RUCIO service for Gamma-ray astronomy projects

RUCIO Workshop, Lancaster University, November 2022

Jordi Delgado (jordidem@pic.es), A. Bruzzese, G. Merino on behalf of the PIC and CTA-ESCAPE teams

OUTLINE

- PORT D'INFORMACIÓ CIENTÍFICA (PIC)
- THE SCAPE DATA LAKE
- CTA-ESCAPE USE CASE (I)
- CTA-ESCAPE USE CASE (II)
- RUCIO DEPLOYMENT AT PIC
- LONG HAUL TRANSFER TESTS
- MAGIC DATA TRANSFER MIGRATION
- SUMMARY

PORT D'INFORMACIÓ CIENTÍFICA (PIC)

Funded in 2003, PIC is a scientific-technological centre specializing in data analysis and management technologies, operated through a collaboration between CIEMAT and IFAE.

Mission: Participate at highest level in LHC Computing (Tier-1) and leverage experience to support other scientific activities.

Value: experience from multidisciplinary collaborations to build scalable and robust common services for multiple experiments: leverage synergies, economies of scale, cross-fertilization.

- Multidisciplinary team to bridge the gap between the scientific needs and the infrastructure - advanced computing services.
- Service providers for data preservation, analysis and sharing.
- Close collaboration with disciplinary researchers to design & build architectures for handling scientific data flows.
    - Agile interaction. Focus on prototype/feedback loop.

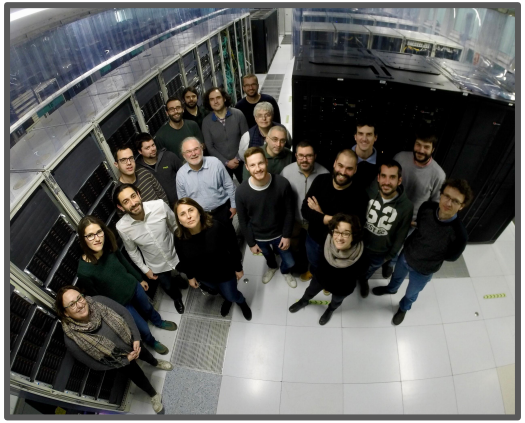

PORT D'INFORMACIÓ CIENTÍFICA (PIC)

200 Gbps connection to Research Networking

- Largest data mover in Spanish academic network: 70 PB in+out per year

Data processing services

- 10.000 CPUs, 18 GPUs (HTCondor cluster)
- 15 PB disk (dCache) + 40 PB tape (Enstore)
- Big data platform (Hadoop/HIVE/Spark)
  - 144 cores/1.5TB RAM/430TB HD
  - Cosmology analysis web portal and data processing cluster

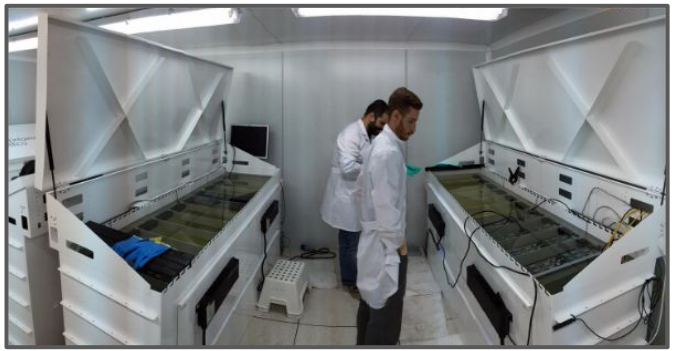

Two machine rooms, total ~150 kW_{IT}:

- ~120 kW in 150 m^2 air-cooled room (high efficiency, renovated 2014-2016)
- ~30 kW in 25 m^2 innovative liquid immersion cooling system (higher efficiency)

PORT D'INFORMACIÓ CIENTÍFICA (PIC)

MAGIC Data Center (Tier 0) from 2009

- Data analysis service to the collaboration: 227 accounts at PIC (168 active)
- Automated data transfers ORM-PIC using FTS: 200TB/yr of MAGIC data transferred
- 2.5 PB of data on tape
- Development of the MAGIC Data Management and Preservation Plan

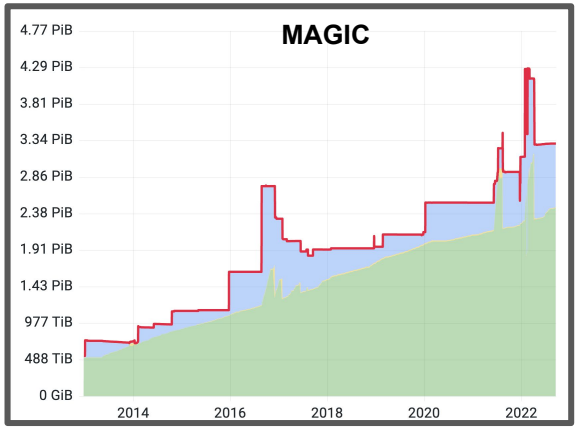

Leading CTA-LST1 off-site coordination tasks and ¼ CTA offsite DCs

Storage for off-site replication of the LST data from 2019

- Automated data transfers ORM-PIC using FTS, now using dedicated 10Gbps circuit: Mirror of the RAW-ZFITS, DL1 and DL2 data
- 2.2 PB of LST data on tape
- 500 TB of disk buffer for data access
- Data distribution via webdav/http

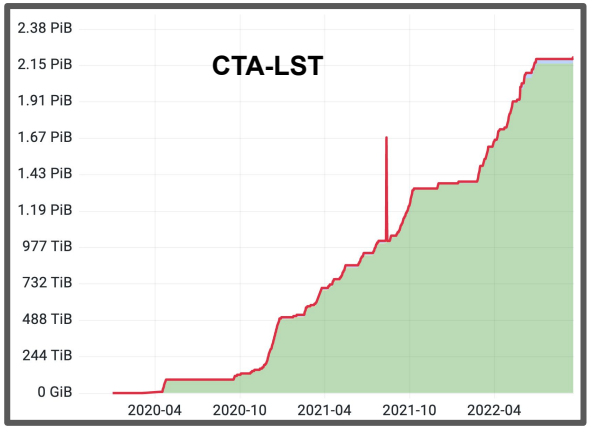

ESCAPE WP2: THE ESCAPE DATA LAKE

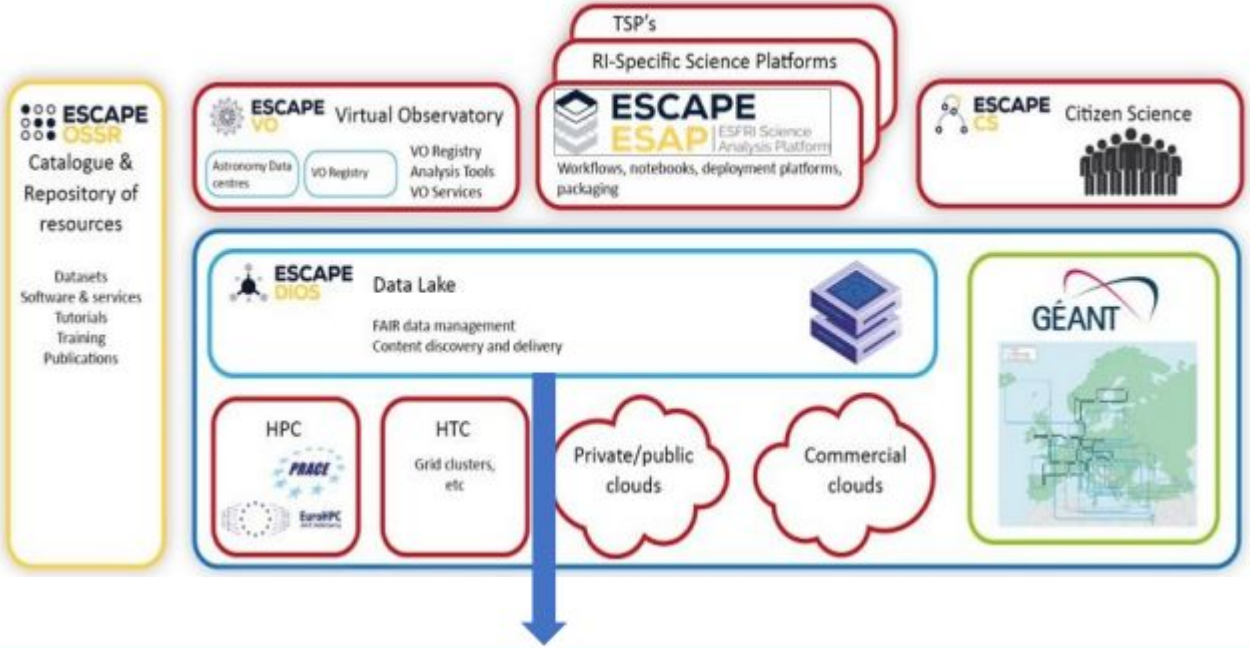

IFAE-PIC as one of the CTA partners in WP2 **ESCAPE Data Lake**

Already presented by Xavier!

Credits: X. Espinal

The ESCAPE Scientific Data Lake is a reliable, policy-driven, distributed data infrastructure. Capable of managing Exabyte-scale data sets, and able to deliver data on-demand at low latency to all types of processing facilities

CTA-ESCAPE WP2 USE CASE (I)

Long haul ingestion and replication

- Imaging Atmospheric Cherenkov Telescopes
- Study the most extreme environments in our Universe
- Ability to produce large volumes of data
- Observatories are by their nature often in remote locations Data transfer, storage and processing are key

MAGIC

CTAO

Credits: G. Hughes

CTA-ESCAPE WP2 USE CASE (I)

Long haul ingestion and replication

Description/Goal: Automatic detection and transfers of data from a remote site (RSE at ‘on-site’) produced from observation at ORM-la Palma, transfer and replication in off-site RSEs (PIC).

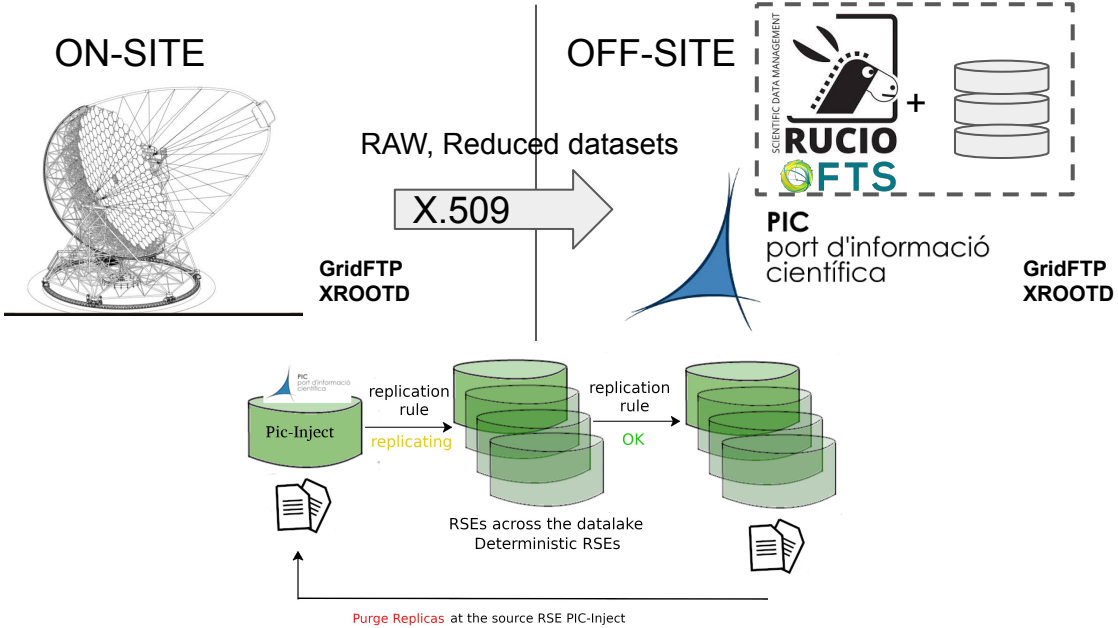

CTA-ESCAPE WP2 USE CASE (I)

Long haul ingestion and replication - **MAGIC Configuration**

Description/Goal: Automatic detection and transfers of data from a remote site (RSE at 'on-site') produced from observation at ORM-la Palma, transfer and replication in off-site RSEs (PIC).

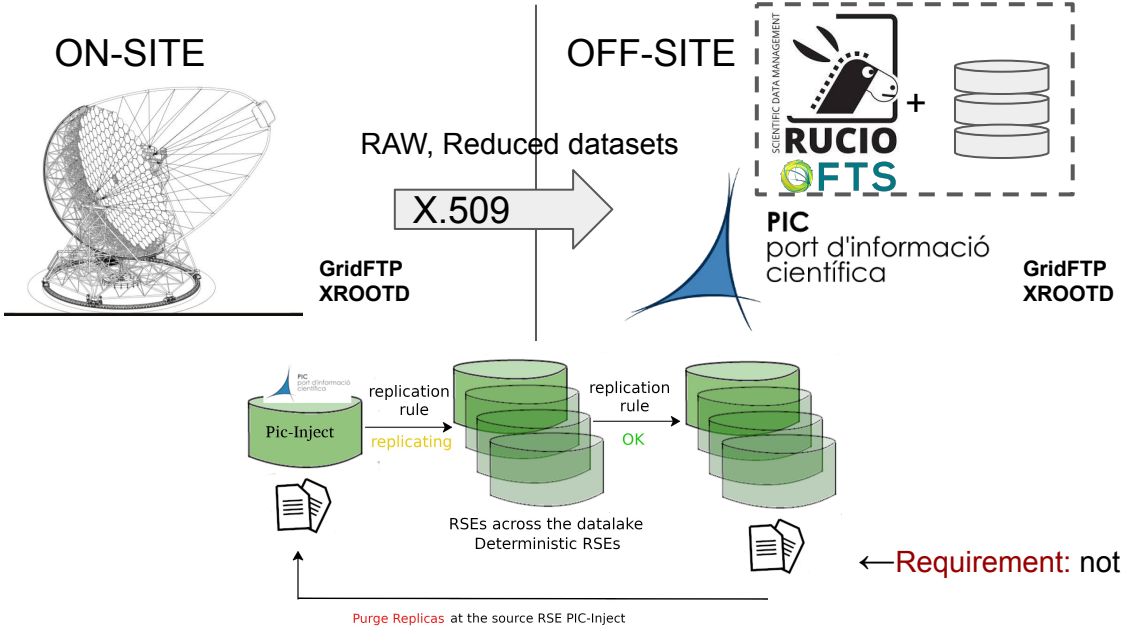

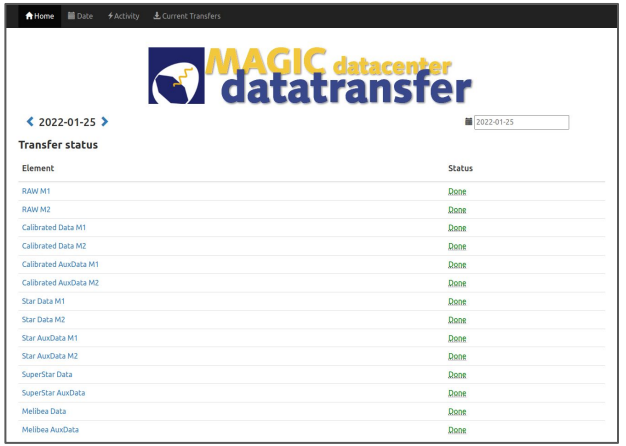

←Requirement: not deleting from the source

CTA-ESCAPE WP2 USE CASE (I)

Long haul ingestion and replication - **CTA Configuration**

Description/Goal: Automatic detection and transfers of data from a remote site (RSE at 'on-site') produced from observation at ORM-la Palma, transfer and replication in off-site RSEs (PIC) and after replication is successfully done, RUCIO triggers the deletion of the files from the origin (La Palma).

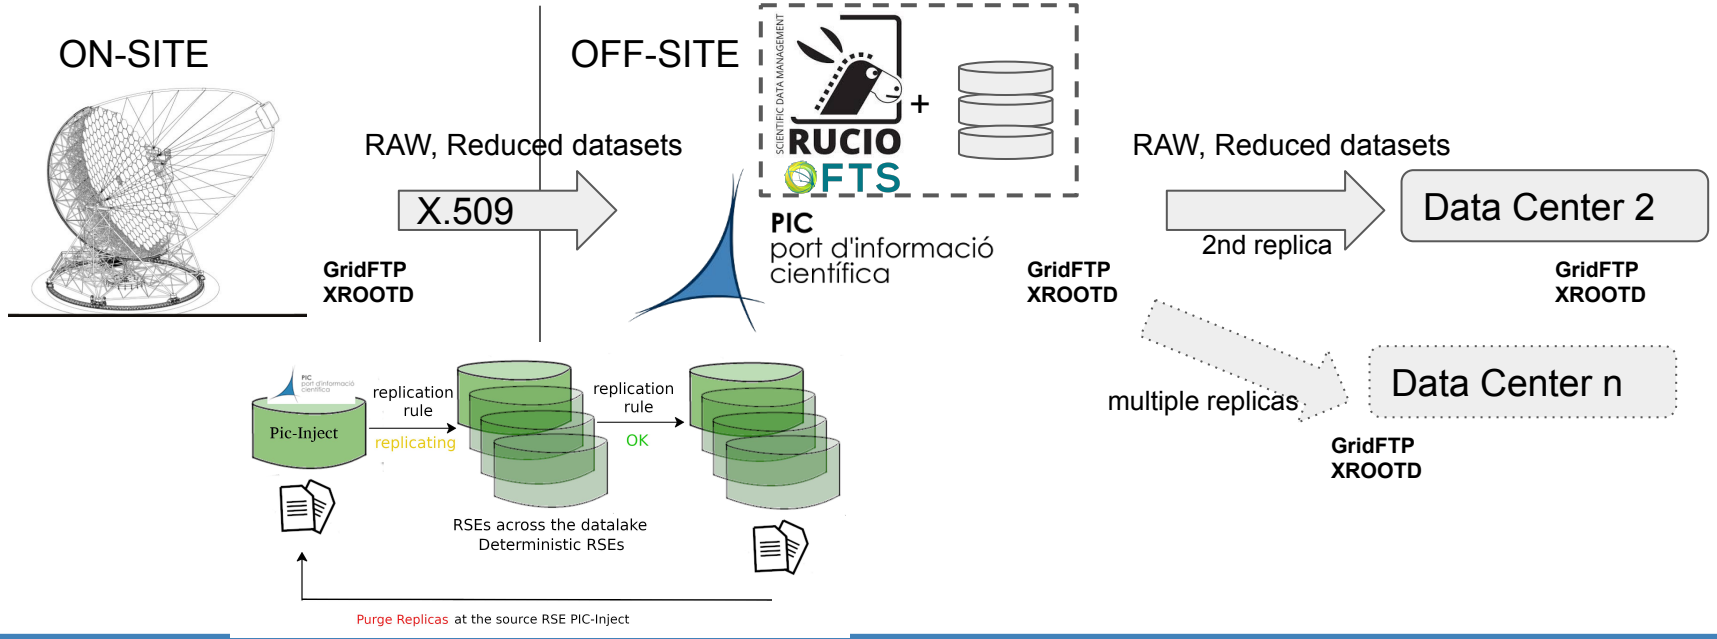

CTA-ESCAPE WP2 USE CASE (II)

DIRAC-RUCIO integration

Description/Goal: DIRAC is the workload management system for CTA. Users and automatic pipelines will be able to submit jobs and stage data through DIRAC. The tests focus on simple operations

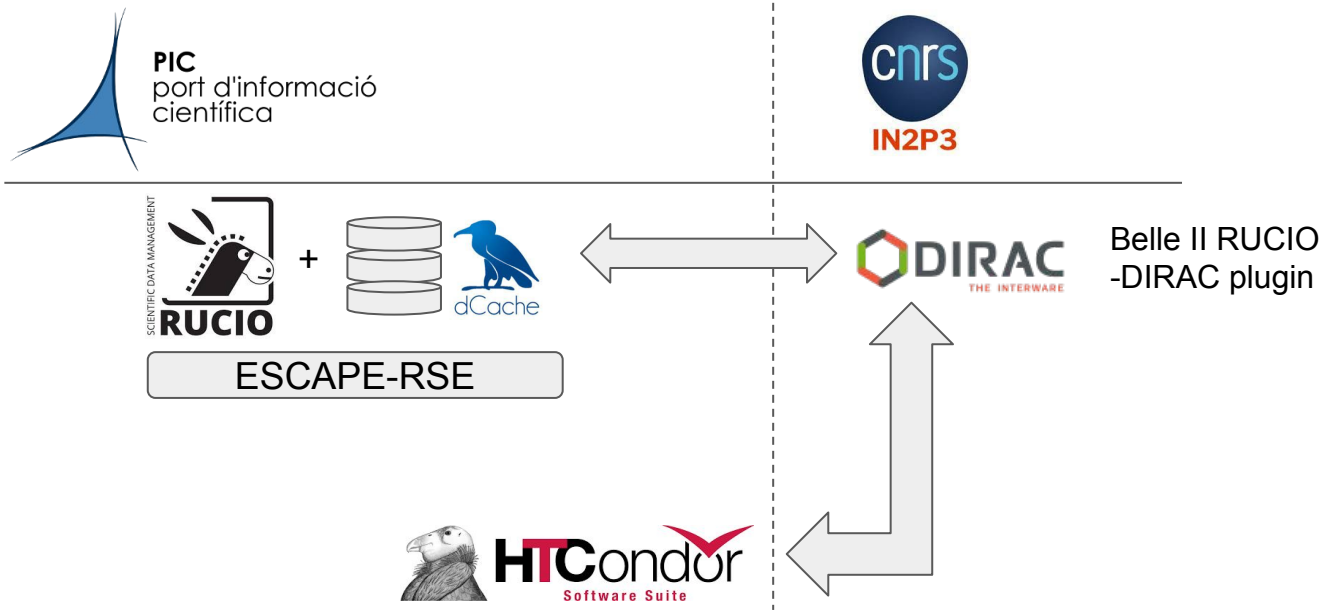

Operations to test

- ingest data using DIRAC
- launch test jobs on worker nodes
- launch CTA production scripts
- read data from tape

RUCIO DEPLOYMENT AT PIC

PIC is hosting three RUCIO (1.26) instances, 1 for MAGIC, 2 for CTA (+ESCAPE) activities:

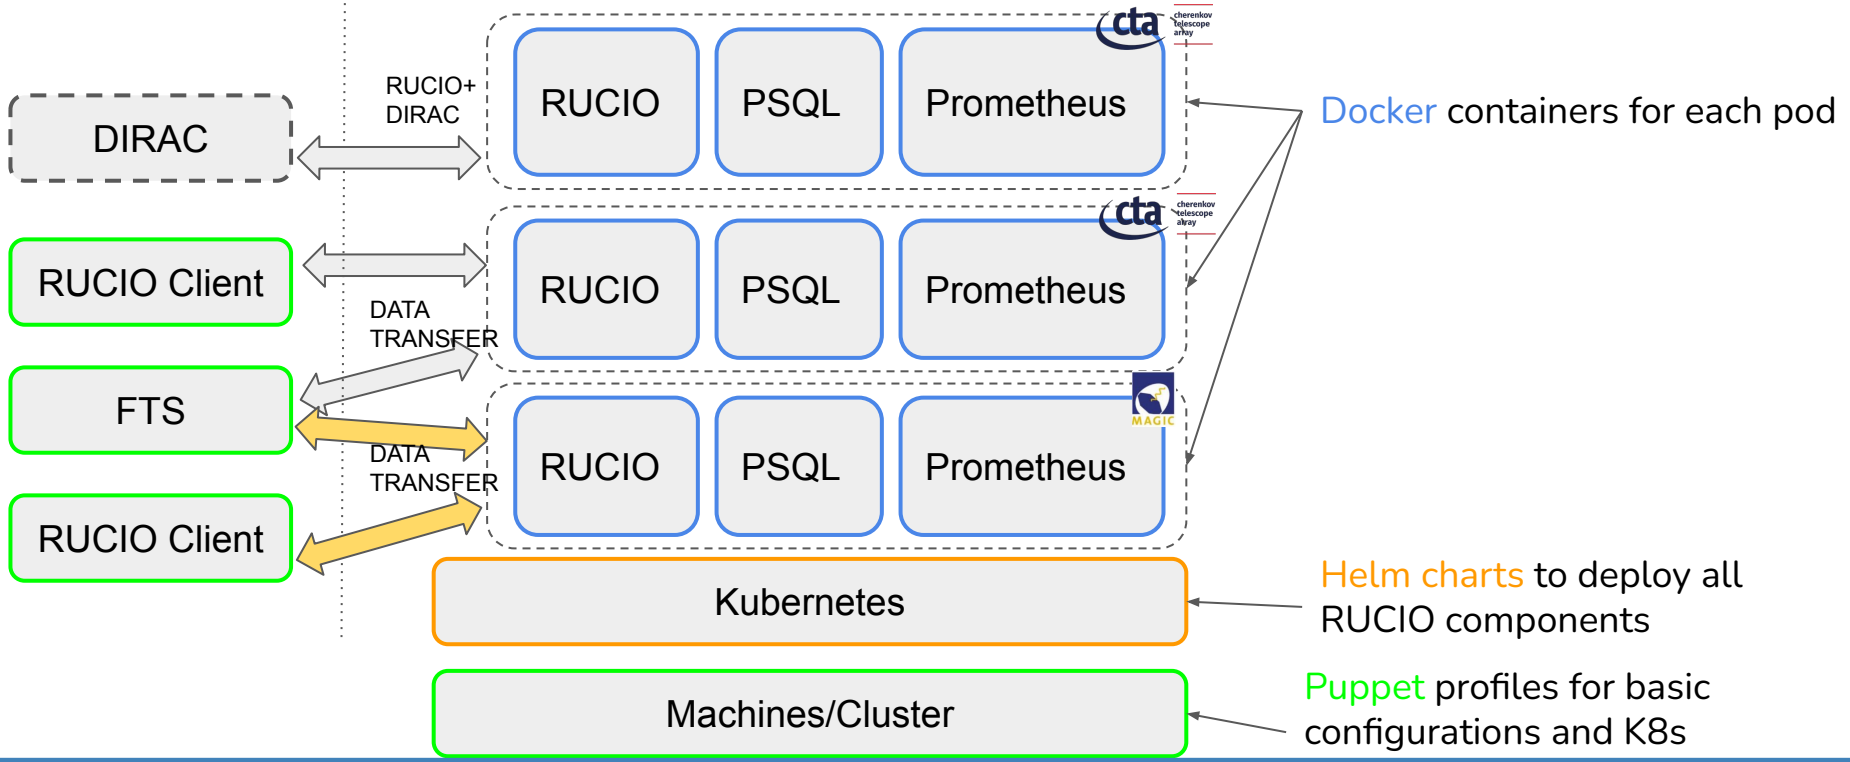

RUCIO DEPLOYMENT AT PIC

PIC is hosting three RUCIO instances, 1 for MAGIC, 2 for CTA (+ESCAPE) activities:

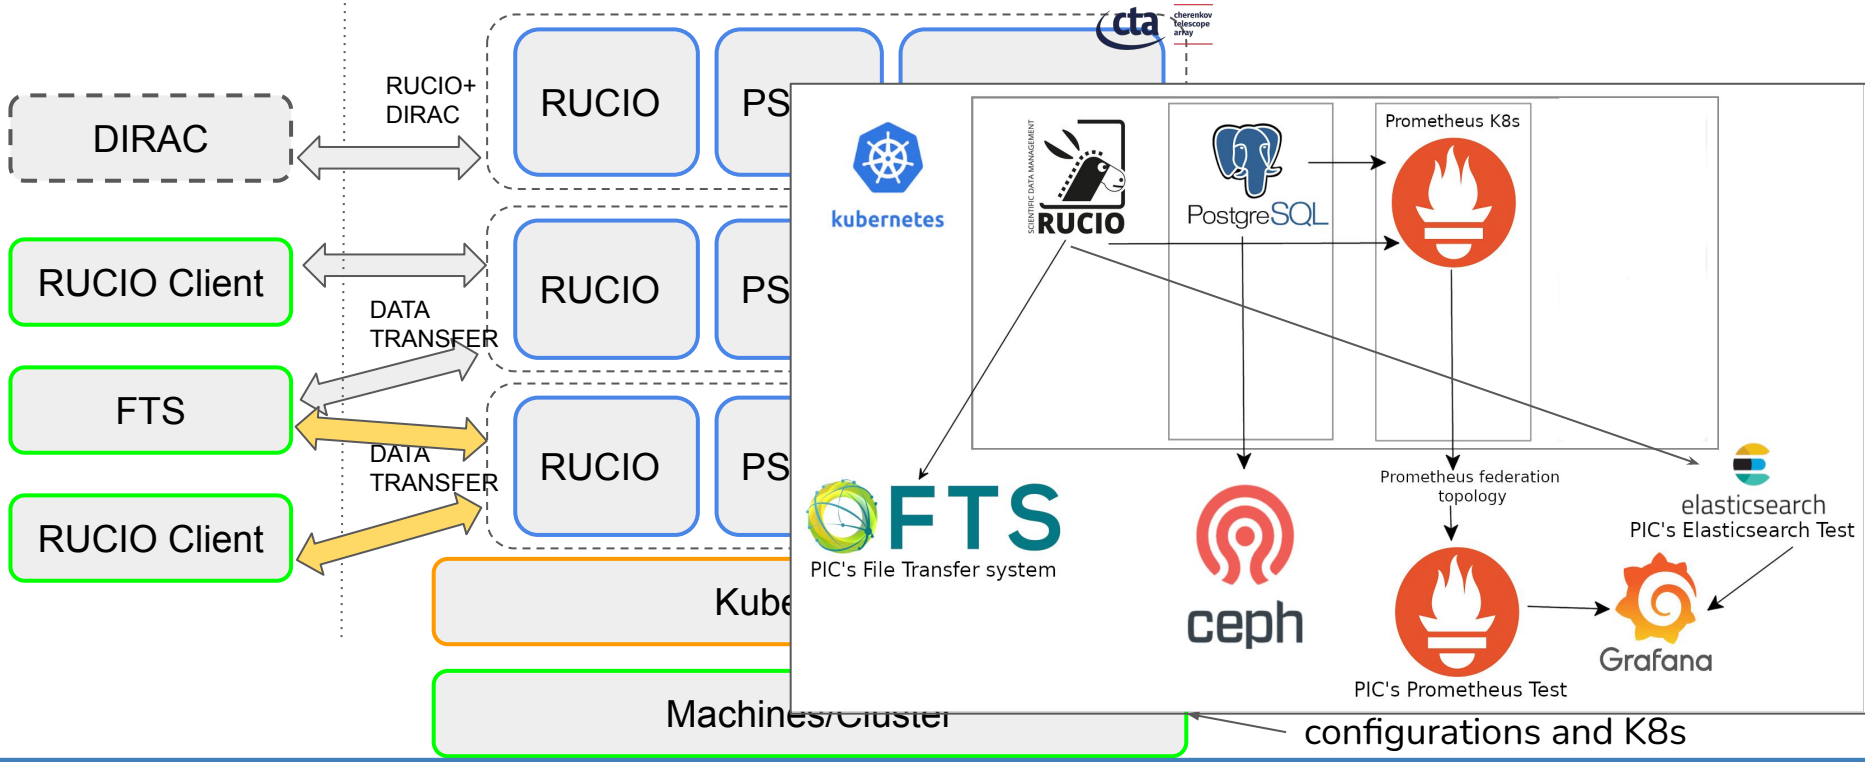

USE CASES RESULTS

Tests of the deployment and configuration for MAGIC and CTA configurations

- Different configurations for long haul transfers
  - GridFTP and XROOTD
  - Namespace
  - File size: from MB to TB
  - Number of files
  - Including 2 sites replication policies
  - Real test conditions (LP to PIC)
    - Dedicated 10Gbps connection
  - Disk and tape storage

- Basic operations RUCIO+DIRAC

- **Successful tests!**

MAGIC DATA TRANSFER MIGRATION

Old system: R.Firpo + J. Delgado (2012-2022)
scripts to orchestrate the transfer pipeline

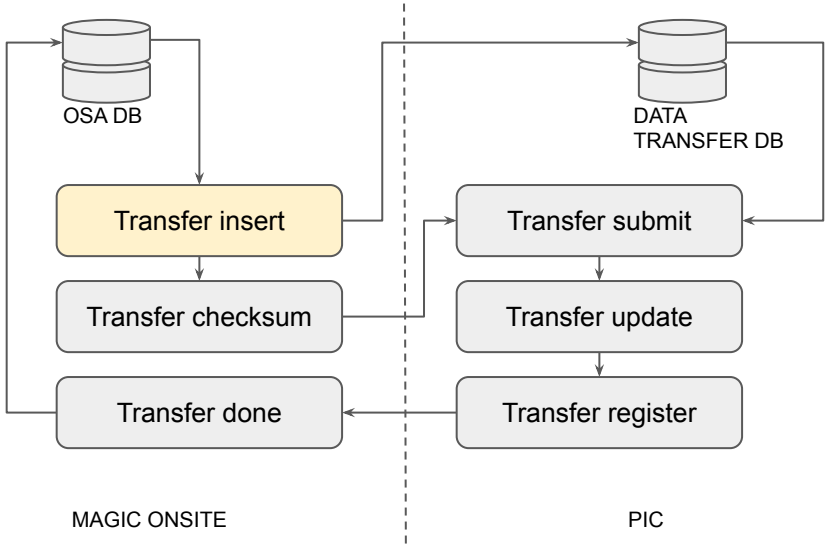

FTS File Transfer Service

New system: A.Bruzzese + J. Delgado (2022-?)
scripts to call RUCIO system

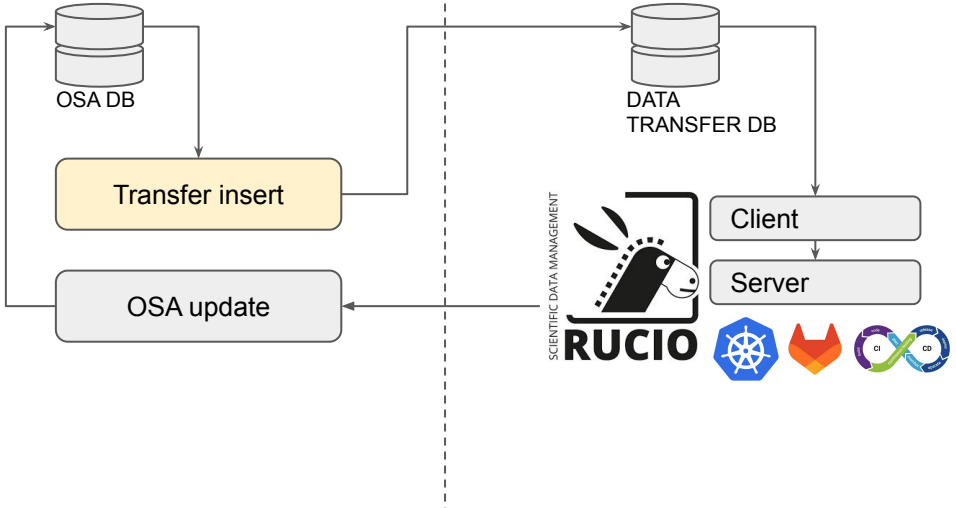

FTS File Transfer Service

MAGIC DATA TRANSFER MIGRATION

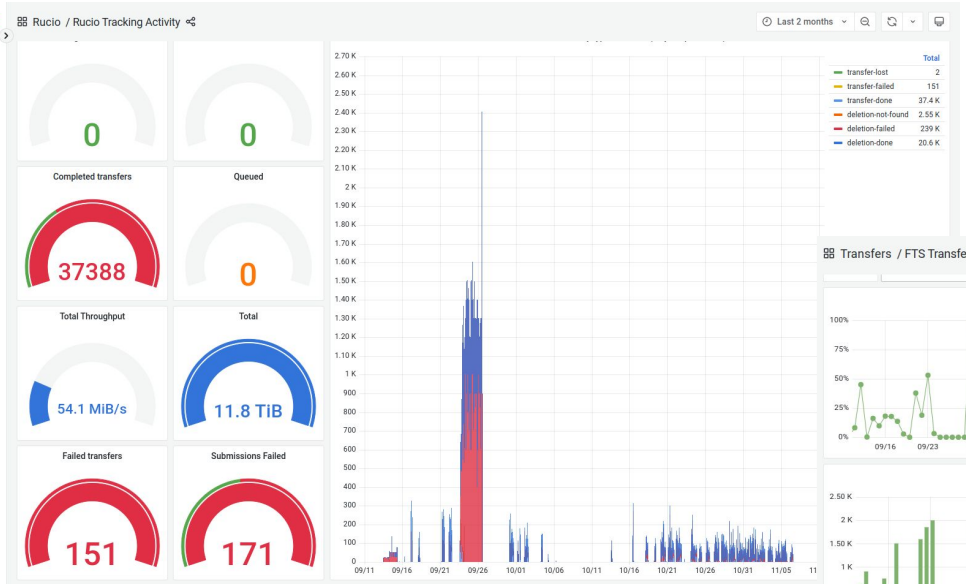

←PIC-RUCIO Dashboard

CERN FTS Dashboard →

SUMMARY

- CTA-ESCAPE activities and support from RUCIO community helped us to learn and succeed
- Willing to learn more and work with Multi-VO
- Post-ESCAPE escenario
  - More use cases to test: transcontinental transfers simulating CTA-South site
  - Integration with CTA on-site pipelines (more requirements will appear)
  - Automatic replications 1-N and N-N for CTA-LST
  - Conclude MAGIC development (still 1 module from the old implementation)
  - User AAI integration
  - Continue DIRAC-RUCIO integration
  - Improve Kubernetes deploys
- PIC is planning to scale RUCIO service to other experiments

RUCIO service for Gamma-ray astronomy projects

RUCIO Workshop, Lancaster University, November 2022

Jordi Delgado (jordidem@pic.es), A. Bruzzese, G. Merino on behalf of the PIC and CTA-ESCAPE teams

BACKUP SLIDE (I)

ON-SITE

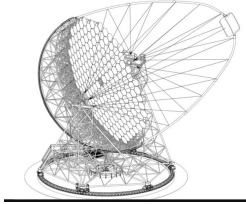

kubernetes

RUCIO

OFF-SITE: **CTA**'s Data lake

FTS

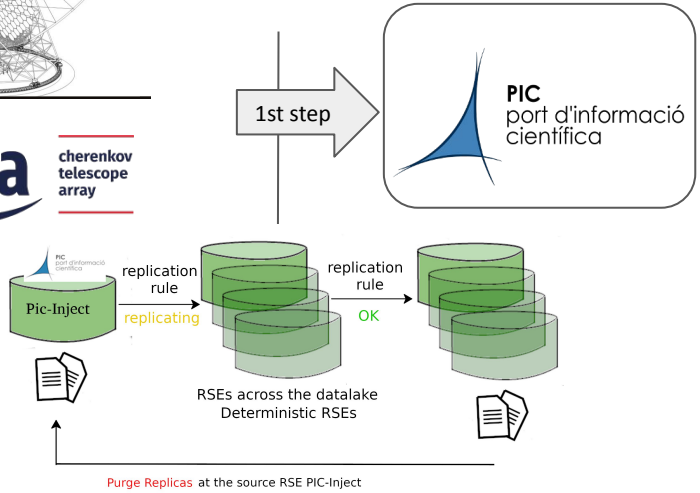

Test name	Protocol	Estimated #Files, data volume	STORAGE
Test 1	GSIFTP	300 GB	Disk
Test 2	XROOTD	300 GB	Disk
Test 3	GSIFTP	10 TB	Disk
Test 4	GSIFTP	40 TB	Disk
Test 5	GSIFTP	10 TB	Tape
Tes 6	GSIFTP	10 TB	Disk

Credits: A. Bruzzese

BACKUP SLIDE (I)

La Palma endpoint:
non deterministic

PIC endpoint :
deterministic

CNAF endpoint :
deterministic

Successful Transfers Percentage

Src	Dst	Bucket Script
CTAN-NON-DET	PIC-CTA-TAPE	98.3%
CTAN-NON-DET	PIC-CTA	99.4%
PIC-CTA	CNAF-DET-TEST	87.9%

Credits: A. Bruzzese