



ESCAPE

European Science Cluster of Astronomy &
Particle physics ESFRI research Infrastructures

Lessons learnt on satellite Rucio and FTS deployments

3rd ESCAPE DIOS Workshop - MAGIC Use Cases

R. Joshi (rohini.joshi@skao.int), A. Bruzzese (bruzzese@pic.es),
R. Bolton (Rosie.Bolton@skao.int), G. Merino (merino@pic.es)

CONTENTS

1. Introduction
2. ESCAPE Deployment
 - a. Architecture
3. PIC Deployment
 - a. Architecture
4. SKAO Deployment
 - a. Architecture
5. Deployment experience
 - a. What worked and lessons learned
 - b. Areas for improvement/areas of concern
6. Looking ahead



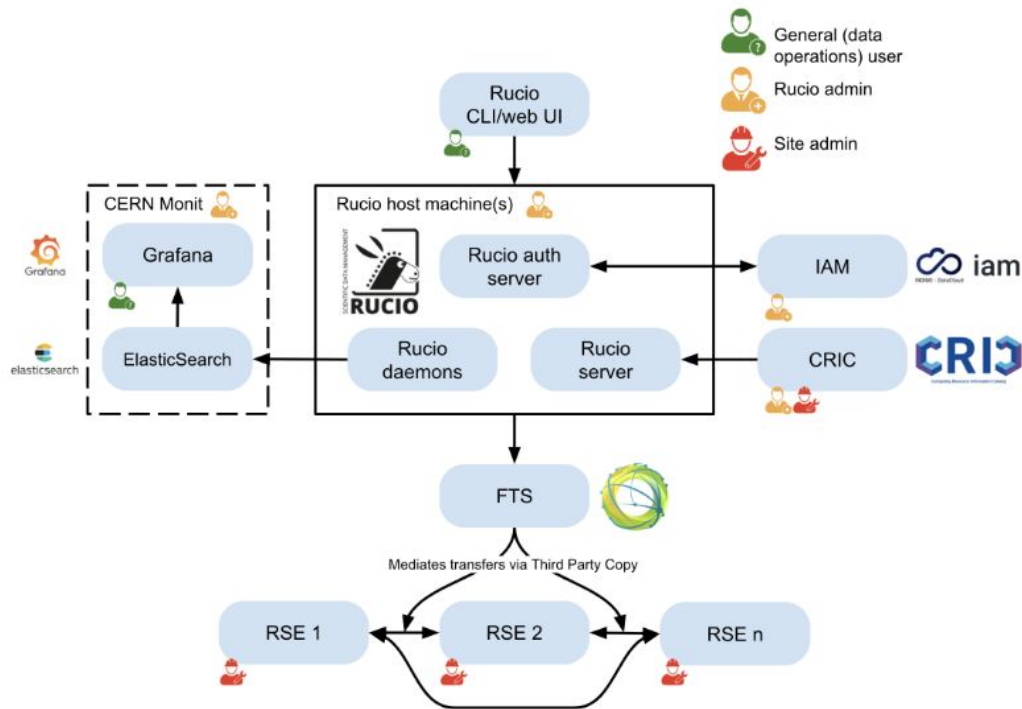
1. Introduction

- Since the beginning of the WP2 activities, the ESCAPE Rucio instance has served several ESFRIS as a first contact in the interaction and use of Rucio.
- On that basis, the experience gained in ESCAPE served to some ESFRI members, SKA and PIC/CTA, to adopt Rucio as the data orchestrator service for their respective experiments.
- This talk is to present the deploy architecture, current state, deploy, difficulties and future directors between SKA and PIC/CTA



ESCAPE: Rucio deployment architecture

- Data Lake endpoints
 - dCache, EOS, XRootD, StoRM, DPM
- Rucio as a data management tool
- FTS as service to coordinate file transfers
- CRIC in order to keep track of the different storage endpoint URLs and the protocols supported at the site



- FTS deployment at PIC
 - FTS is the GRID transfer service. It connects two endpoints and controls that the transfer takes place.
 - Traditionally, it was regularly used by Magic, CTA and for Rucio.
 - Mysql DB
- We used CERN's grafana in order to monitor it's activity
 - Link: <https://monit-grafana.cern.ch/login>

Instancia	Rol Servidor	Hostname
Producción	fts	fts01.pic.es
Producción	fts	fts02.pic.es
Test	fts-test	fts01-test.pic.es

Current State

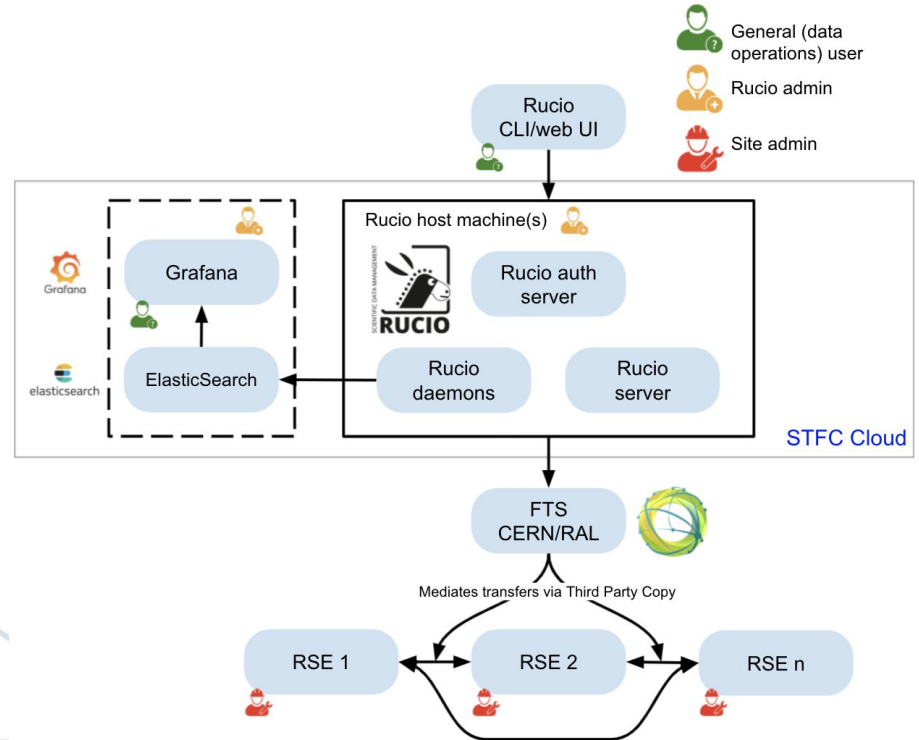
- **PIC/CTA**

- Results obtained during DAC21 have shown that the Rucio test instance on the PIC/CTA allows efficient orchestration of data transfer from the onsite MAGIC/CTAN to the offsite MAGIC/CTAN for long distances. Therefore, the next phase will be the migration to a production instance of both K8s and Rucio.
- Two Rucio instance:
 - A dedicated Rucio instance for Magic experiments. Currently we have cronjobs continuously running transfers.
 - PIC/CTA recently moved to Rucio as Data Management service integrated with DIRAC. For this integration, the Belleii DIRAC permission and schema was used, as a first approach.
 - Use of custom policy module necessitates separate Rucio instance



SKAO: Rucio deployment architecture

- Helm based deployment on a k8s cluster on STFC Cloud
- SSL traffic routed via a proxy machine
- Dev and Prod k8s clusters with one head node and 2 workers
- Rucio version 1.26, Postgres DB
- No web UI deployed
- RAL + CERN pilot FTS
- Grafana monitoring, regular functional tests submitting transfers.
- <https://gitlab.com/ska-telescope/sr/ska-rucio-prototype>
- Token integration currently being trialed on the Dev instance



Deploying Rucio: what worked and lessons learned

- Community aspect of the project was very helpful in getting started
- New docs are much better (<http://rucio.cern.ch/documentation/>)
- Helm version vs rucio version inconsistencies - resolved now which is appreciated
- SKAO: We used gdocs to log everything as we went, generating 100+ pages



Deploying Rucio: Current Limitations, Areas for improvement

- General
 - Upgrading major versions is not a particularly user friendly experience; database upgrade with alembic is a very manual process (documented [here](#))
- Token-based Rucio
 - Demonstrated use only for a specific, assumed architecture, didn't work out the box for us (e.g. routing through proxy, [issue #5220](#) and a smaller [issue #5219](#))
 - Not many communities currently using tokens, so software rot could be a concern
 - Assumed knowledge in documentation is large - but we understand that tokens is an area in flux
- DIRAC integration
 - Rucio has some constraints/features that are not available in DIRAC that needs to be taken into account: for eg, unique scopes
 - Interaction with data must be via DIRAC
 - More here ([Link](#))



Looking ahead

- Run token-based Rucio functional tests
 - Upgrade to Rucio v1.28
- RSE deployment:
 - Deploy and document how to setup a lightweight RSE (StoRM)
 - Ceph/S3 RSE integration (following [this](#))
- Continued Rucio-DIRAC integration
 - Explore suitability of MultiVO deployment



Looking ahead

How can we be best supported to do what is next on our plate?

- What is the timeline for ATLAS to switch to an OIDC compliant Rucio stack?
- We need refresh token functionality in order to get token-based Rucio functional tests running - is this correct?
- What spaces and forums would be available after end of ESCAPE to continue collaborating on satellite Rucio/FTS deployments?
- Can a multi-VO Rucio deployment have one VO integrated with DIRAC and one operating without?



Acknowledgements

Thanks to the people involved in each of the experiments.

SKA team:

James Collinson (SKAO),
Rob Barnsley (SKAO),
Rosie Bolton (SKAO)

PIC/CTA team:

Frederic Gillardo (LAPP),
Luisa Arabito (IN2P3),
Gareth Hughes (CTAO),
Nadine Neyroud (LAPP),
Jordi Delgado (PIC),
Gonzalo Merino (PIC)

And also other people from ESCAPE partners

Alba Vendrell
Rizart Dona
Riccardo di Maria
Xavier Espinal



The logo features a white circle containing a blue orbital path with a starburst at one end and a yellow dot at the other. The word "ESCAPE" is written in large, bold, black capital letters across the center of the circle.

ESCAPE

European Science Cluster of Astronomy &
Particle physics ESFRI research Infrastructures

Thanks for listening!
Questions?