

RUCIO/experiments mini-workshop

Tuesday 4 May 2021, 14:00 → 17:00 Europe/Paris

Chairs: Martin and Alba, Secretary : Xavi

Present: Rosie, Arturo, Marek, Nadine, Rosie, Martin, Alba, Xavi, Riccardo, Marcelo, Florian, Benedikt Ziemons, Aleem Sarwar, Agustin, Aleem, Fabio, Florian Uhlig, Gareth, Gonzalo, Guido, James Collinson, Jutta, Maisam, Matthias Fuessling, Nazaret Bello, Pandey, Peter Caligari, Radu Carpa, Rizart, Rob Barnsley, Rohini, Thomas Beermann, Taras Yakobchuk. Marco Guenter, Lucia, Frederic Gillardo, Federica Agostini, Berkay Turk, Tamas Gal, Cedric Serfon, Raymond Oonk, Ghita Rahal, Mario

Meeting notes (you're all welcome to contribute)

Introduction:

- Alba introduced the workshop, emphasising the discussion oriented style.
- Martin presented shortly RUCIO project organisation, communication, release frequency, community and documentation, etc. [\[link\]](#)
 - 3 to 4 releases per year
 - Rucio is a community project; objective is to foster long-term community involvement
 - 30 different contributors in 2020. 19k out of 35k lines of code coming outside from ATLAS.

Discussion Forum:

- ***Envisaged use of metadata in RUCIO?***
 - CTA attach metadata to files and dataset. Special interest in
 - Timestamps (to ease time range query for datasets and files). Query using time ranges works with some operators (=) but not (<, >)
 - Coordinates (x,y) sky pointing positions of the telescopes and searchable over operators (<,>)
 - Two kinds of data suitable to be browsed through metadata: a) Archive (easiest queries done by experiment experts) and b) High level data archive for scientists, more open to any kind of queries.
 - 200k file test looking good for '=' operator, more high performance tests to be done.
 - SKA metadata driven subscription: metadata in a file directs the association into a dataset, then the flags associated with these types of data are honored.
 - LOFAR granular access (data access granted to few people)
 - Containers might be used to "silo" the data but Access Control is probably better addressed in the ACLs/Embargo strategy
 - External metadata interface to RUCIO in place, eg. DUNE.

- Question whether this can be used to place subscriptions, eg. metadata triggered replication or rule honoring. Subscriptions are evaluated at the “data injection”, but this can be addressed a posteriori through queries and then trigger actions.

- **ACLs and Embargo data needs**

- Delicate disconnection by design: sites storage permission not coupled to RUCIO, makes difficult to delegate ACL enforcing in RUCIO.
- The model proposed in WP2 (namespace based ACLs) can be a good strategy
- OAuth/OIDC token will make this simpler and open new possibilities, this is ongoing work
 - RUCIO scope based authorization, the token will come with a pass, and more fine grained auth will be possible based on scopes.
- Suggestion to start slowly and see the token based auth as a global orchestration “service”, need probably to address and cover the current cases: power users, plain users, no access in all levels of the data orchestration zoo (RUCIO, FTS and storages)
- Is some embargo data mechanism working right now? proof of concept done in the context of CMS data, a demonstrator with quite “manual” need in the sense it needs the storage system to set permissions on a path/namespace level.

- **Data organisation, File volumes, Container sizes**, etc. no issues/comments.

- **RUCIO instances besides the CERN/ESCAPE, deployment plans, issues, questions**

- **CTA** dedicated RUCIO instance at PIC, running with no main issues. Monitoring tools are scarce. Deployed with k8s. Tried also with Docker but felt easier with k8s. Documentation ok
- **Vera C. Rubin/LSST** deployed a RUCIO instance at IN2P3 for educational purposes. Joint FTS and RUCIO package deployed. Some conflicts with the python version needed by RUCIO vs. LSST software.
 - Suggest not to constrain people to use k8s, easier with Docker or single executable. This will lower the barrier and foster easiness to try.
- **SKA** deployed a RUCIO instance. Used k8s, good knowledge in SKA already on k8s. Understanding/creating the secrets somehow felt tricky.

- **Data Lifecycles**

- Current support in RUCIO covers quite a lot of range on that.
- SKA
 - Data lifecycle needed on future time triggers. This is implemented or imminent (need to be checked)
 - What is the data popularity current level of implementation status in RUCIO?
 - Popularity component seldomly used, re-thinking this module at the moment.
 - Machinery in place., eg. daemon store traces each time a file is accessed. This info might be used/interfaced with WMS. Enforcing rules/actions based on popularity possible but not

planned at the moment (e.g. QoS transitions: disk to tape offload in case of no access after some time)

- RUCIO as a community effort
 - Consider joining the community, some of the dev/ops meetings and workshops
 - Opportunity for students in Computing Science or related fields, opportunities to set up projects in the data management framework.
 - Are there specific areas where RUCIO need the contribution of the community;
 - Many. To name some: metadata extension, token based auth, QoS implementation
 - And some simple stuff, 200 open tickets in Github, small contributions appreciated as well
- Misc
 - Long haul network transfers, have other different transfer tools been considered, tested, compared e.g. globus.
 - FTS based in WLCG but some tests done and used in prod, eg. HPCs in the US want(need) globus. ATLAS and CMS communities are the one experienced.
 - From the network perspective gridftp is widely used and has been widely tested. Nevertheless the movement to webdav (mandatory as globus is moving to \$\$\$\$) is still to be tested at the same scale.
 - Evaluation of the Token integration effort and maintain the legacy x509
 - For RUCIO this means running both components, storage is a different story and not under RUCIO control
- RUCIO desktop application
 - Summer student work, the feeling is that it should be working. To be checked.
 - <https://github.com/rucio/documentation>
 - Future plans about this foreseen to be driven by user needs.

Identified needs/issues:

- Support for operators (e.g. =,<,>) on metadata queries, main reasons: time ranges, sky spatial coordinates, etc.
- SKA metadata driven subscription drives the association of files into a dataset
- ACLs is a hot topic, tokens will ease and open the floor for more fine grained access control. Work in progress in ESCAPE and WLCG/TPC.
- RUCIO deployment, understanding of the secrets scope and needs are difficult to figure out.
- Suggested to have a collection of links to all documentation and wikis, it is found hard for beginners to find all the information out there that goes beyond the official readthedocs (<https://github.com/rucio/documentation>)

Follow-up

- Multi-RUCIO WP2 initiative will be tracking down the deployment of several RUCIO instances in some of the ESCAPE ESRIs
- AAI/token/ACLS WP2 initiative will be tracking the progress towards the embargo data needs and the full token support in the datalake orchestration habitat: FTS, RUCIO, storage providers, etc.