PIC
port d'informació
científica

RUCIO

Rucio at Port d'Informació Científica (PIC)

7th Rucio Community Workshop, 1st October 2024

Francesc Torradeflot

Outline

- Introduction
- Deployment
  - Helm Chart
  - Client
- Projects:
  - MAGIC
  - ICFO
  - InCAEM
  - Others
- Conclusions

Introduction

- PIC stands for **Port d'Informació Científica**
- Founded in 2003, collaboration between IFAE and CIEMAT. Located near Barcelona in the UAB campus.
- **Tier-1 node of the WLCG** with the mission to transfer this knowledge and technologies to other activities
- Infrastructure
  - 2x100 Gbps to Academic Network
  - 100 PB in+out per year
  - Disk - dCache: 20 PB (+Ceph 3.5 PB raw)
  - Tape - Enstore: 63 PB
  - Computing - HTCondor: 12k cores

Introduction

- PIC has been traditionally involved in Physics experiments:
  - Particle Physics
  - Cosmology
  - Gamma-ray Astronomy
  - …
- We started exploring Rucio within the scope of ESCAPE in 2019
- Rucio has become part of our portfolio and we aim to use it in other research fields

Deployment

- Deploy Rucio with all necessary components for data transfer and orchestration.
- Ensure deployment is as simple and automated as possible
- We will eventually want to launch many instances (one per experiment)
- Integration of Rucio deployments into PIC’s monitoring / alerting systems
- Automatic detection and transfers of data from a remote site (RSE at ‘on-site’) produced from observation at and replication in off-site RSEs (PIC).

Deployment - Helm Chart

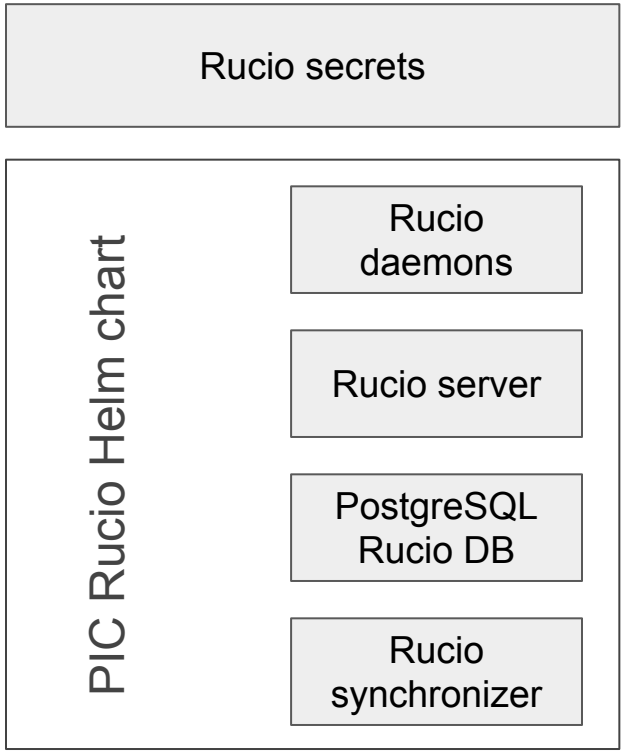

TLS certificate & key, CA certificates, X509 certificate & key

Fork of the official Rucio Helm Chart to:

- Add custom permissions, schema, lfn2pfn (?), etc
- Tweak monitoring / hermes
- …

Persistent PostgreSQL DB on k8s with initialized schema (currently without HA)

Set up monitoring, add RSE and account info, add proxy renewal CronJob, handle race conditions

Deployment - Client

- File discovery at origin RSE
  - External DB
  - CSV file
  - List directory content
- Get file info
  - name
  - checksum & size
  - metadata
- Register replica at origin RSE
- Create datasets / containers
- Add replication rule
- Remove from origin

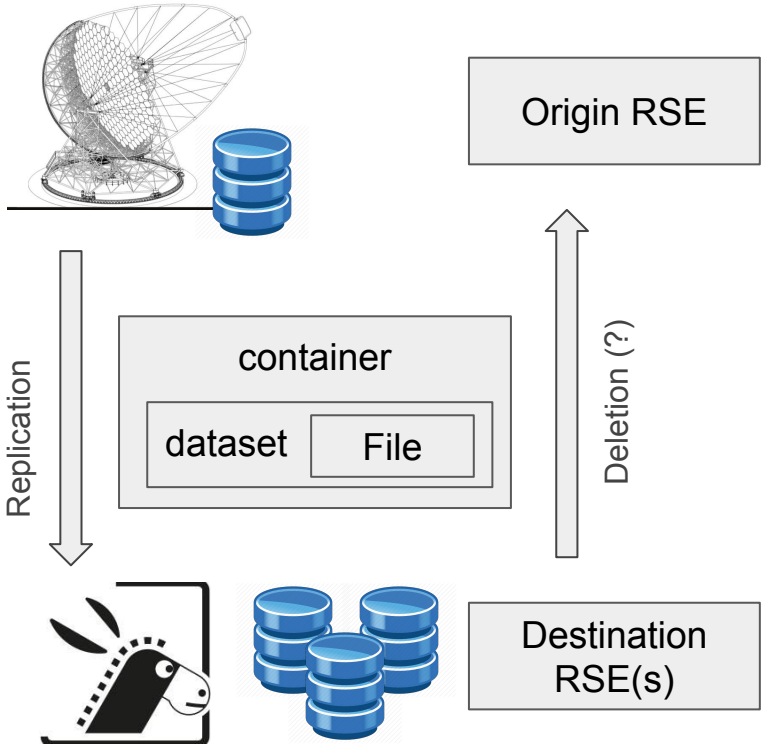

Deployment - Client

20241001_M1_05117132.005_I_LHAASOJ0056+-W0.60+090.root

Origin RSE pfn	gsiftp://grid.magic.iac.es:2811//data/Other/rucio_tmp/ESCAPE-DC22/transfers/20241001_M1_05117132.005_I_LHAASOJ0056+-W0.60+090.root
name	pic.es/data/magic/Data/Star/v1/LHAASOJ0056+/2024_10_01/20241001_M1_05117132.005_I_LHAASOJ0056+-W0.60+090.root
adler32	65520df3
bytes	4565276
metadata	night=2024_10_01, run_number=5117132, datatype=Star, etc

dataset	2024_10_01
container 1	5117132
container 2	2024_10_01
container 3	RAW_M1

Projects - MAGIC

Our first production use case (2022)

Goals:

- Replace scripting layer to automate data transfers from MAGIC observatory to PIC, deletion of data at source once replicated
- Continue feeding the legacy multivac DB.

Problems:

- Only using Rucio as a thin wrapper on top of FTS. The truth is in the Multivac DB
- File deletion strategy (with nested rules) is not working anymore
- Moving data to tape is done by hand

Future:

Minimum effort. Clean Rucio DB and make file deletion work again (how?)

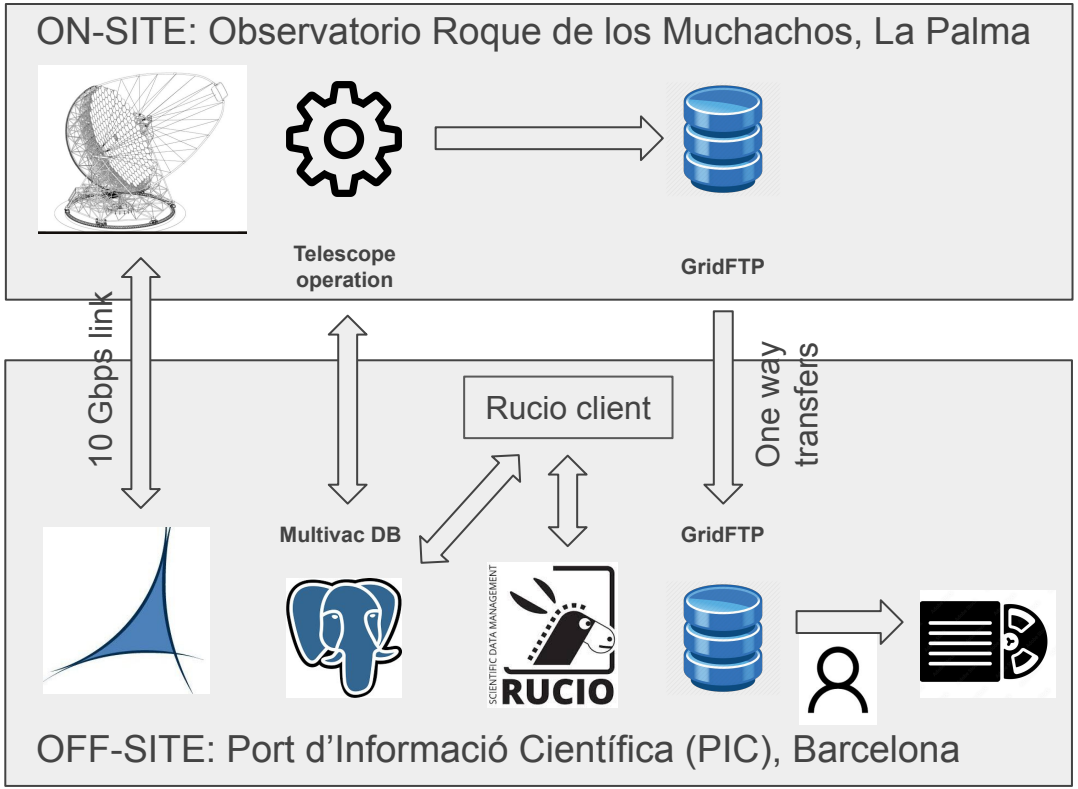

Projects - ICFO

In production since 2023

Goals:

- Automate data transfers from ICFO to PIC
- Automate tape archival
- Follow metadata standards

Problems:

- Poor network connection
- Very open, everything to-be-defined

Future:

- Define and implement metadata DB - Rucio integration (ideas?)

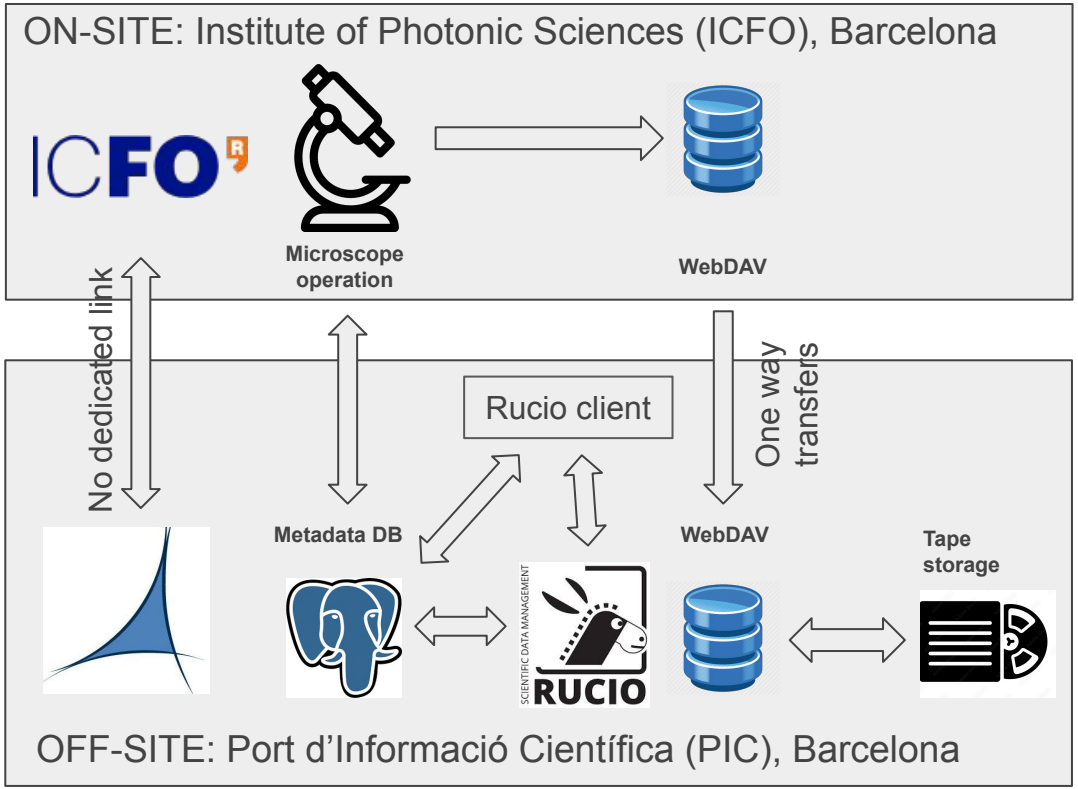

Projects - InCAEM

Exploratory phase

Goals:

- Data management for the InCAEM project
- Integration with standard Metadata catalog (ICAT?)
- Preserve AAI from ALBA
- Provide POSIX access to data

Problems:

- Potentially thousands of users/groups per year
- Rucio + FTS might be an overkill. Exploring alternatives: Globus, rsync
- Still very open

Future:

- Choose a tool to orchestrate data transfers

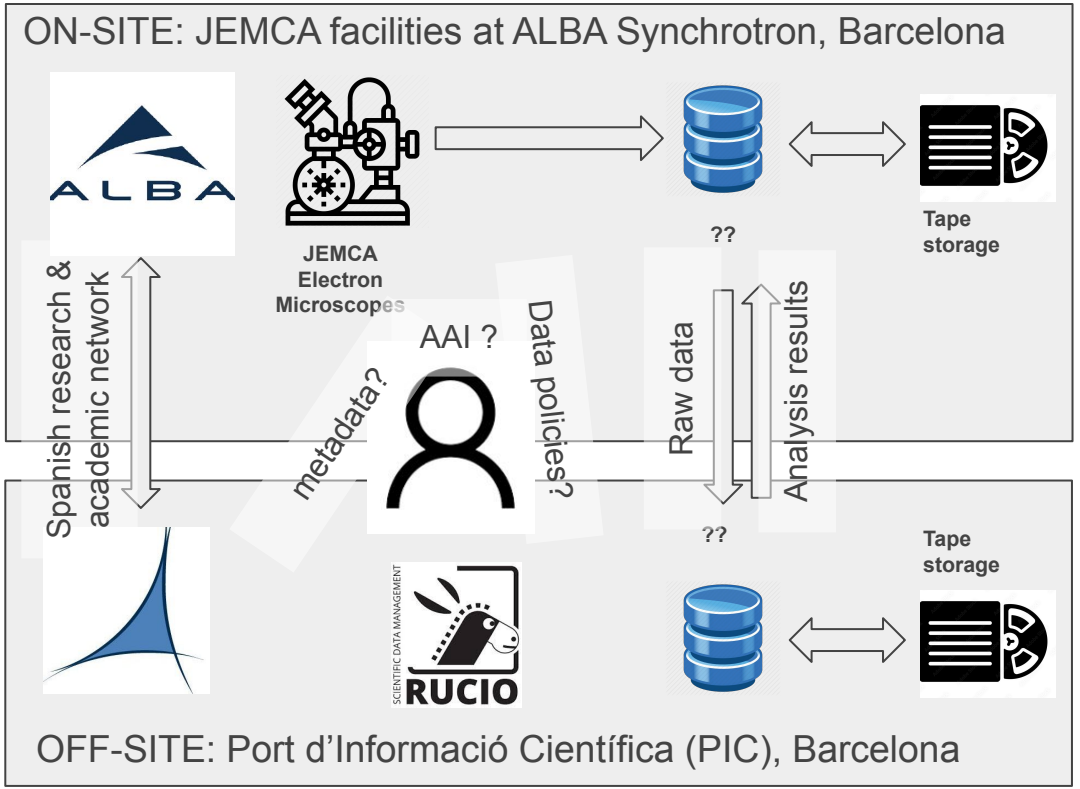

Projects - Others

We are considering adopting Rucio for the data management of other projects

- LST:
  - Orchestrate data transfers from La Palma to PIC and then to other LST datacenters
  - CTA has already chosen Rucio as a component of the DPPS
  - Currently setting up Rucio and RSEs
- Ciemat:
  - Data management between Ciemat's datacenters: Xula, Ceta, Ciemat-T2 and PIC
  - Easier access to data and computing resources for Ciemat scientists
  - Currently working on network setup, RSEs, etc
- Spanish Supercomputing Network (RES)

Conclusions

- Semi-automated deployment, 100% on in-house k8s, for SMEs
- Proved to be a valid tool to automate data transfers
- Steep learning curve. It might be an overkill if used only for data transfers
- TODO: use all of Rucio's capabilities
  - Define A&A policies
  - Metadata integration
  - Tape storage management
- More TODOs:
  - Set up new instances
  - Update versions of current instances
  - Explore new functionalities: tokens, webUI, etc
- A lot of work and limited manpower (we'll see …)

Thank you