


# Xtra Large Data Platform

AIDEL Osman

# Who I am ?

## Database service manager at CC-IN2P3

• Mariadb



• PostgreSQL



• MySQL



• Oracle

ORACLE®

• MongoDB



• ElasticSearch



• XLDP



# Plan



Context

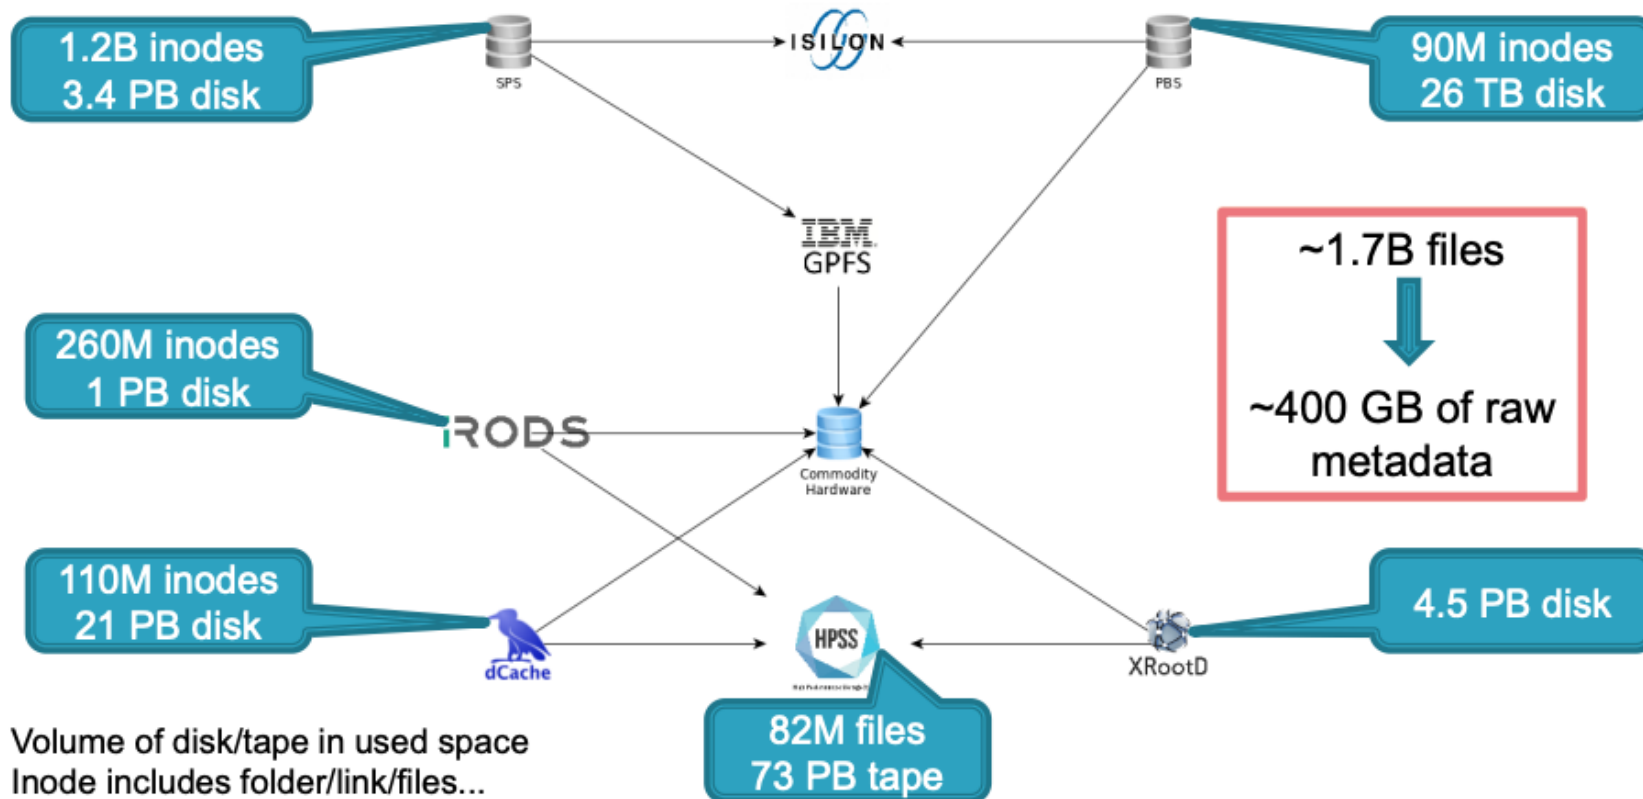
Objectives

XLDP overview

Prospectives

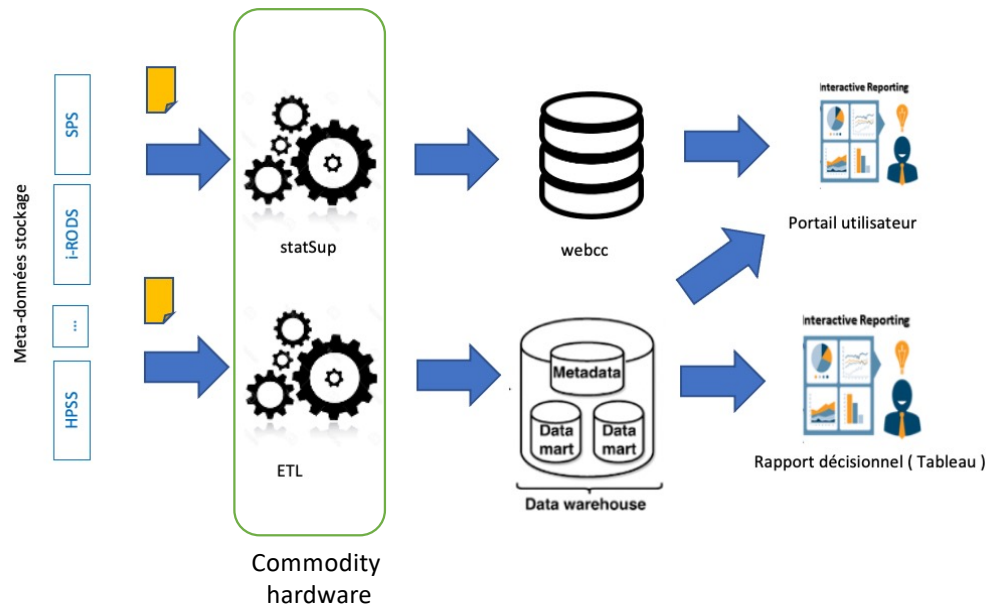


# Context



# Context

Storage systems provide metadata such as file size, owner, last access time, last modification time ...  
Those metadata are processed by some internal tools in order to build storage statistics.



Many views :

- Services
- Experiments
- Users

For many customers :

- CC-Users
- Czar : Experiment manager
- Service expert
- Team manager / Direction


# Context



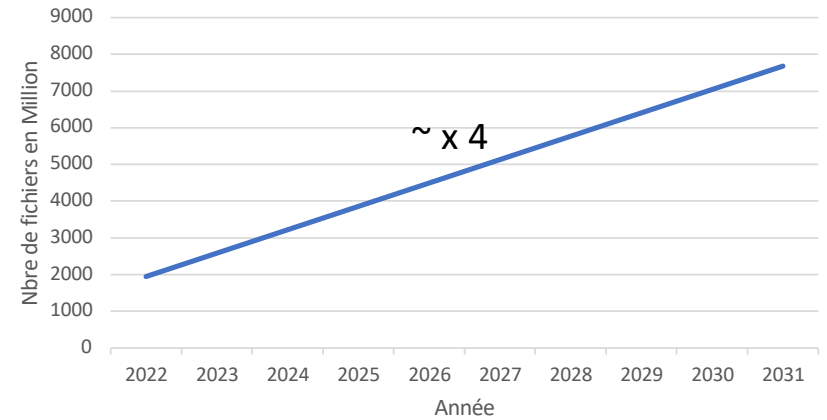
## Current metadata volume by storage service

	Compressed size in GB	Files number in million
Dcache	9	119
SPS	13	1569
PBS	1,5	150
HPSS	2,9	109
Xrootd	5,3	
<b>Total</b>	<b>33,2 (uncompressed ~200)</b>	<b>~ 1947</b>

## Estimated volumetry at 10 years



## Estimated number of files at 10 years



HOW TO PROCESS STORAGE STATISTICS IN 2030 ?

# Objectives

Reproducing storage statistics regardless of the metadata volume.

Collecting and centralizing heterogeneous data sources.

Standardizing and centralizing metadata processing :

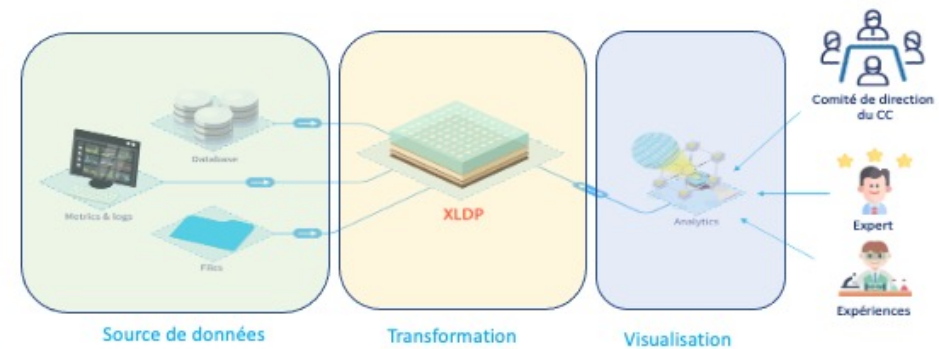
- Data cleaning
- Data consolidation with OpenIDM

Performing simple and complex processing.

Secure access to the platform: Compliance with [General Data Protection Regulation](#).

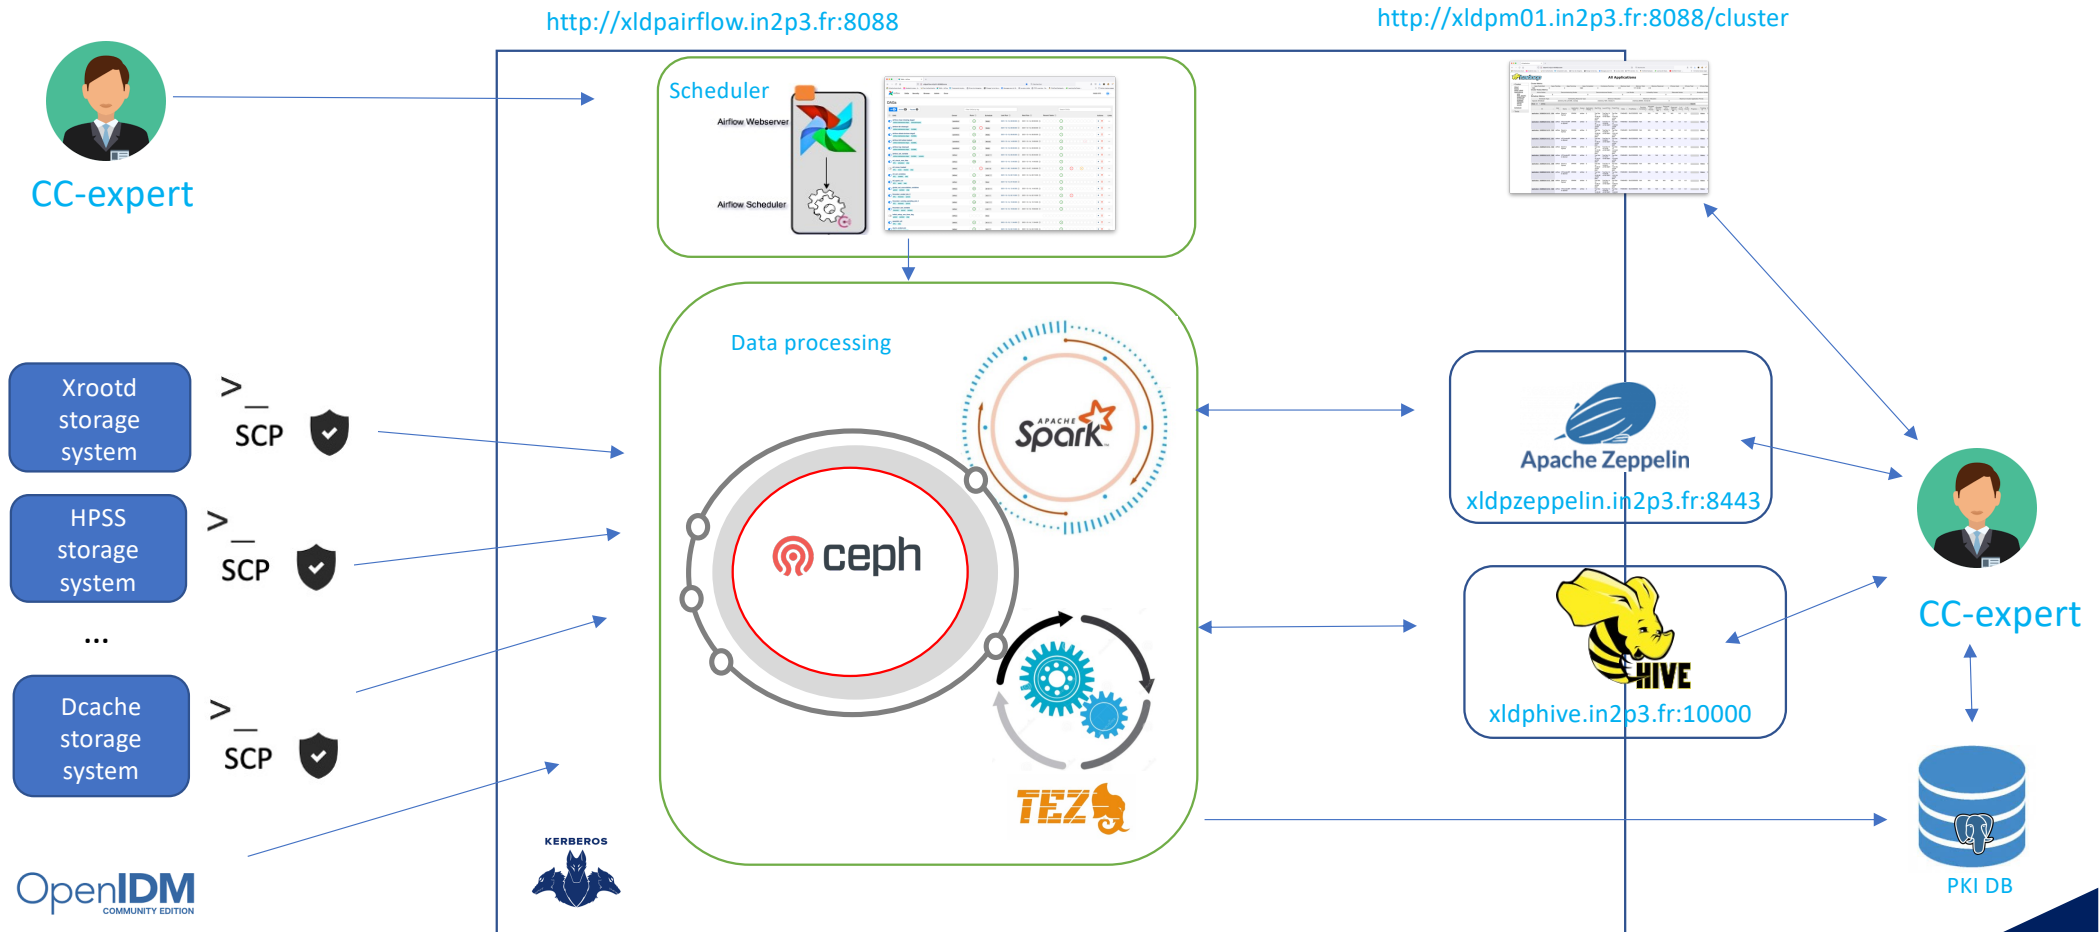
Facilitate the integration of new Key Performance Indicators.

Providing access to data sources if needed.



# XLDP

## XLDP





# Zeppelin

Zeppelin Notebook Job Search andubois

## Scala-SQL guide

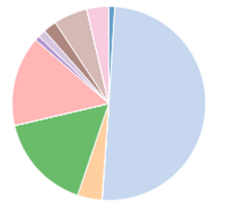
Données vis SPARK JOB FINISHED

### camembert

```
%spark.sql
SELECT og.groupname, sum(d.size) as
total_size
FROM dcache_file d
JOIN `group` og on d.gid= og.gid
GROUP BY og.groupname
```

settings

- lbno cta fgrilles pauger
- agata t2k root xenon
- juno lsst hess belle2
- dteam ilc



Took 24 sec. Last updated by andubois at April 13 2022, 3:52:42 PM. (outdated)

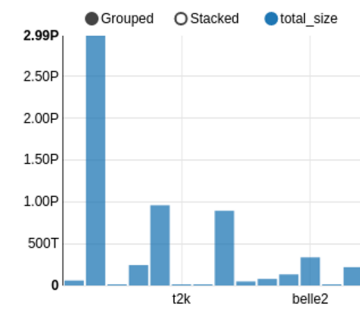
Données vis SPARK JOB FINISHED

### histogramme

```
%spark
z.show(
  dcache_df.join( group_df, dcache_df.col
    ("gid").equalTo( group_df.col("gid") )
  )
  .groupBy( group_df.col
    ("groupname") )
  .agg( sum( dcache_df.col( "size" )
    ).alias("total_size") )
)
```

settings

Grouped Stacked total\_size



Took 26 sec. Last updated by andubois at April 13 2022, 3:52:42 PM. (outdated)

Données bru SPARK JOB FINISHED

```
%spark
z.show( spark.sqlContext.sql("SELECT og
.groupname, sum(d.size) as total_size FROM
dcache_file d JOIN `group` og on d.gid= og
.gid GROUP BY og.groupname") )
```

settings

groupname	total_size
lbno	59452535154261
cta	2987849890157364
fgrilles	990042195485
pauger	243536059786404
agata	959442771111419
t2k	5320
root	22678
xenon	893573849317986

Took 21 sec. Last updated by andubois at April 13 2022, 3:52:49 PM. (outdated)

# Airflow

The screenshot displays the Airflow web interface. At the top, there is a navigation bar with the Airflow logo and menu items: DAGs, Security, Browse, Admin, and Docs. The current time is 16:46 (+01:00) and the user is identified as OA.

### DAGs

Summary: All 34, Active 27, Paused 7. Filter DAGs by tag: [ ] Search DAGs: [ ]


DAG	Owner	Runs	Schedule	Last Run	Next Run	Recent Tasks	Actions	Links
<a href="#">airflow-clear-missing-dags2</a> <small>airflow-maintenance-dags   teamclalrvoyant</small>	operations	44	@daily	2023-01-30, 01:00:00	2023-01-31, 01:00:00	1	[▶] [🗑]	...
<a href="#">airflow-db-cleanup2</a> <small>airflow-maintenance-dags</small>	operations	1	@daily	2023-01-30, 01:00:00	2023-01-31, 01:00:00	13	[▶] [🗑]	...
<a href="#">airflow-delete-broken-dags2</a> <small>airflow-maintenance-dags   GLOBAL</small>	operations	44	@daily	2023-01-30, 01:00:00	2023-01-31, 01:00:00	1	[▶] [🗑]	...

2023-01-31T12:45:16+0 [Runs: 25] Run: etl-pbs\_home-2023-01-30 [Layout: Left > Right] [Update] Find Task...

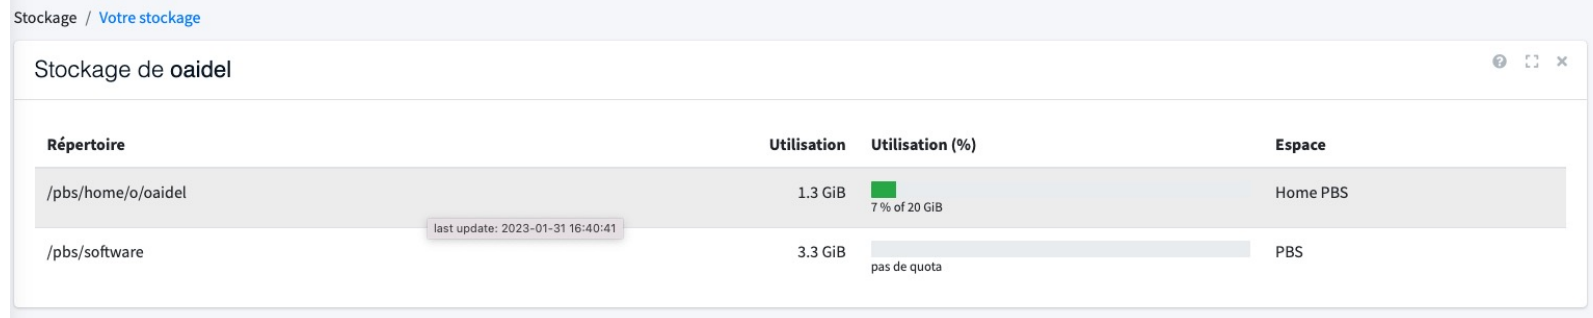
Operator tags: BetterSparkOperator, PythonOperator, ShortCircuitOperator. Status legend: queued, running, success, failed, up\_for\_retry, up\_for\_reschedule, upstream\_failed, skipped, scheduled, deferred, no\_status.

Auto-refresh: [ ]

```
graph LR; A[check_configuration] --> B[load_configuration]; B --> C[check_file_not_used]; C --> D[set_spark_resources_task]; D --> E[start_etl]; E --> F[validate_parquet_file]; F --> G[remove_old_csv]; F --> H[update_hive_external_table]; G --> I[move_input_file]; H --> I;
```



# KPI : File extensions



Home

Jobs

Stockage

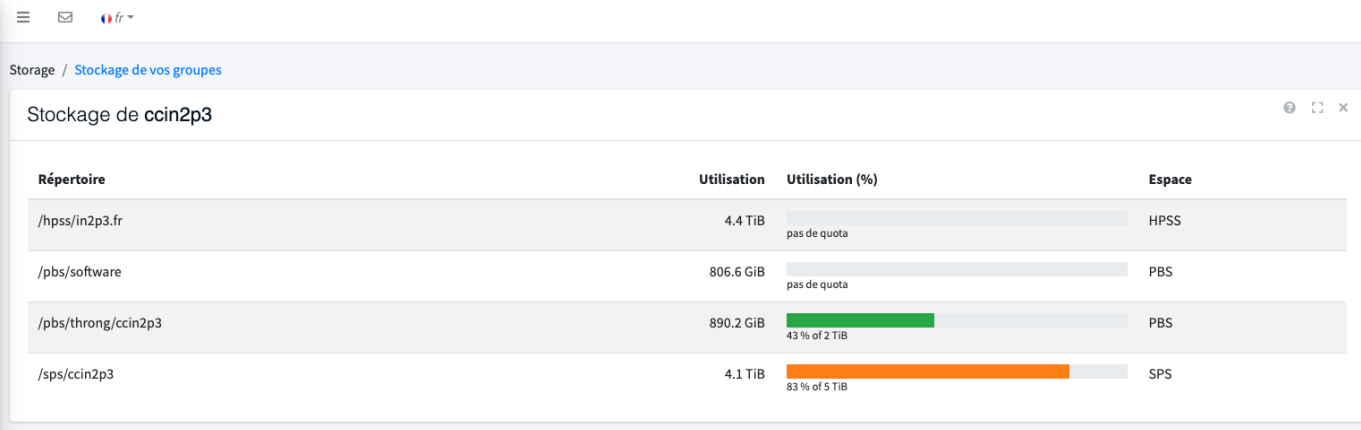
Votre stockage

Stockage de vos groupes

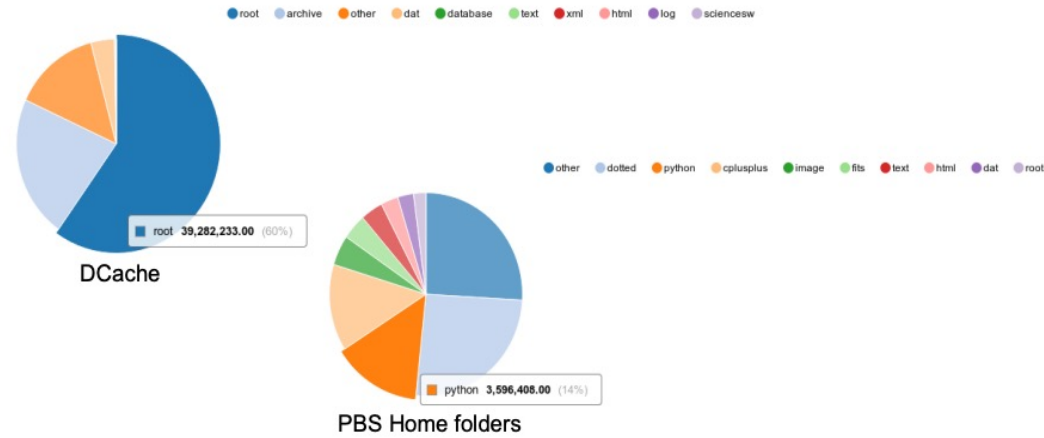
Rapports

Ressources

Tickets





# KPI example



File extensions repartition in # of file

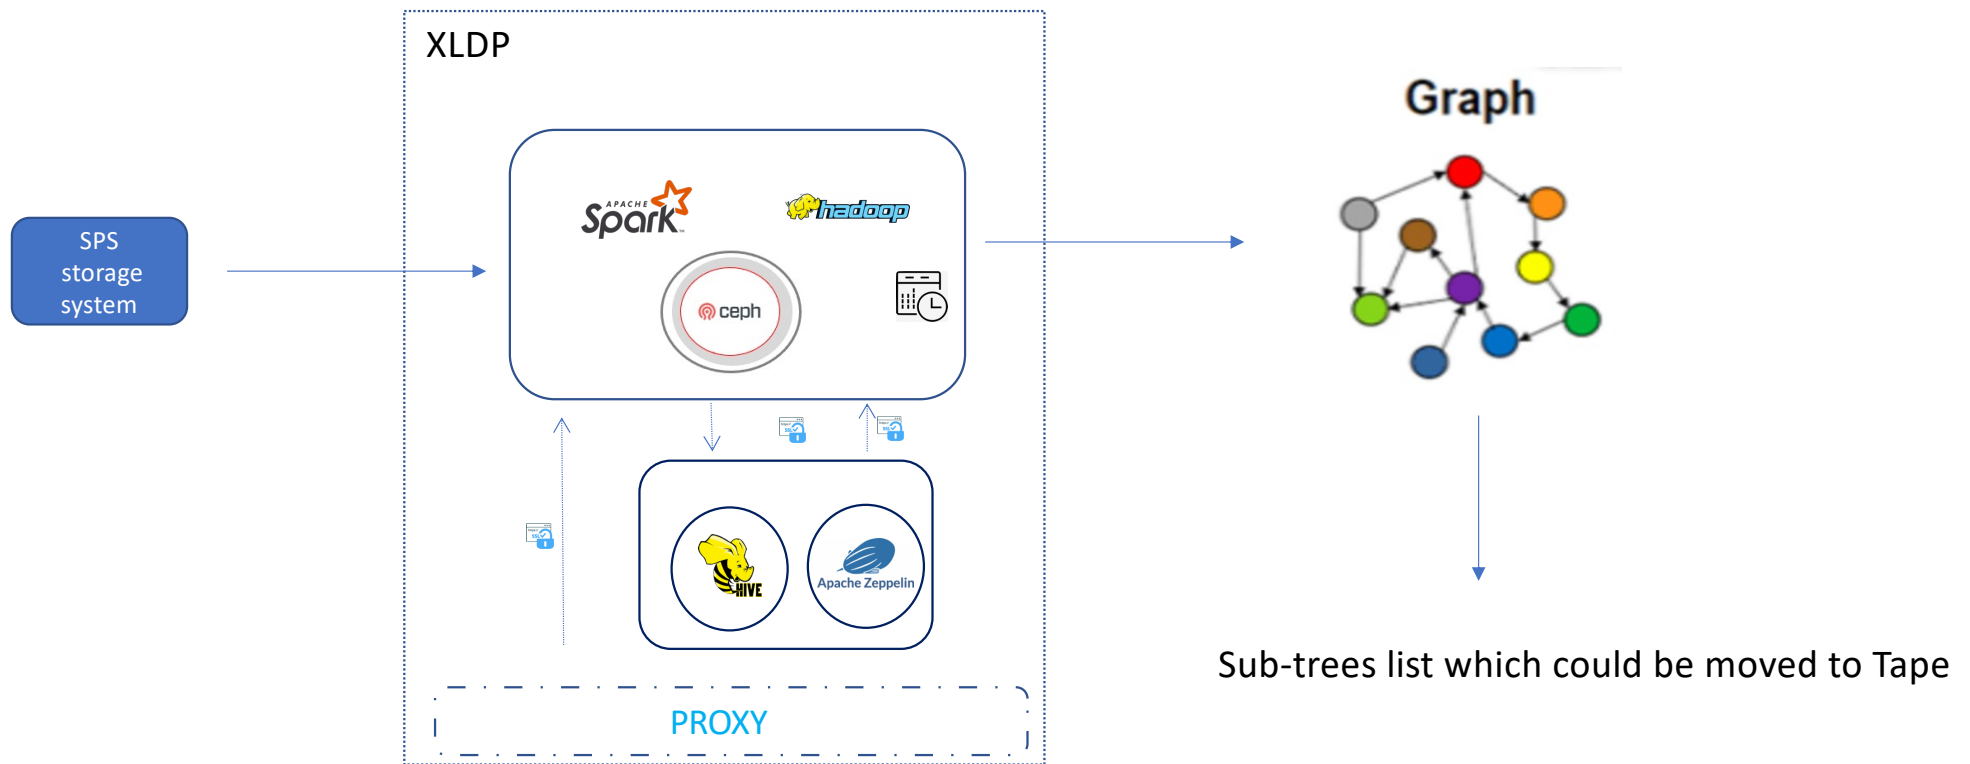
## File age per access time



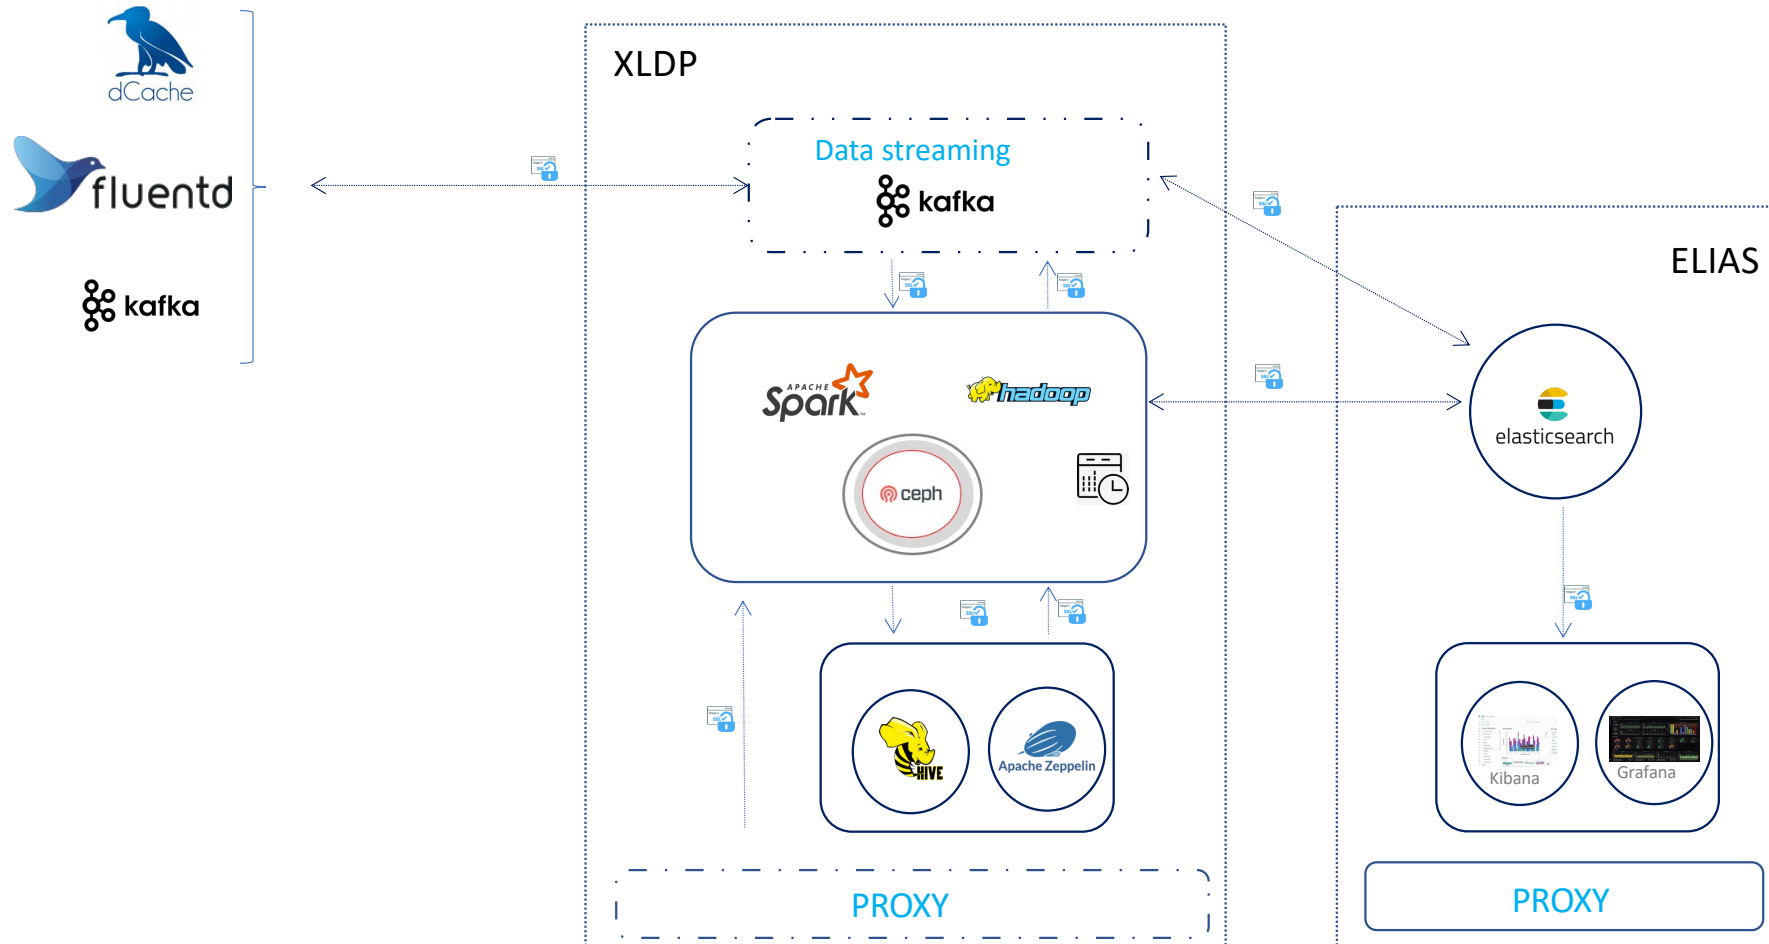
# Prospectives

Tapes are cheaper than disks.

To save money, we try to identify dead sub-trees no more accessed on disks for a long time.



# Prospectives



# Questions

