

Optimising data access analytics: Integrating dCache *BillingDB* with PIC's scalable big data platform

J. Flix, M. Santamaría, E. Planas, P. Tallada, J. Carretero

PIC
port d'informació
científica



HEPiX Autumn 2024

Norman, OK (USA)

04-08 Nov 2024

Introduction

Hadoop/Hive platform at PIC deployed **since Oct. 2016** for Cosmo/Astro projects (cosmohub.pic.es)



The image is a promotional graphic for Cosmo Hub, set against a dark, starry background. At the top center, the logo 'COSMO HUB' is displayed, with 'COSMO' in blue and 'HUB' in white. Below the logo, the text 'Build your own Universe' is written in white, followed by the subtitle 'Interactive data analysis of massive cosmological data without any SQL knowledge'. At the bottom, there are four blue icons: a star, a paper airplane, a square with a smaller square inside, and a bar chart. Each icon is accompanied by a short description of a feature.

COSMO HUB

Build your own Universe

Interactive data analysis of massive cosmological data without any SQL knowledge

- Billions of observed and simulated galaxies
- Superfast queries means superfast results
- Features to make you work faster and easier
- Online plotting preview and data download

Introduction

Hadoop/Hive platform at PIC deployed **since Oct. 2016** for Cosmo/Astro projects (cosmohub.pic.es)

Currently, it is **deployed on 32 nodes**:

- 768 cores, 16 TiB RAM
- 60 TB NVMe (for cache)
- 4.3 PiB raw storage (3 PiB usable)

It supports **multiple projects**: DES, PAUS, Euclid, MICE, LST, Gaia, LSST...

This platform offers **interactive exploration** (visualization) of catalogues, and capability to export results in multiple formats (Parquet, CSV, FITS, ASDF)

Stats: ~150 raw catalogs, ~2000 users registered, 17K custom catalogs
~300 tables, $3.27 \cdot 10^{11}$ NumRows (!), ~160 TB rawDataSize

>75% of all queries in <3 min - easy integration in **Jupyter Notebooks**

CosmoHub: in-house software stack

Software stack

- 1th gen: HDP 2.6.5
- 2nd gen: HDP 3.1.4
- **3rd gen: Shepherd (in-house)**

Hardware architecture

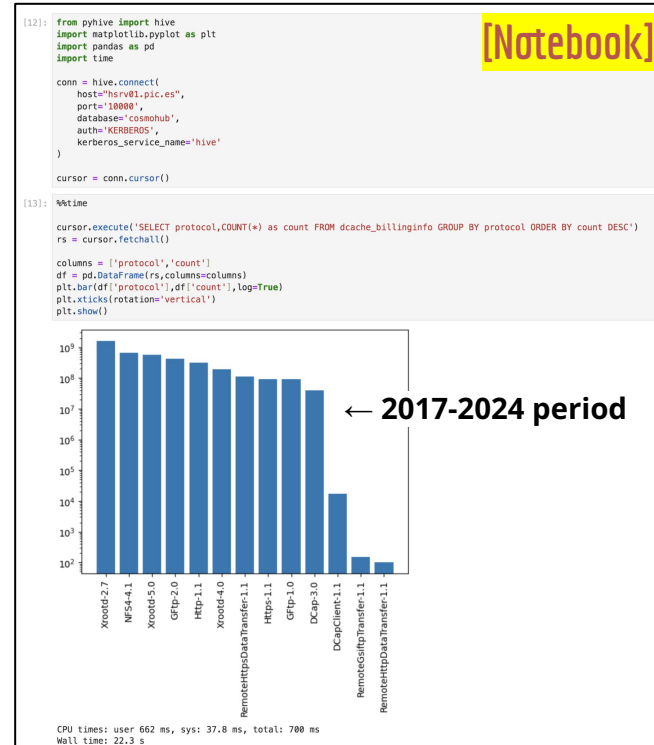
- 1st gen: obsolete nodes
- 2nd gen: 4 dual-twin
- 3rd gen: 12 nodes DIY
- **4th gen: 20+12 nodes**



Hive at PIC: door to Big Data Analytics

PIC historical **dCache billingDB** and **ESnet High-Touch for Data Challenge '24** inserted
Allows quick data analysis and exploration through **Jupyter Notebooks**

Table	NumRows	rawDataSize (TB)	totalSize (GB)
dcache_billinginfo	4.09E+09	5.05	138.32
dcache_doorinfo	5.11E+09	5.84	147.52
dcache_hitinfo	3.07E+09	1.75	49.88
dcache_storageinfo	6.64E+07	0.05	2.37
<i>dCache total</i>	<i>1.23E+10</i>	<i>12.69</i>	<i>338.08</i>
esnetHT_dc24	2.30E+09	6.60	125.53
All (287 tables)	3.27E+11	158.36	99194.37



Storage resources towards HL-LHC

Currently expanding and adapting the World-Wide LHC Computing Grid (WLCG) to accommodate **increased data processing demands expected at the HL-LHC era**

- Emphasis on the need for **cost-effective solutions** to manage the growing volume of data
- **Possible consolidation of storage** resources in fewer sites
- Focus on deploying lightweight storage systems (**data caches**) supporting traditional (Grid) and opportunistic (Cloud/HPC) compute resources
- **Boost task execution performance** by implementing efficient data caching mechanisms in close proximity to end users [\[CHEP talk\]](#)
- But **how are we using today the 'traditional' storage resources?**

A case study for the CMS experiment

CMS changed his Data Management system from **PhEDEx** to **Rucio** around 2020

Historical dCache data has been analyzed at PIC: **CMS access patterns, lifespans** and **other relevant file statistics**, in the period from 2017 to 2024, and comparing metrics pre/post PhEDEx

Previous data accesses studies at PIC storage system were performed on a replica dCache *billingDB* PostgreSQL, but the analysis of the data was slow and only one year worth of data was analysed at the time [link to CHEP]

This recent study has used the **Hive platform** over $1.25 \cdot 10^{10}$ rows and ~ 13 TB rawDataSize - it was carried out by **M. Santamaría** (last year Physics degree student)

A case study for the CMS experiment

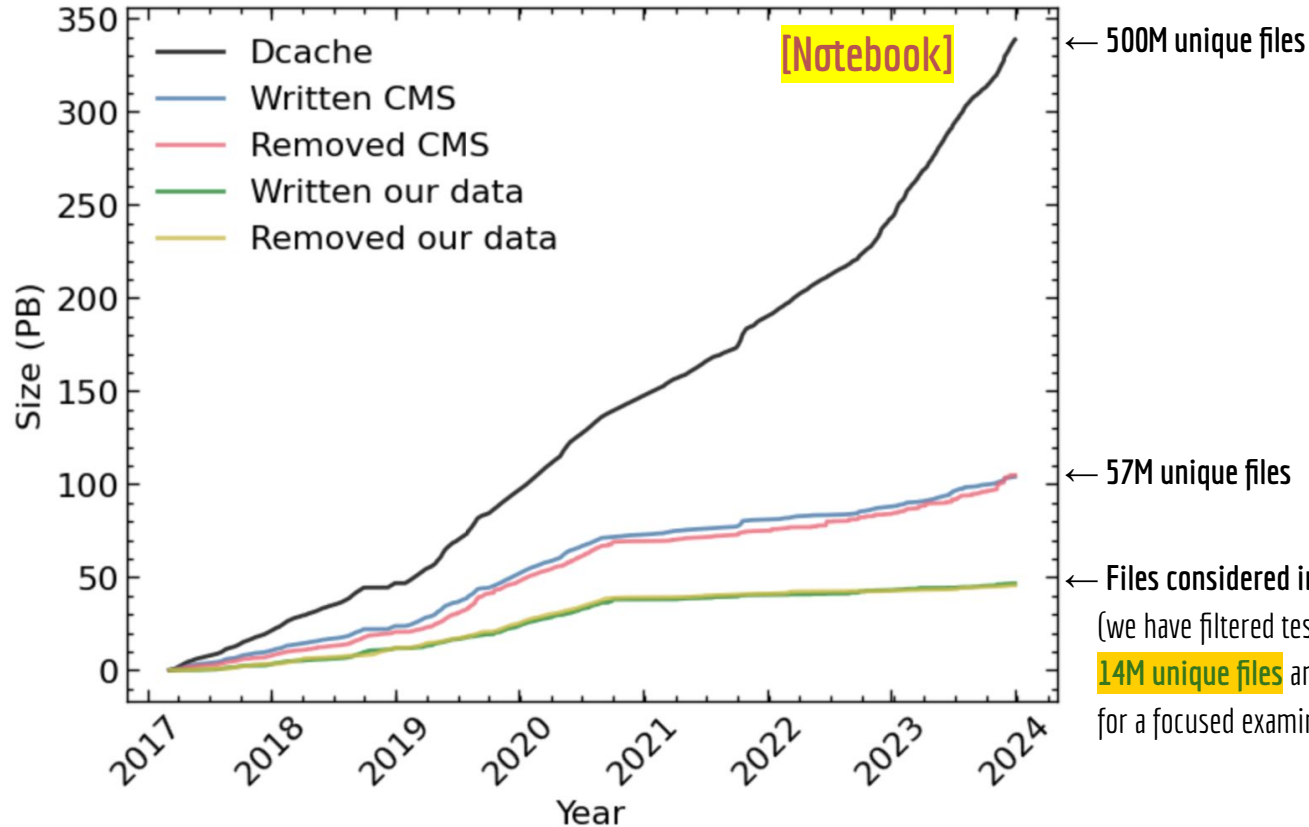
PIC currently deploys **~20 PB of disk storage**, which is managed by **dCache**

The LHC data that is stored in PIC typically **saturates the system every 3 months**, so the files in the storage system are **very dynamic**

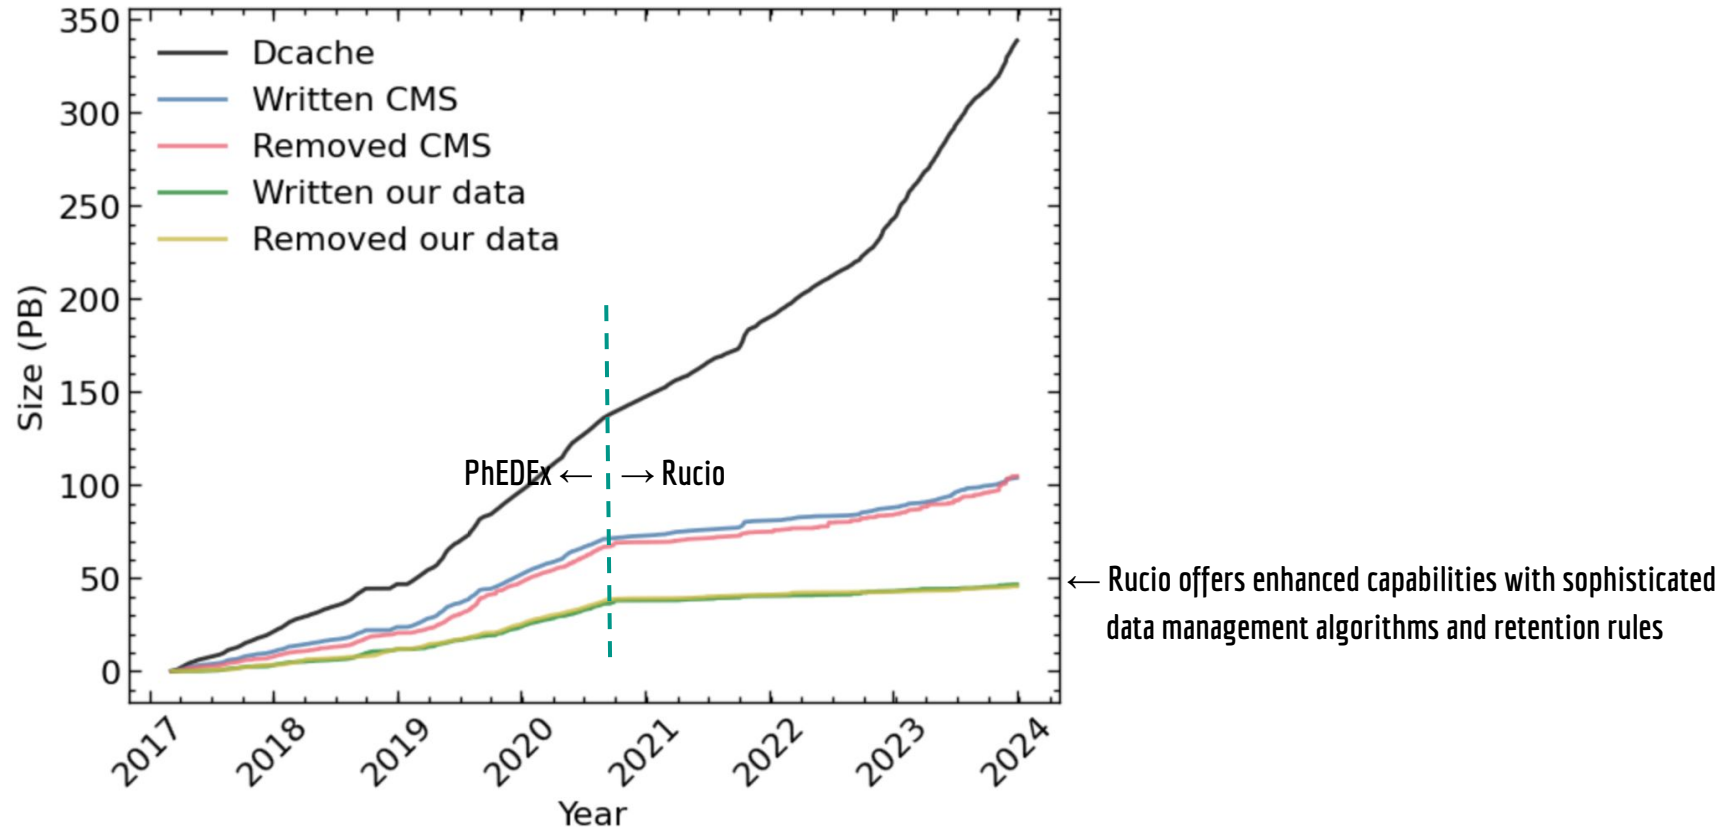
This is also seen in the number of **file deletions that occur often** to allow space for new files to be stored in the PIC storage system

The CMS experiment has written in PIC storage **~100 PB since 2017 (57M files)**


A case study for the CMS experiment




A case study for the CMS experiment



A case study for the CMS experiment



The total records that have been used in this study span over 1 billion database records, and Big Data analysis techniques have been applied. It is remarkable that the deployed methodology has provided results ('a graph') in less than 2', and it is one of the first studies ever made at PIC applying these techniques on dCache monitoring data



← Rucio offers enhanced capabilities with sophisticated data management algorithms and retention rules

Creating an unified table for analytics

The primary information is stored in two distinct PostgreSQL tables within the dCache billingDB: ***dcache_doorinfo*** and ***dcache_billinginfo***

To avoid costly SQL joins, **we merged these tables into an unified table**, which involved intricate SQL operations to trace and keep all data accesses of files that have been managed by dCache: creation - accesses - deletion (if any) - re-creations...

This process **demanded significant effort** and **careful attention** to ensure that all relevant information was correctly aligned and integrated into a single, cohesive table suitable for analysis within the Hadoop framework

This process has been **documented**, so it can be replicated considering new data or different periods of time (and even works in a PostgreSQL context)

Creating an unified table for analytics

Example illustrating the **lifecycle of a file** as represented in the *fulltable*:

fulltable.dclient	fulltable.bclient	fulltable.access_count	fulltable.reaccess_count	fulltable.action	fulltable.creation_date	fulltable.data	fulltable.data_desv
fulltable.protocol	fulltable.fullsize	fulltable.read_data	fulltable.corrected_path	fulltable.pnfsid	fulltable.id_access		
[REDACTED]	[REDACTED]	0	1	request	1579113829	1.57911382928E9	0.0
Gftp-2.0	3881827254	3881827254	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				000019
109AB82124BFFB0307C05D13FFB70	1579109561868000						
[REDACTED]	[REDACTED]	1	1	request	1579113829	1.5809505081229715E9	1287.5535606057604
Xrootd-2.7	3881827254	1328316269	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				000019
109AB82124BFFB0307C05D13FFB70	AAWd3Fewybg:						
[REDACTED]	[REDACTED]	2	1	request	1579113829	1.5809509886487913E9	877.2298492033507
Xrootd-2.7	3881827254	1317845693	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				000019
109AB82124BFFB0307C05D13FFB70	AAWd3KpQ31A:						
[REDACTED]	[REDACTED]	3	1	request	1579113829	1.5809515834984841E9	1159.1698642535305
Xrootd-2.7	3881827254	1340721592	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				000019
109AB82124BFFB0307C05D13FFB70	AAWd3K6XEJg:						
[REDACTED]	[REDACTED]	4	1	remove	1579113829	1.582741042951E9	0.0
NULL	3881827254	NULL	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				000019
109AB82124BFFB0307C05D13FFB70	1582737442951-2617388						

Creating an unified table for analytics

Example illustrating the **lifecycle of a file** as represented in the *fulltable*:

fulltable.dclient	fulltable.bclient	fulltable.access_count	fulltable.reaccess_count	fulltable.action	fulltable.creation_date	fulltable.data	fulltable.data_desv
fulltable.protocol	fulltable.fullsize	fulltable.read_data	fulltable.corrected_path		fulltable.pnfsid	fulltable.id_access	
[REDACTED]	[REDACTED]	0	1	request	1579113829	1.57911382928E9	0.0
Gftp-2.0	3881827254	3881827254	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	1579109561868000						
[REDACTED]	[REDACTED]	1	1	request	1579113829	1.5809505081229715E9	1287.5535606057604
Xrootd-2.7	3881827254	1328316269	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3Fewybg:						
[REDACTED]	[REDACTED]	2	1	request	1579113829	1.5809509886487913E9	877.2298492033507
Xrootd-2.7	3881827254	1317845693	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3KpQ31A:						
[REDACTED]	[REDACTED]	3	1	request	1579113829	1.5809515834984841E9	1159.1698642535305
Xrootd-2.7	3881827254	1340721592	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3K6XEJg:						
[REDACTED]	[REDACTED]	4	1	remove	1579113829	1.582741042951E9	0.0
NULL	3881827254	NULL	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	1582737442951-2617388						

File creation
[FTS transfer]

Creating an unified table for analytics

Example illustrating the **lifecycle of a file** as represented in the *fulltable*:

fulltable.dclient	fulltable.bclient	fulltable.access_count	fulltable.reaccess_count	fulltable.action	fulltable.creation_date	fulltable.data	fulltable.data_desv
fulltable.protocol	fulltable.fullsize	fulltable.read_data	fulltable.corrected_path		fulltable.pnfsid	fulltable.id_access	
[REDACTED]	[REDACTED]	0	1	request	1579113829	1.57911382928E9	0.0
Gftp-2.0	3881827254	3881827254	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	1579109561868000						
[REDACTED]	[REDACTED]	1	1	request	1579113829	1.5809505081229715E9	1287.5535606057604
Xrootd-2.7	3881827254	1328316269	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3Ffewybg:						
[REDACTED]	[REDACTED]	2	1	request	1579113829	1.5809509886487913E9	877.2298492033507
Xrootd-2.7	3881827254	1317845693	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3KpQ31A:						
[REDACTED]	[REDACTED]	3	1	request	1579113829	1.5809515834984841E9	1159.1698642535305
Xrootd-2.7	3881827254	1340721592	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3K6XEJg:						
[REDACTED]	[REDACTED]	4	1	remove	1579113829	1.582741042951E9	0.0
NULL	3881827254	NULL	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	1582737442951-2617388						

Multiple XRootD accesses
(resolving IPs we know if they were
local/remote accesses)

Creating an unified table for analytics

Example illustrating the **lifecycle of a file** as represented in the *fulltable*:

fulltable.dclient	fulltable.bclient	fulltable.access_count	fulltable.reaccess_count	fulltable.action	fulltable.creation_date	fulltable.data	fulltable.data_desv
fulltable.protocol	fulltable.fullsize	fulltable.read_data	fulltable.corrected_path	fulltable.pnfsid	fulltable.id_access		
[REDACTED]	[REDACTED]	0	1	request	1579113829	1.57911382928E9	0.0
Gftp-2.0	3881827254	3881827254	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				000019
109AB82124BFFB0307C05D13FFB70	1579109561868000						
[REDACTED]	[REDACTED]	1	1	request	1579113829	1.5809505081229715E9	1287.5535606057604
Xrootd-2.7	3881827254	1328316269	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				000019
109AB82124BFFB0307C05D13FFB70	AAWd3Fewybg:						
[REDACTED]	[REDACTED]	2	1	request	1579113829	1.5809509886487913E9	877.2298492033507
Xrootd-2.7	3881827254	1317845693	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				000019
109AB82124BFFB0307C05D13FFB70	AAWd3KpQ31A:						
[REDACTED]	[REDACTED]	3	1	request	1579113829	1.5809515834984841E9	1159.1698642535305
Xrootd-2.7	3881827254	1340721592	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				000019
109AB82124BFFB0307C05D13FFB70	AAWd3K6XEJg:						
[REDACTED]	[REDACTED]	4	1	remove	1579113829	1.582741042951E9	0.0
NULL	3881827254	NULL	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				000019
109AB82124BFFB0307C05D13FFB70	1582737442951-2617388						

File Removal
[if any]

Creating an unified table for analytics

Example illustrating the **lifecycle of a file** as represented in the *fulltable*:

fulltable.dclient	fulltable.bclient	fulltable.access_count	fulltable.reaccess_count	fulltable.action	fulltable.creation_date	fulltable.data	fulltable.data_desv
fulltable.protocol	fulltable.fullsize	fulltable.read_data	fulltable.corrected_path		fulltable.pnfsid	fulltable.id_access	
[redacted]	[redacted]	0	1	request	1579113829	1.57911382928E9	0.0
Gftp-2.0	3881827254	3881827254	/pnfs/picc.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	1579109561868000						
[redacted]	[redacted]	1	1	request	1579113829	1.5809505081229715E9	1287.5535606057604
Xrootd-2.7	3881827254	1328316269	/pnfs/picc.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3Fewybg:						
[redacted]	[redacted]	2	1	request	1579113829	1.5809509886487913E9	877.2298492033507
Xrootd-2.7	3881827254	1317845693	/pnfs/picc.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3KpQ31A:						
[redacted]	[redacted]	3	1	request	1579113829	1.5809515834984841E9	1159.1698642535305
Xrootd-2.7	3881827254	1340721592	/pnfs/picc.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3K6XEJg:						
[redacted]	[redacted]	4	1	remove	1579113829	1.582741042951E9	0.0
NULL	3881827254	NULL	/pnfs/picc.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/0000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	1582737442951-2617388						

Access count
[# actions on a file
File popularity]

Re-access count
[if files have been deleted
and placed again in dCache]

← Re-creations are traced with
different PNFSids, so we trace
identical files by filenames

Creating an unified table for analytics

Example illustrating the **lifecycle of a file** as represented in the *fulltable*:

fulltable.dclient	fulltable.bclient	fulltable.access_count	fulltable.reaccess_count	fulltable.action	fulltable.creation_date	fulltable.data	fulltable.data_desv
fulltable.protocol	fulltable.fullsize	fulltable.read_data	fulltable.corrected_path		fulltable.pnfsid	fulltable.id_access	
[REDACTED]	[REDACTED]	0	1	request	1579113829	1.57911382928E9	0.0
Gftp-2.0	3881827254	3881827254	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/00000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	1579109561868000						
[REDACTED]	[REDACTED]	1	1	request	1579113829	1.5809505081229715E9	1287.5535606057604
Xrootd-2.7	3881827254	1328316269	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/00000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3Fewybg:						
[REDACTED]	[REDACTED]	2	1	request	1579113829	1.5809509886487913E9	877.2298492033507
Xrootd-2.7	3881827254	1317845693	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/00000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3KpQ31A:						
[REDACTED]	[REDACTED]	3	1	request	1579113829	1.5809515834984841E9	1159.1698642535305
Xrootd-2.7	3881827254	1340721592	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/00000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	AAWd3K6XEJg:						
[REDACTED]	[REDACTED]	4	1	remove	1579113829	1.582741042951E9	0.0
NULL	3881827254	NULL	/pnfs/pic.es/data/cms/disk/store/data/Run2018B/Tau/RAW/v1/000/317/392/00000/3ADC1ECE-D767-E811-BEC2-02163E01A136.root				0000019
109AB82124BFFB0307C05D13FFB70	1582737442951-2617388						

This table consolidates and distills all the essential information from *dcache_billinginfo* and *dcache_doorinfo*. It takes a few hours to be created (for the full CMS records from 2017 to 2024), creating a few intermediate and auxiliary tables. The results eases all of the posterior analysis

Creating an unified table for analytics


```
CREATE TABLE jflix.OneTable STORED AS ORC AS (  
SELECT *,  
ROW_NUMBER() OVER (PARTITION BY corrected_path, pnfsid ORDER BY data) - 1 AS access_count,  
DENSE_RANK() OVER (PARTITION BY corrected_path ORDER BY creation_date) - 1 AS reaccess_count  
FROM (  
SELECT  
a.client AS dclient,  
a.action,  
a.data,  
CASE  
WHEN a.data_desv!=a.data_desv THEN 0.0  
ELSE a.data_desv  
END AS data_desv,  
a.pnfsid,  
a.corrected_path,  
MAX(b.fullsize) AS fullsize,  
id_access,  
MIN(UNIX_TIMESTAMP(b.datestamp)) AS creation_date  
FROM  
jflix.AccDoor AS a  
JOIN  
dcache_billinginfo AS b ON a.pnfsid = b.pnfsid  
WHERE  
b.isnew = TRUE  
AND b.datestamp > '2016-05-18'  
AND b.errorcode = 0  
GROUP BY  
a.client, a.action, a.data, a.data_desv, a.pnfsid, a.corrected_path, id_access  
) AS s  
);
```

```
CREATE TABLE jflix.fulltable_prev STORED AS ORC AS  
SELECT  
o.*,  
b.bclient,  
b.protocol,  
b.transfersize AS read_data  
FROM  
jflix.OneTable AS o  
LEFT JOIN  
jflix.accbill AS b ON (o.pnfsid = b.pnfsid AND o.id_access = b.id_access);
```

```
CREATE TABLE jflix.fulltable stored as orc AS  
SELECT  
f.dclient,f.bclient,  
CASE WHEN f.pnfsid IN (  
SELECT DISTINCT f.pnfsid  
FROM (  
SELECT id_access, pnfsid  
FROM fulltable_prev  
WHERE access_count = 0  
) AS f  
JOIN (  
SELECT initiator, pnfsid  
FROM cosmohub.dcache_billinginfo  
WHERE isnew = true  
) AS b  
ON f.pnfsid = b.pnfsid  
AND REGEXP_EXTRACT(b.initiator, 'door:([A-Za-z@]{1,})-?:{.*}[0-9-]{1,}', 2) != f.id_access  
)  
THEN access_count + 1  
ELSE access_count  
END AS access_count,  
f.reaccess_count,f.action,f.creation_date,f.data,f.data_desv,f.protocol,f.fullsize,f.read_data,f.corrected_path,f.pnfsid,f.id_access  
FROM fulltable_prev as f;
```

Data popularity


Data



1. **Data:** All files containing information about real collision events.

- (a) RAW: Only includes RAW data.
- (b) RECO: Includes ALCARECO, RAW-RECO, and RECO.
- (c) AOD: Includes AOD, MINIAOD, and NANO AOD.

MonteCarlo



2. **MonteCarlo:** Pertains to data generated from simulated collisions.

- (a) RAW: Includes GEN-SIM-RAW and GEN-SIM-DIGI-RAW.
- (b) RECO: Includes GEN-SIM-RECO, ALCARECO, and GEN-SIM-RAW-RECO.
- (c) AOD: Includes NANO AODSIM, AODSIM, and MINIAODSIM.

Data popularity

It is noteworthy that a **significant number of files have no recorded accesses**



	DATA			MC		
	RAW	RECO	AOD	RAW	RECO	AOD
Before Rucio	0.35	0.44	0.52	0.37	0.33	0.52
With Rucio	0.20	0.56	0.39	0.61	0.54	0.44

Comparison of percentage of non-accessed files with respect to total for each type of file before and after Rucio

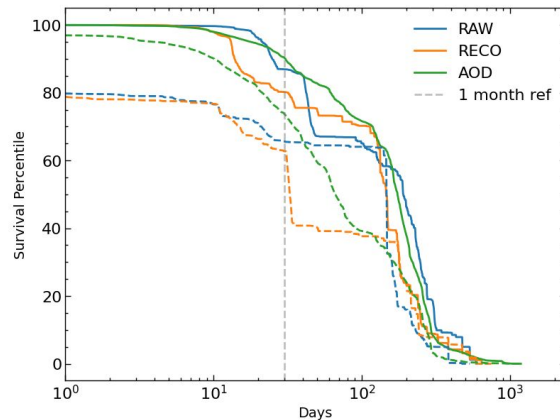

How long these files lived in the storage? Of course, recent files have possibility of not being accessed, but we observe many files that were never accessed, so the file lifetime is a relevant parameter to assess the real impact for these type of files

Lifetimes for accessed/non-accessed files

Files that were accessed
tended to remain in
storage for a longer
period →



Files persist longer
in storage with
Rucio compared to
the PHEEx era →

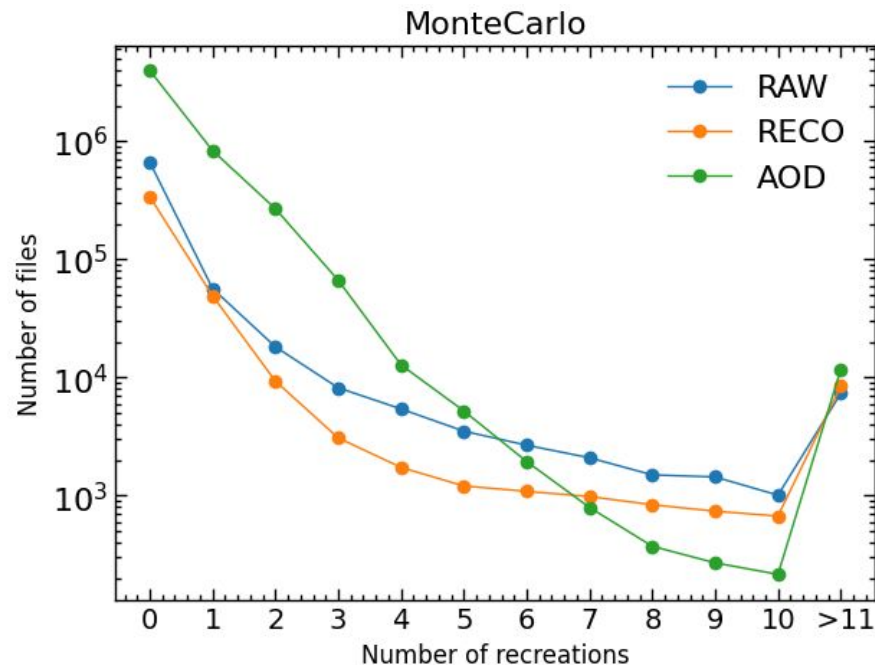
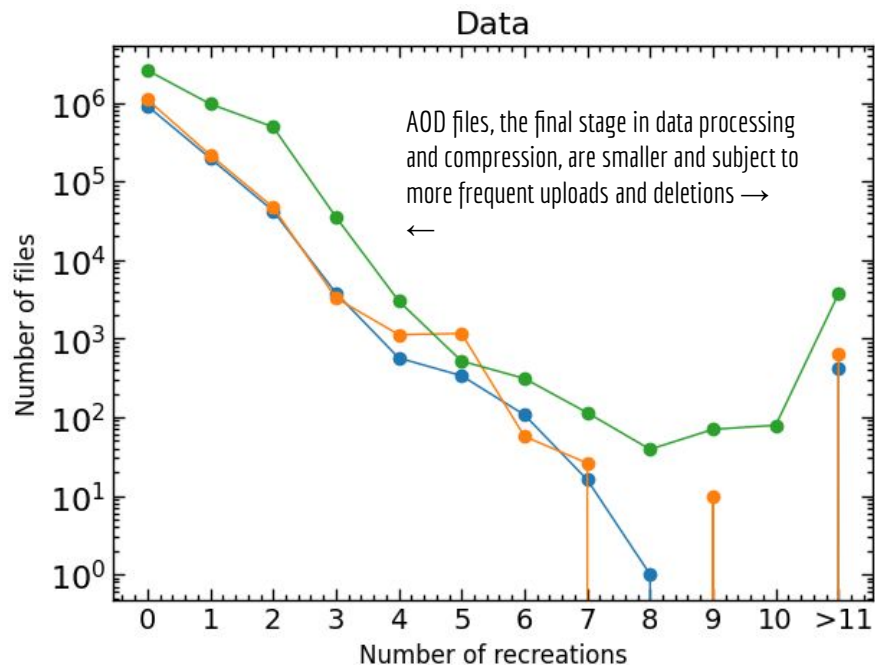


Lifetimes for accessed/non-accessed files

```
# Loop for Data with access_count > 1
i = 0
for key, value in Data.items():
    sql = """
    SELECT u.diff as bin, count(u.diff) as freq
    FROM (
        SELECT CAST((data-creation_date)/100 as int) as diff
        from fulltable
        where action='remove'
        and access_count>1
        and creation_date<1593561600
        and corrected_path rlike 'pnfs/pic.es/data/cms/disk/store/\(data\)/.\*?/\(.\*\)/'
    ) AS u
    where u.diff is not null
    GROUP BY u.diff
    ORDER BY u.diff
    """ .format(value)
    cursor.execute(sql)
    rs = cursor.fetchall()
    df = pd.DataFrame(rs)
    df.columns = ['bins', 'freq']
    accum = np.array(np.cumsum(df['freq']))
    neg_perc_accum = 100. - accum / accum[-1] * 100.
    ax[0][0].plot((100 * df['bins']) / 3600. / 24., neg_perc_accum, '-', color=line_colors[i])
    i += 1
```

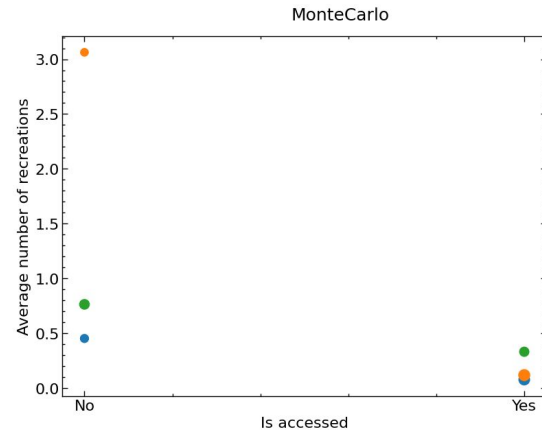

File re-creations

Most files are stored on disk only once and, once deleted, are not re-uploaded.
However, **some files may be created, deleted, and recreated multiple times**, each time with a new *pnfsid*


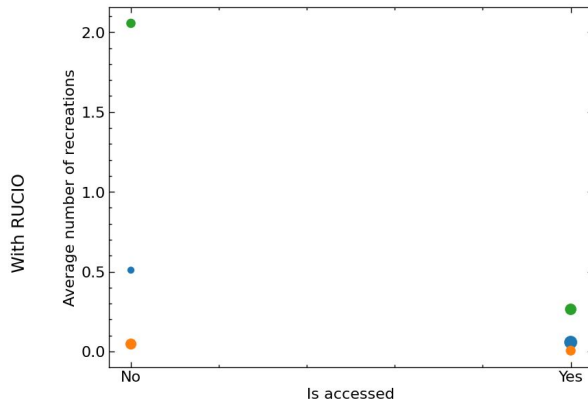


Re-creations for accessed/non-accessed files

Files that were more re-created
tended to be the ones with no
accesses at all →




Files that were more re-created
tended to be the ones with no
accesses at all →



Re-creations for accessed/non-accessed files

Re-creations weighted by file
lifetimes →

(non-accessed files had shorter
lifespans before Rucio)



Conclusions - next steps

Demonstrated **large-scale data analysis** on dCache metadata using Hive at PIC

Analyzed **long-term file characteristics for CMS across multiple years**

Observed significant impact of **non-accessed files**, aligning with findings from other centers (e.g., BNL)

Currently establishing **automated injection** of *billingDB* records into Hive, creating a streamlined process for data extraction and analysis to support future studies across multiple VOs, which include deployment of custom **grafana views**

New possibilities for integrating **additional data sources**, such as job submissions from PIC's **HTCondor** batch system, or **ESnet High-Touch** for DC'24 and other network mini-challenges...

Acknowledgements

The authors of this work express their gratitude to the PIC and CIEMAT teams for their support in these studies and for deploying novel cache services for the CMS experiment in the Spanish region. This project is partially financed by the Spanish Ministry of Science and Innovation (MINECO) through grants FPA2016-80994-C2-1-R, PID2019-110942RB-C22 and BES-2017-082665, which include FEDER funds from the European Union. It has also been supported by the Ministerio de Ciencia e Innovación MCIN AEI/10.13039/501100011033 under contract PID2020-113614RB-C21, the Catalan government under contract 2021 SGR 00574, and the Red Española de Supercomputación (RES) through the grant DATA-2020-1-0039.

CosmoHub has been developed by PIC (maintained by IFAE and CIEMAT) in collaboration with ICE-CSIC. It received funding from the Spanish government (grant EQC2021-007479-P funded by MCIN/AEI/10.13039/501100011033), the EU NextGeneration/PRTR (PRTR-C17.I1), and the Generalitat de Catalunya.



Thanks!
Questions?