

Building a multi-VO Rucio instance at RAL

Ian Johnson, Alastair Dewhurst

Project Scope

- Rucio is a Distributed Data Management service
 - Used by ATLAS, CMS, exploration from SKA, DUNE and LIGO.
- We believe that Rucio is a sensible choice to manage user data between the various IRIS storage resources.
 - Currently Rucio is limited to a single Experiment per instance.
- The project is to build a production-quality, multi- experiment Rucio instance at RAL.
 - The Tier-1 has Rucio expertise from its Tier-1 Liaisons.
 - The Tier-1 had already set up a Rucio test instance for SKA to evaluate.
 - **Attract.eu project to test network connectivity between South Africa and Europe using this instance.**



Current Status

- Ian Johnson has spent 3PM on this work.
 - 1.1 Deploy a resilient database backend.
 - MI.1 Database migration failed. We had to re-populate the database manually.
 - DI.1 Resilient PostgreSQL database deployed.
 - 1.2 Deploy a Rucio UI and enable monitoring.
 - DI.2 UI has been setup.
 - MI.2 Monitoring is waiting on stable Elastic Search cluster.
 - 1.3 Improve resilience and scalability of service.
 - MI.3 Migrated service to OpenStack instance and documented deployment steps (easier in future).
 - DI.3 Provisioned ActiveMQ machine, but no additional Rucio servers yet.



Future work

- Add Multi-VO support to Rucio, to support multiple experiments in the same instance:
 - 2.1 Schema adaptation - Add VO field to all database tables.
 - 2.2 Authentication and Base client - Add VO specific authentication
 - 2.3 Functionality - Add VO field to all Rucio functions.
- Will recruit 6 month graduate from either STFC (Scientific Computing Dept.) or Centre for Doctoral Training.
 - Ian will supervise during this time.
- Expected completion September 2019.
 - We would want to start engaging with potentially interested experiments then.
 - We would expect to request further funding for digital assets to continue to expand the service.



Backup

