
  - Documentation
    - Glossary, tutorials, configuration, examples
  - Software
    - Virtual machines, containers
- **Current Size**
  - > 5 PB
  - > 70k entries
  - > 2.5 million files




# Recent updates in content

**ATLAS releases 65 TB of open data for research**  
2024-07-01 by ATLAS Collaboration

**Explore over 7 billion LHC collision events - from home**  
The ATLAS Experiment at CERN has made two years' worth of scientific data available to a collision energy of 13 TeV. This is the first time that ATLAS has released data on this scale.

**LHCb releases the entire Run I data**  
2023-12-20 by LHCb Collaboration

**CMS releases 13 TeV data**  
2024-04-02 by CMS Collaboration



have been made available to the global community within just ten years.

...ing purposes only. The content may be erased from time to time. Please use

...ed events. It allows to view events as they turn up in the detector, examine tracks, start up this tool.

...e it for testing purposes only. The content may be erased from time to time. Please use

...a from the JADE detector (DESY-PETRA)

...9 to 1986 with e+e- collision energies ranging from 12 to 46.7 GeV. JADE was one of the

# Data volumes

- Taken from J. Boyd at [4th DPHEP workshop](#)
  - Table updated since originally created, and consistent with previous estimates

	ALICE (TB)	ATLAS (TB)	CMS (TB)	LHCB (TB)	Sum (TB)	Cumulative (PB)
2022	15	0	1089	200	1304	1.3
2023	20	150	872	4600	5642	6.9
2024	105	0	1436	0	1541	8.5
2025	105	0	1768	0	1873	10.4
Total	245	150	5165	4800	10360	

# Cold Storage support

- **Given:**

- Amount of data: 5 PB, and increasing
- Long term data preservation and data immutability
- Accepted latency
- Multiple copies of data (?) (beware: data ownership)
- Business Continuity and Disaster Recovery

- **Are there any more efficient ways of long term archival?**

- **Possibilities:**


- Computing instead of storage
- Cold Storage. Move part of the data to cheaper storage (Tape, experiment frameworks)

P. Nogga: [LHCb Open Data Ntupling Service](#)

# Current use case



opendata  
CERN



opendata  
CERN

HIAIIPhysics primary dataset in RECO format from the 2.76 TeV Pb-Pb run of 2010 (/HIAIIPhysics/HIRun2010-ZS-v2/RECO)  
/HIAIIPhysics/HIRun2010-ZS-v2/RECO, CMS collaboration

Cite as: CMS collaboration (2020). HIAIIPhysics primary dataset in RECO format from the 2.76 TeV Pb-Pb run of 2010 (/HIAIIPhysics/HIRun2010-ZS-v2/RECO). CERN Open Data Portal.  
DOI:10.7483/OPENDATA.305G.POEC

There is one publication referring to these data

Dataset Collision Heavy-ion physics CMS 2.76TeV PpPb CERN-LHC

Description

HIAIIPhysics primary dataset from the 2.76 TeV Pb-Pb run of 2010.

The list of validated runs, which must be applied to all analyses, either with the full validation or with the muons only validation.

Validated runs, full validation

Validated runs, muons only

Dataset characteristics

75482193 events. 28954 files. 151.0 TiB in total.

Dataset characteristics

75482193 events. 28954 files. 151.0 TiB in total.

## Options


### 2. Download file index

### 1. Download files locally





# Request staging



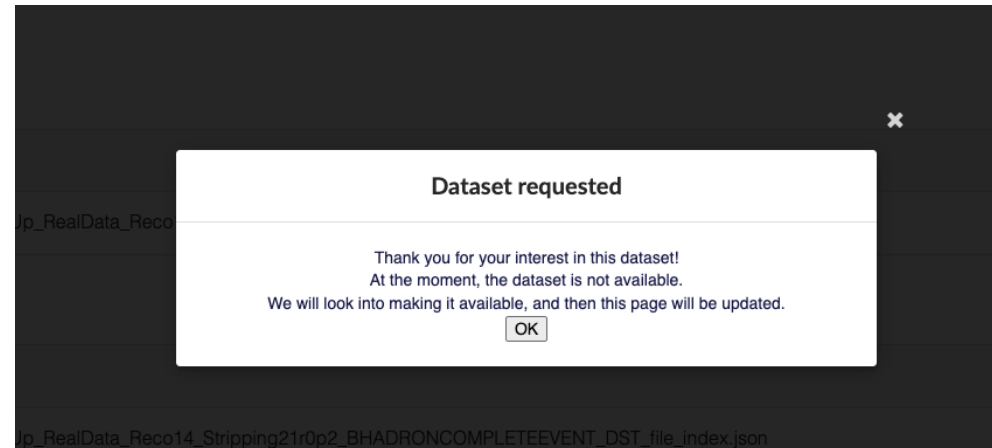
**Ensure that anonymous users can download data**



**Protect system from abuse**



**Notify users when data are available**





# File and record status

- **File can be:**

- Only on hot
- Only on cold
- Both

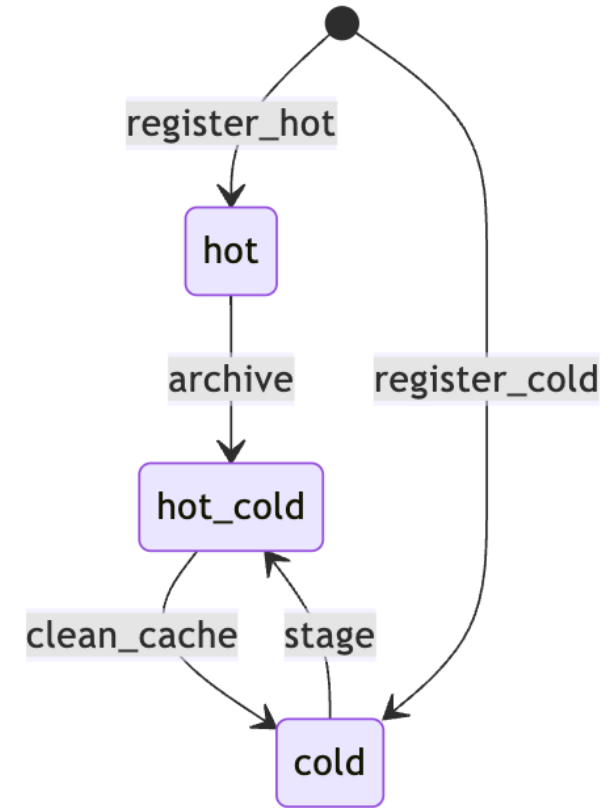
- **Record can have multiple files**

- All files on hot
- Some files on hot
- All files only on cold

```
$ cernopendata cold  
Usage: cernopendata cold [OPTIONS] COMMAND [ARGS]...
```

Commands:

archive Move a record to cold.  
clear-hot Delete the hot copy of a file that has a cold copy.  
list Print the urls for an entry.  
stage Move a record from cold.



# CLI for cold storage

- **Commands for the service managers**
  - **With the goal of automating them later on...**

## **\$ cernopendata cold list 15227**

The files referenced in '15227' are:

\* Hot copy: root://eospublic.cern.ch//eos/opendata/cms/...

\* Hot copy: root://eospublic.cern.ch//eos/opendata/cms/...

...

Summary: 10 files (3.062 GiB), with 10 hot copies (3.062 GiB) and 0 cold copies (0 bytes)

## **\$ cernopendata cold archive 15227**

Moving the record 15227 to cold:

Checking if the file CMS... transfer has been issued!

Checking if the file CMS... transfer has been issued!

...

10 transfers have been issued.

Record 15227 moved.

## **\$ cernopendata cold list 15227**

The files referenced in '15227' are:

\* Hot copy: root://eospublic.cern.ch//eos/opendata/cms/...

\* Cold copy: https://eosctapublic.cern.ch/eos/ctapublic/archive/opendata/cms/...

\* Hot copy: root://eospublic.cern.ch//eos/opendata/cms/...


\* Cold copy: https://eosctapublic.cern.ch/eos/ctapublic/archive/opendata/cms/...

Summary: 10 files (3.062 GiB), with 10 hot copies (3.062 GiB) and 10 cold copies (3.062 GiB)

# Cold storage challenges


## 1. Monitoring

1. Identify candidates for cold storage
2. Keep track of cold storage requests




## 2. Transfer files

- A. CERN Tape Archive: FTS
- B. (Experiment frameworks: Dedicated plugins per experiment)



## 3. Data sovereignty

1. Experiment still responsible for data (!)



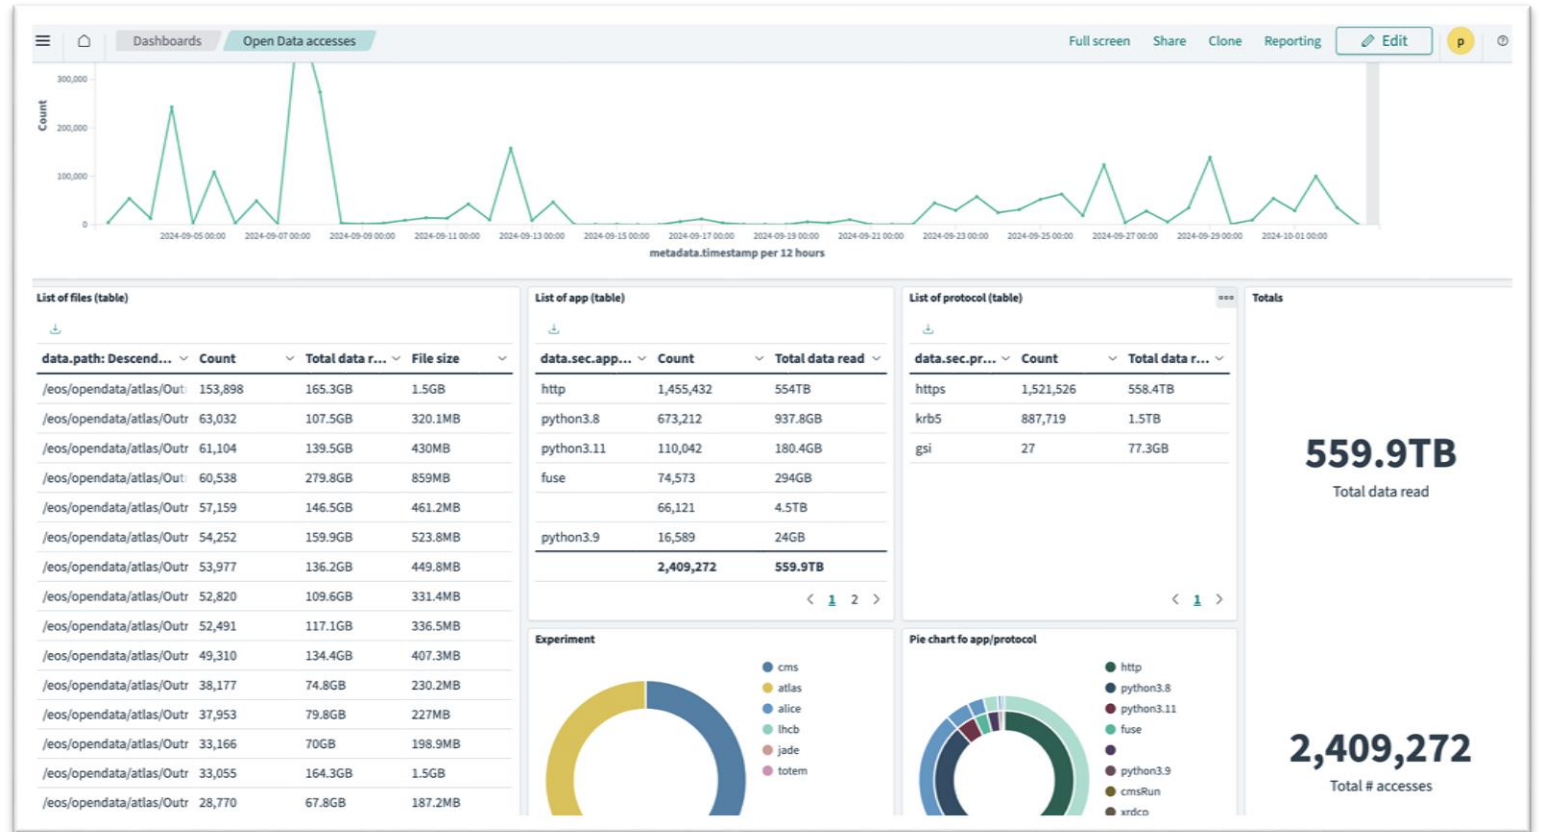
## 4. Multiple copies of datasets



# Internal dashboard for Service Admin

## Created early this year:

- Using logs from eos
- # accesses
- Volume transferred
- Breakdown per experiment
- Last 30-days
- **Ongoing:**
  - Keep long term stats
  - Map from file to record
  - Open it up
  - More breakdown categories
  - Prototype with MONIT-GRAFANA



# Cold storage: current status



1. Create tools to archive/stage/remove



2. Store information of multiple copies



3. Agree on disk/tape quotas



4. Modify UI: display/search hot copies; request cold copies



5. Automatic staging




(Optional) Implement other cold storage alternatives



# CERN Open Data Portal

- Digital repository with 5 PB of high energy physics data
  - Curated data: datasets, documentation, software...
  - 3 more in the process:



- Adding cold storage support:
  - Files might be stored on cheaper media, not immediately accessible.
  - Interface with FTS and CTA to archive/stage files
  - First round with manual operations triggered by user/curators requests
    - With the goal of automating it later on
  - Ongoing effort on improving the monitoring



[opendata-team@cern.ch](mailto:opendata-team@cern.ch)

<http://opendata.cern>



[home.cern](https://home.cern)